

# LeOPaRd: Towards Practical Post-Quantum Oblivious PRFs via 2HashDH Paradigm

Muhammed F. Esgin<sup>1</sup>, Ron Steinfeld<sup>1</sup>, Erkan Tairi<sup>2\*</sup>, and Jie Xu<sup>1</sup>

<sup>1</sup> Faculty of Information Technology, Monash University, Australia

<sup>2</sup> University of California, Berkeley

name.surname@monash.edu, erkan.tairi@berkeley.edu

**Abstract.** In this work, we introduce a more efficient post-quantum oblivious PRF (OPRF) design, called LeOPaRd. Our proposal is round-optimal and supports verifiability and partial obliviousness, all of which are important for practical applications. The main technical novelty of our work is a new method for computing samples of MLWE (Module Learning With Errors) in a two-party setting. To do this, we introduce a new family of (interactive) lattice problems, called *MLWE-PRF with re-use* (MLWE-PRF-RU). Here, the adversary is given a mix of MLWE and PRF samples where each PRF error is dependent on an *adversarially-chosen* matrix and the MLWE error. We rigorously study the hardness of MLWE-PRF-RU and provide a reduction from the standard MLWE to MLWE-PRF-RU, establishing a strong security foundation. We believe MLWE-PRF-RU problem family and the intermediate security reductions we introduce along the way can be of independent interest for other interactive protocols. LeOPaRd exploits this MLWE-PRF-RU assumption to realize a lattice-based OPRF design without relying on heavy machinery such as noise flooding and fully homomorphic encryption used in earlier works. LeOPaRd can feature around 137 KB total communication, compared to 300+ KB in earlier works, while also achieving about  $2\times$  reduction in *online* communication compared to the prior state-of-the-art result. We also identify gaps in some existing constructions and models, and propose appropriate fixes.

**Keywords:** Oblivious PRF · Post-Quantum · Lattice · MLWE with Hints · 2HashDH

## 1 Introduction

An oblivious pseudorandom function (OPRF) extends a standard PRF in the following manner. In a standard PRF, a party in possession of a secret key  $k$ , can evaluate the PRF  $F$ , on any given input  $x$  to obtain the output  $y = F(k, x)$ . The pseudorandomness of a PRF ensures infeasibility of distinguishing the outputs of this function from those of a random function. Now, in OPRF, the goal is to compute the PRF output  $y$  in a two-party server-client model, where the server holds the secret key  $k$  and the client has the input  $x$ . After they run the two-party OPRF protocol, we want the client to learn the PRF output  $y$  without learning anything about the secret key  $k$ , and the server does not learn any information on the input  $x$  or the output  $y$ . We say that an OPRF is *verifiable* OPRF (VOPRF) if the client is guaranteed that the output received is indeed evaluated under the committed key. As discussed in [CHL22, TCR<sup>+</sup>22], many applications of OPRFs such as for password-authenticated key exchange, checking compromised credentials and spam detection require a part of the client’s input to be public to allow for domain separation for the PRF computation. In this case, we divide the client’s input into a private part,  $x$ , and a public part,  $t$ , called the tag, and we call the resulting primitive a *partial* OPRF (POPRF).

(VP)OPRFs have become an important tool for privacy-preserving protocols and have numerous applications, including but not limited to private lightweight authentication mechanisms [DGS<sup>+</sup>18], private set intersection for checking compromised credentials [LPA<sup>+</sup>19, TPY<sup>+</sup>19], secure data de-duplication [KBR13], password-protected secret sharing [JKK14, JKX16] and password-authenticated key exchange (PAKE) [JKX18].

---

\* Work done while the author was at ENS Paris.

Scheme	Assumption	Rounds	Communication	Model
[ADDS21]	RLWE + SIS	2	>128GB	strong, QROM
[ADDG24]	heuristic	2	15MB + 633KB	strong, ROM
[BDFH25]	Legendre PRF	9	911KB	strong, ROM
[YBH <sup>+</sup> 25]	Power Residue PRF	$\geq 5$	$\geq 1,600$ KB	strong, ROM
[AG24], $Q = 2^{16}$	RLWE	2	108KB + 189KB	weak, ROM
[AG24], $Q = 2^{32}$	RLWE	2	114KB + 198KB	weak, ROM
LeOPaRd, $Q_{x,t} = 2^{16}$	MLWE-PRF-RU	2	(10 + 30)KB + 97KB	strong, ROM
LeOPaRd, $Q_{x,t} = 2^{32}$ (Section 4)	+ MLWE + MSIS	2	(21 + 58)KB + 103KB	strong, ROM

Table 1: Comparison of post-quantum (publicly) verifiable (P)OPRF constructions in the full malicious setting (without trusted setup). The summands in communication (when provided) indicate an offline cost and an online cost per query, respectively. For LeOPaRd, the offline cost is further broken down into the one-time (first term) and per-query (second term) costs while the offline cost is one-time for others (when provided).  $Q$  denotes a *global* upperbound on the number of all OPRF queries; while  $Q_{x,t}$  denotes the maximum number of OPRF queries *per fixed* PRF input/tag pair  $(x, t)$ . Enforcing the bound  $Q_{x,t}$  is easy by counting the number of queries under the same *public* tag  $t$ .

Despite the widespread use of (VP)OPRFs, the existing constructions are either efficient but based on classical assumptions [FIPR05, JL09, JKK14, TCR<sup>+</sup>22], or based on (plausibly) post-quantum assumptions [ADDS21, ADDG24, BDFH25, AG24] but practically inefficient. Given the recent standardization efforts by the IETF<sup>3</sup> for the DH-based OPRFs of [JKK14, TCR<sup>+</sup>22] and by NIST for other post-quantum primitives<sup>4</sup>, it is imperative to find efficient post-quantum OPRF constructions, which can pave the way for future post-quantum OPRF standards. We summarize the state of the art in post-quantum verifiable (P)OPRFs in Table 1 (extending a table from [AG24]). All schemes here provide security against malicious clients and servers (as needed in many applications). We also distinguish between the security models used in the schemes as “strong” if a stronger model as in [TCR<sup>+</sup>22] or UC framework is used<sup>5</sup>, and as “weak” for otherwise weaker models.

## 1.1 Our Contributions

**Novel post-quantum VPOPRF construction.** Our first contribution in this work is a construction of a lattice-based VPOPRF, called LeOPaRd, that is more efficient than the current state-of-the-art. The comparison of our construction with the existing post-quantum VPOPRFs is given in Table 1, where we present our parameters for  $\approx 95$ -bit security to match the security level for prior lattice-based schemes [AG24]. A more thorough performance analysis of our scheme is provided in Section 5, along with parameters for 128-bit security in Table 3. The efficiency gains of our construction come not only from utilizing the state-of-the-art lattice-based NIZK proof systems, such as LaBRADOR [BS23], but also by using a novel assumption that we introduce, and cleverly combining this with some additional lattice-based techniques that we explain in Section 1.3. The vast majority ( $\sim 90\%$ ) of the online communication cost for LeOPaRd in Table 1 arises from the underlying well-formedness NIZKs (performed by client and server). Excluding the NIZKs, LeOPaRd’s core online OPRF communication cost is only around 0.5 KB for the server (vs 1.8 KB in [AG24]) and 7-13 KB for the client (vs  $\sim 73$  KB in [AG24]); see Tables 3 and 4 for more details.

We prove our construction secure in the random oracle model in Section 4.1. Specifically, we show that our construction achieves the strong security definitions from [TCR<sup>+</sup>22]: pseudorandomness, which provides

<sup>3</sup> <https://datatracker.ietf.org/doc/rfc9497/>

<sup>4</sup> <https://csrc.nist.gov/projects/post-quantum-cryptography>, <https://csrc.nist.gov/projects/pqc-dig-sig>

<sup>5</sup> We note that the UC model of OPRF is even stronger than the one given in [TCR<sup>+</sup>22].

security in the presence of malicious clients (POPRF security), and request privacy against malicious servers (POPRIV2 security). Moreover, in Appendix E.2 we show that our construction also achieves uniqueness, which in the verifiable setting ensures to the clients that the server honestly and consistently performs the blind evaluations. Furthermore, as described in Remark 4, our construction easily extends to  $n$ -out-of- $n$  threshold setting, as in [AG24], thanks to its key-homomorphic property.

**New MLWE-like assumption.** We introduce a family of (interactive) assumptions, called *Module Learning with Errors & PRF with Re-Use* (MLWE-PRF-RU), and use it to prove the security of our construction (against malicious clients). We believe introduction of such an interactive MLWE-like assumption is a natural next step given the interactive nature of protocols we deal with. The new assumption allows us not only to prove the security of our OPRF construction, but also to circumvent the shortcomings of the prior lattice-based OPRF constructions [ADDS21, ADDG24], such as removing the need for noise flooding or fully homomorphic encryption. We also note that our assumption is not “one-more” type (unlike earlier proposals, e.g. [JKK14, TCR+22, BDFH25]), and supports *adaptive* queries.

The MLWE-PRF-RU problem can be seen as a family of assumptions, which roughly state that a mix of MLWE and PRF samples are indistinguishable from uniform, even when the PRF samples have “deficient” errors that depend on *adversarially-chosen* (small-norm) matrices and MLWE errors. Crucially, to circumvent “trapdoor” attacks, the small matrix coefficients must be committed by the attacker *before* seeing the MLWE and PRF samples. This prevents the attacker from meaningfully biasing the “deficient” PRF samples. To establish strong confidence in the security of MLWE-PRF-RU, in Section 3, we provide a series of reductions that eventually prove that MLWE-PRF-RU is as hard as the standard MLWE problem (under appropriately chosen parameters). We note that the MLWE-PRF-RU assumption and the intermediate assumptions we introduce along the way might be of independent interest and find use in proving the security of other types of interactive protocols, such as blind signatures, as an alternative to one-more type assumptions.

**Issues with existing constructions and models.** As an additional contribution, we uncover flaws in the security proofs of two existing lattice-based OPRF constructions, [ADDS21] and [ADDG24]. We refer to Appendix E.1 for a more detailed exposition of these flaws and their potential fixes. We also observe that a reduction made in [TCR+22], that correctness and POPRIV2 security together imply uniqueness, does not necessary hold unless we have a one-to-one correspondence between the OPRF secret and public keys. In Appendix E.2, we patch this reduction by introducing a new property, called *key binding*, and show that correctness, key binding and POPRIV2 are sufficient for uniqueness. Along the way we also prove that LeOPaRd achieves key binding.

## 1.2 Related Work

The first post-quantum, and also lattice-based, VOPRF construction was given in [ADDS21], which although relied on a standard RLWE and SIS assumptions, was totally impractical as it required communication costs in the order of GBs. This was improved in [ADDG24] by combining fully homomorphic encryption with the weak PRF given in [BIP+18], but the reported bandwidth per query was in the order of MBs. Moreover, in order to achieve verifiability they needed to rely on a heuristic assumption about the hardness of evaluating deep circuits in an FHE scheme supporting only shallow circuits, which was recently broken [CJ25]. Recently, Beullens et al. [BDFH25] provided a VOPRF construction that is based on the Legendre PRF and which makes use of oblivious transfer (OT) and ZK proofs that can be instantiated from lattice assumptions, resulting also in a plausibly post-quantum secure VOPRF. However, this construction also has communication costs in the order of MBs.

**Comparison with concurrent works** [AG24], [YBH+25], [HKL+25] and [DDT25]. In a concurrent and independent work published on IACR’s ePrint just a few days before we uploaded the first version of this paper to ePrint, Albrecht and Gur [AG24] provided a new lattice-based ( $n$ -out-of- $n$ ) VOPRF construction. Their construction significantly improves, both assumption and performance wise, the earlier lattice-based VOPRF given in [ADDG24] as shown in Table 1. Similar to ours, their starting point is the OPRF construction given in [ADDS21], however, unlike us, they take a different technical route. They improve the construction given in [ADDS21] by first removing the reliance on the 1D-SIS assumption by borrowing a

trick from [GdKQ<sup>+</sup>24] and making use of Rényi divergence to avoid noise flooding, which together remove a superpolynomial factor from their modulus  $q$ . However, switching to Rényi divergence forbids them from using an indistinguishability-based security model, and in turn forces them to use a *weaker* security model than the one we use, which is the well-established OPRF security model given by Tyagi et al. [TCR<sup>+</sup>22]. Specifically, the model given in [TCR<sup>+</sup>22] is a simulation-based indistinguishability model that provides a careful and granular tracking of the oracle calls and comes close to the UC security model for blind protocols given in [JKK14]. On the other hand, the security model used by Albrecht and Gur [AG24] is akin to the one-more unforgeability definition used in the context of blind signatures, and is comparatively weaker than the model of Tyagi et al. [TCR<sup>+</sup>22]. Moreover, the use of Rényi divergence forces them to use a *global* bound on number of all OPRF queries, which can potentially hinder the practicality. Our proposal, on the other hand, relies on a bound  $Q_{x,t}$  on the number of queries *per input/tag pair*  $(x, t)$ . Given the tag  $t$  is *public*, it is straightforward for the server to count the number of OPRF queries made per tag  $t$ . For example, if the tag is a fixed username (unique per client) as in PAKE protocols [JKX18], we can set  $Q_{x,t} = 2^{24}$  *per client* and support a practically unlimited number of clients, say,  $2^{64}$ ; whereas [AG24] would need to set  $Q = 2^{88}$  to support the same setting, significantly degrading efficiency. Note that exhausting  $Q_{x,t} = 2^{24}$  queries would require a client to continuously make queries *every minute* for about *32 years* (1 year  $\approx 2^{19}$  minutes).

Apart from the aforementioned improvements over the prior work [ADDS21], Albrecht and Gur [AG24] also make use of the newer lattice-based proof systems, specifically LNP22 [LNP22] and LaBRADOR [BS23], analogous to us. However, in order for their security proof to go through, they require the underlying proof system to be straight-line extractable. This can be done for LNP22 using the Katsumata transform [Kat21] by incurring a significant (4-8 $\times$ ) performance penalty [NO25]. However, it is an open problem how to obtain straight-line extractability for LaBRADOR. Importantly, we do not require straight-line extractability from the underlying NIZK proof systems for our security proofs. Lastly, the additional overhead that comes from the straight-line extractability, along with the fact that they aim only for 90-100 bit security level, are not properly reflected in their reported numbers, which are also shown in Table 1. Given that they [AG24] rely on (fully structured) Ring LWE problem, they have less flexibility in choosing their parameters, and would need to incur almost  $2\times$  performance penalty if their parameters aimed at 128-bit or higher security level. To match the security level of [AG24], our reported results in Table 1 also aim at 90-95 bits of security, and we provide parameters and performance results at 128-bit security level in Table 3. Overall, accounting for (i) a stronger security model, (ii) straight-line extractability for NIZKs, and (iii) at least 128-bit security level would lead to a significant cost increase in [AG24]. Our proposal LeOPaRd and its performance analysis readily account for all these features and even partial obliviousness.

Another concurrent work, Gold OPRF [YBH<sup>+</sup>25], takes a completely different technical approach, using a generalization of Legendre PRF, namely Power Residue PRFs. In the full malicious setting (without trusted setup), their communication cost (without amortization) is significantly higher than ours ( $> 10\times$ ) while also requiring at least 5 rounds of communication (compared to our round-optimal proposal). On the other hand, Gold OPRF can support efficient batching on both the server and client sides, but the amortization starts to make a significant effect only after batch sizes above thousands, which may not be realistic in real-life applications. Moreover, Gold OPRF primarily targets the pre-processing model, which does not fit some important applications of OPRFs exhibiting a password-only setting such as password-authenticated key exchange (PAKE) protocols. Our LeOPaRd proposal, however, can fit such password-only settings.

Another concurrent work, LEAP, in [HKL<sup>+</sup>25] is considered only in the semi-honest model for both the client and server (therefore, also is not included in Table 1). Even in this model, it requires 6 rounds of communication, 23 KB online communication, and almost 800 KB of preprocessing communication. Moreover, it is based on a heuristic security assumption, and we are aware that LEAP is actually insecure [Hei25].

In a very recent work, Davidson et al. [DDT25] extended LEAP [HKL<sup>+</sup>25] by removing its heuristic assumption, and achieving round-optimal online phase by utilizing preprocessing. However, they also only achieve semi-honest security (and hence, is not included in Table 1). Their online communication is 11.9 KB, but the preprocessing requires between 790 KB and 1.9 MB communication depending on parameter choices (as highlighted in [DDT25, Table 7]).

### 1.3 Technical Overview

**VPOPRF construction.** The 2HashDH OPRF paradigm [JKK14] (and its extension to partially oblivious setting in [TCR+22]) has been shown to be highly effective in the discrete-log setting. Therefore, our goal is to efficiently realize this high-level paradigm in the lattice setting. An initial attempt was already done in [ADDS21], which serves as a starting point for our scheme. For simplicity, our discussion here focuses on the regular OPRF setting (without partial obliviousness).

The high-level idea of the 2HashDH paradigm is to first hash the PRF input  $x$  to some group element  $b_x = H_1(x)$ . This term is then blinded via a randomness  $r$  and the resulting value  $c_x = H_1(x)^r$  is sent to the server. The server then blindly performs PRF evaluation using its key  $k$  as  $u_x = c_x^k$ . The blinded evaluation result  $u_x$  is sent to the client, who unblinds the value using their randomness  $r$  to obtain a PRF value  $z = H_1(x)^k$ . The final PRF output is computed via a second hashing as  $y = H_2(x, z)$ . The blind/unblind operations rely on the homomorphic properties of the computations performed in the protocol.

To realize this idea in the lattice setting over a ring  $R_q$  with modulus  $q$ , we want to encrypt the matrix  $\mathbf{B}_x = H(x)$  using homomorphic encryption. We can do this as follows.

1. Client computes  $\mathbf{C}_x = \mathbf{R}\mathbf{A}_r + \mathbf{B}_x$  for randomness  $\mathbf{R}$  and public matrix  $\mathbf{A}_r$ .
2. Upon receiving  $\mathbf{C}_x$ , the server can compute  $\mathbf{u}_x = \mathbf{C}_x\mathbf{k} + \mathbf{e}'_s$  using its key  $\mathbf{k}$  with some error vector  $\mathbf{e}'_s$ . Observe that  $\mathbf{u}_x = \mathbf{R}\mathbf{A}_r\mathbf{k} + \mathbf{B}_x\mathbf{k} + \mathbf{e}'_s$ .
3. Given  $\mathbf{u}_x$  and additionally  $\mathbf{v}_k = \mathbf{A}_r\mathbf{k} + \mathbf{e}_s$ , the client can compute  $\mathbf{u}_x - \mathbf{R}\mathbf{v}_k = \mathbf{B}_x\mathbf{k} + \mathbf{e}_f$  for some final error term  $\mathbf{e}_f$ .

Provided the final error is small enough, the client can round it off to arrive at  $\lfloor \mathbf{B}_x\mathbf{k} \rfloor_p$  where  $\lfloor \cdot \rfloor_p$  denotes multiplying each coefficient by  $p/q$  and rounding to the nearest integer. Both parties additionally use NIZK proofs to show the correctness of their computations to provide security and verifiability. What we have discussed so far is effectively the high-level blueprint in [ADDS21].

The main technical difficulty in realizing this idea *efficiently* is that we need to argue that  $\mathbf{u}_x$  does not leak information about the server's key  $\mathbf{k}$ . Observe that even though  $\mathbf{u}_x$  looks like an MLWE sample, the matrix  $\mathbf{C}_x$  is controlled by the (malicious) client, and therefore, is not guaranteed to be uniformly random. To get around this problem, [ADDS21] relies on noise flooding (a.k.a. smudging) with an exponentially large error  $\mathbf{e}'_s$  with coefficients of size  $O(2^\lambda)$  to make sure that no secret information is leaked. However, then the final error term that we need to get rid of is also of size  $O(2^\lambda)$ , meaning that we need  $q = O(2^{2\lambda})$  to arrive at the correct PRF output with overwhelming probability. Noting that the overall communication size is *quadratic* in  $\log q$ , such a large modulus leads to an inefficient construction<sup>6</sup>.

Beyond using state-of-the-art NIZK proofs in our scheme, the main technical novelty of our LeOPaRd proposal is that we introduce a new method to argue server security *without* relying on noise flooding. In particular, we want to force the client into committing to its randomness  $\mathbf{R}$  *before* seeing the public matrix  $\mathbf{A}_r$ . To this end, the client first commits to  $\mathbf{R}$ , resulting in  $\mathbf{c}_r$ , and uses a hash function (modeled as a random oracle) to derive the matrix  $\mathbf{A}_r$ . One can immediately observe that the client's control in  $\mathbf{C}_x$  is now significantly limited and intuitively revealing  $\mathbf{u}_x$  as above should be secure based on the randomness of  $\mathbf{A}_r$  and the fact that the client in any case obtains a noisy vector close to  $\mathbf{B}_x\mathbf{k}$  in their final step.

**Formal security analysis via MLWE-PRF-RU assumption.** We go beyond intuitive thinking and analyze the security of the construction rigorously. The client's commitment  $\mathbf{c}_r$  to  $\mathbf{R}$  can be seen as the client submitting an oracle query with input  $\mathbf{R}$  to obtain back a random matrix  $\mathbf{A}_r$ . To realize this in the security reduction while adhering to client security, we use an *extractable* commitment scheme that enables extraction/decryption of the committed message using a trapdoor (that the reduction knows in the security analysis). We then use our new assumption, MLWE-PRF-RU, to prove the pseudorandomness property (server security) of LeOPaRd. In MLWE-PRF-RU, the adversary is given access to ( $\text{SampMLWEcom}, \text{SampPRF-RU}$ ) oracles that, on input  $(\mathbf{R}, x)$  returns a random matrix  $\mathbf{A}_r$ , an MLWE sample  $\mathbf{c}$  under  $\mathbf{A}_r$ , and a PRF-like

<sup>6</sup> We note here that [ADDS21] also used earlier NIZK proof systems, which are very inefficient compared to more state-of-the-art proofs of today. However, even without considering NIZKs, [ADDS21] incurs communication in the order of MBs.

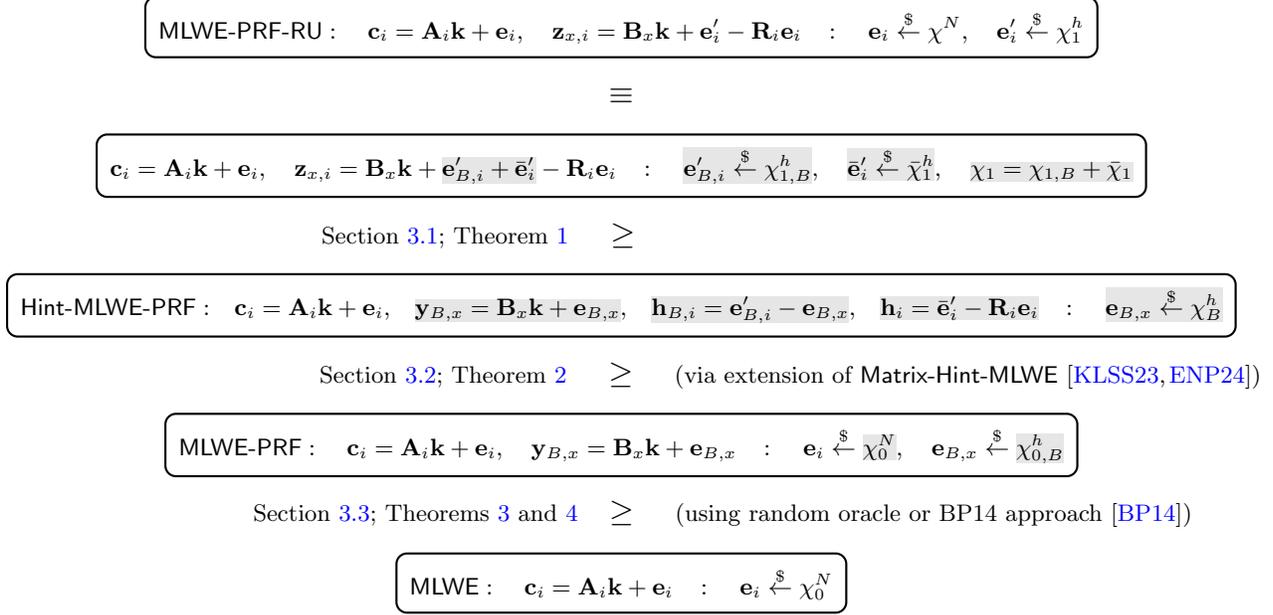


Fig. 1: High-level (simplified) summary of our reductions with changes marked in gray boxes. We denote by “≡” equivalence, and by “≥” reduction. In all experiments (where relevant),  $(\mathbf{R}_i, x)$  is an (adaptively chosen) input by an adversary, and  $\mathbf{A}_i \stackrel{\$}{\leftarrow} R_q^{N \times m}$ ,  $\mathbf{k} \stackrel{\$}{\leftarrow} \chi_k^m$  and  $\mathbf{B}_x := \mathbf{H}(x)$  for a “message mapping”  $\mathbf{H}$ .

sample  $\mathbf{z}$  that the client is able to derive in the OPRF protocol. Given these values, the term  $\mathbf{u}_x$  above is uniquely determined as  $\mathbf{u}_x := \mathbf{z} + \mathbf{R}_x \mathbf{v}_k$ . The MLWE error is re-used in both  $\mathbf{c}$  and  $\mathbf{z}$ , and hence the naming of the assumption. The client security is easier to argue and is proven based on MLWE.

**Security reduction for MLWE-PRF-RU from MLWE-PRF and MLWE.** We introduce a series of reductions to establish strong security foundations for MLWE-PRF-RU as depicted in Figure 1. First, we aim to get around the deficiency in the PRF samples and move to a more standard assumption, we call MLWE-PRF, where both MLWE and PRF samples have *independent* and proper small errors following a Gaussian distribution. That is, the MLWE-PRF assumption is that the randomized function  $x \mapsto \mathbf{H}(x) \cdot \mathbf{k} + \mathbf{e}$  with small Gaussian error  $\mathbf{e}$  is pseudorandom even when given additional MLWE samples for  $\mathbf{k}$  with respect to uniformly random matrices  $\mathbf{A}$ . This MLWE-PRF  $\rightarrow$  MLWE-PRF-RU reduction step is the main technically challenging part that we elaborate on further below. The latter step of arguing security of MLWE-PRF is easier. First of all, if we model  $\mathbf{H}$  as a random oracle, then the (randomized) PRF samples simply become MLWE samples, resulting immediately in a reduction from MLWE to MLWE-PRF, and hence to our main MLWE-PRF-RU assumption. This result demonstrates that the only way an attacker can break MLWE-PRF is by exploiting the structure of the “message mapping”  $\mathbf{H}$  (or breaking standard MLWE). We therefore consider other natural instantiations of  $\mathbf{H}$  as in earlier lattice-based PRF proposals [BLMR13, BP14], which already proved that certain algebraic instantiations of  $\mathbf{H}$  still maintains pseudorandomness based on MLWE. We extend on the results in [BP14] to base the hardness of MLWE-PRF-RU on the standard MLWE problem (see Section 3.3).

Let us now discuss the MLWE-PRF to MLWE-PRF-RU reduction at a high level. In the MLWE-PRF-RU problem<sup>7</sup>, each of the attacker’s oracle queries  $(x, \mathbf{R}_i)$  is answered with MLWE-PRF samples  $(\mathbf{A}_i, \mathbf{c}_i = \mathbf{A}_i \mathbf{k} + \mathbf{e}_i, \mathbf{z}_{x,i} = \mathbf{B}_x \mathbf{k} + \mathbf{e}'_i - \mathbf{R}_i \mathbf{e}_i)$  where  $\mathbf{B}_x := \mathbf{H}(x)$  and  $\mathbf{e}_j, \mathbf{e}'_j$  are some Gaussian errors, reminiscent of the Hint-MLWE problem [KLSS23]. Accordingly, our security reduction from MLWE-PRF to MLWE-PRF-RU proceeds in two steps. In the first step, we reduce a variant of the Hint-MLWE problem called Hint-MLWE-PRF to

<sup>7</sup> Our description here is simpler than the actual MLWE-PRF-RU assumption in Section 3. In the real definition, we consider even a stronger adversary that can adaptively request PRF samples after seeing the MLWE samples.

MLWE-PRF-RU. In the Hint-MLWE-PRF problem, besides the MLWE-like samples w.r.t.  $\mathbf{B}_x$  and  $\mathbf{A}_j$  matrices (i.e.,  $\mathbf{y}_B = \mathbf{B}_x \mathbf{k} + \mathbf{e}_B$  and  $\mathbf{c}_j = \mathbf{A}_j \mathbf{k} + \mathbf{e}_j$  as before), the adversary is also provided with appropriate linear combinations of the error vectors as hints of the form  $\mathbf{h}_j := \mathbf{e}'_j - \mathbf{R}_j \mathbf{e}_j$  and  $\mathbf{h}_{B,j} := \mathbf{e}'_{B,j} - \mathbf{e}_B$ . These hints allow the reduction from MLWE-PRF to MLWE-PRF-RU to simulate the ‘deficient’ PRF samples as  $\mathbf{z}_{x,j} := \mathbf{y}_B + \mathbf{h}_{B,j} + \mathbf{h}_j = \mathbf{B}_x \mathbf{k} + \bar{\mathbf{e}}'_j - \mathbf{R}_j \mathbf{e}_j$ , where  $\bar{\mathbf{e}}'_j := \mathbf{e}'_j + \mathbf{e}'_{B,j}$  is the simulated error term for  $\mathbf{z}_{x,j}$ .<sup>8</sup> Note that in Hint-MLWE-PRF each  $\mathbf{h}_j$  hint contains a *linear combination of several* error vectors (rather than hints containing only *scalar multiples* of the error vector as in [KLSS23]). Hence Hint-MLWE-PRF can be viewed as a variant of the *Matrix* Hint-MLWE [ENP24] that generalizes Hint-MLWE [KLSS23] to the case of hint matrices  $\mathbf{R}_j$ . Our Hint-MLWE-PRF variant further includes a HintPRF oracle that provides the PRF samples  $\mathbf{y}_B$  along with the hints  $\mathbf{h}_{B,j}$  on the PRF error. In the second step, we present a reduction from MLWE-PRF to Hint-MLWE-PRF by extending the MLWE to Matrix Hint-MLWE [ENP24] reduction.

Overall, we achieve a reduction from MLWE to MLWE-PRF-RU, provided that the necessary parameter constraints are satisfied. Our MLWE to MLWE-PRF-RU reduction is essentially tight in the sense that our reduction’s parameter condition on the underlying error width parameter is *necessary* to prevent an existing *attack* on MLWE-PRF-RU. In particular, our reduction requires the “masking ratio” lower bound  $\sigma_1/\sigma \geq \sqrt{Q_x^\infty}$  for large  $Q_x^\infty$ , where  $\sigma_1$  and  $\sigma$  denote the width (std dev) of the masking errors  $\mathbf{e}'_j$  and MLWE errors  $\mathbf{e}_j$ , respectively, and  $Q_x^\infty$  denotes an upper bound on the oracle samples queried by the adversary *per input*  $x$  (corresponding to input/tag pair  $(x, t)$  in our POPRF protocol). On the other hand, this lower bound is necessary to protect against a simple error “averaging” attack on MLWE-PRF-RU (and our OPRF protocol as well as those using the same blueprint [ADDS21, AG24]), which we describe in Appendix B.

## 2 Preliminaries

**Notation.** We denote the security parameter by  $\lambda \in \mathbb{N}$ , and  $\kappa \in \mathbb{N}$  denotes the statistical correctness parameter. We will aim to obtain the correct PRF value except with probability at most  $2^{-\kappa}$ . For  $n \in \mathbb{N}$ , set  $[n] = \{1, 2, \dots, n\}$ . We denote by  $x \stackrel{\$}{\leftarrow} X$  the uniform sampling of the variable  $x$  from the set  $X$ , and we denote the uniform distributions on a set  $X$  by  $\mathcal{U}(X)$ . We write  $x \leftarrow \mathbf{A}(y)$  to denote that a probabilistic polynomial time (PPT) algorithm  $\mathbf{A}$  on input  $y$ , outputs  $x$ . When we want to make the randomness  $r$  used by  $\mathbf{A}$  explicit, we write it as  $x \leftarrow \mathbf{A}(y; r)$ . If  $\mathbf{A}$  has oracle access to a procedure  $\mathbf{O}$ , we represent this by superscript  $\mathbf{A}^{\mathbf{O}}$ . If  $\mathbf{A}$  is a deterministic polynomial time (DPT) algorithm, we use the notation  $x := \mathbf{A}(y)$ . We use the same notation for the projection of tuples, e.g., we write  $\sigma := (\sigma_1, \sigma_2)$  for a tuple  $\sigma$  composed of two elements  $\sigma_1$  and  $\sigma_2$ . We define *polynomial* functions as  $\text{poly}(\lambda) = \bigcup_{d \in \mathbb{N}} O(\lambda^d)$  and *negligible* functions as  $\text{negl}(\lambda) = \bigcap_{d \in \mathbb{N}} o(\lambda^{-d})$ . We write the convolution of two distributions  $D_1$  and  $D_2$  as  $D_1 + D_2$ .

We denote by  $R_q = \mathbb{Z}_q[X]/(X^d + 1)$ . We write  $R_{\text{bin}}$  to denote the set of polynomials in  $R_q$  that have binary coefficients. For  $a \in \mathbb{Z}^+$ , we use  $\mathbb{S}_a$  to denote the set of polynomials in  $R_q$  with infinity norm at most  $a$ . The rounding operation  $[a]_p : R_q \rightarrow R_p$  is defined as multiplying  $a$  by  $p/q$  and rounding each resulting coefficient to the nearest integer. When  $\mathbf{a}$  is a vector, then by  $[\mathbf{a}]_p$  we consider coordinate-wise rounding. For a matrix  $\mathbf{R}$ , we denote by  $\|\mathbf{R}\|$ ,  $\|\mathbf{R}\|_1$  and  $\|\mathbf{R}\|_{\infty, M}$  its matrix 2-norm (largest singular value), matrix 1-norm and matrix  $\infty$ -norm, respectively and by  $\sigma_{\min}(\mathbf{R})$  and  $\sigma_{\max}(\mathbf{R})$  the smallest (resp. largest) singular values of  $\mathbf{R}$ . We denote by  $\|\mathbf{R}\|_\infty$  the entry-wise maximum absolute value of  $\mathbf{R}$ .

As the NIZK relations to be proven get more complex for various protocols, we believe it is imperative to have a more explicit and concise notation. We introduce the notation “ $\mathcal{D}(x); \mathcal{K}(w)$ ” to describe a NIZK relation to mean “given a statement  $x$  and knowledge of witness  $w$ ”.

Next, we provide some of the preliminaries needed for the rest of the paper, and refer the reader to Appendix A for additional preliminaries.

<sup>8</sup> We remark that this reduction could also directly use a single combined hint of the form  $\bar{\mathbf{h}}_j := \mathbf{h}_{B,j} + \mathbf{h}_j = \bar{\mathbf{e}}'_j - \mathbf{R}_j \mathbf{e}_j - \mathbf{e}_B$ . However, such a combined variant of Hint-MLWE makes it difficult to deal with *adaptive* oracle  $\mathbf{R}_j$  queries in the second step of reducing from MLWE-PRF to our Hint-MLWE-PRF problem. Hence, we split the hint into the two separate hints  $\mathbf{h}_j$  and  $\mathbf{h}_{B,j}$ .

## 2.1 (Partially) Oblivious PRF

We use a slightly modified version of the game-based definitions for (round-optimal) OPRF with preprocessing given in [DDT25], which in turn extends the definition given in [TCR+22]. An OPRF is a protocol between a server  $S$  that has a private key  $k$  and a client  $C$  that wants to obtain an evaluation of PRF  $F_k$  on inputs of its choice. We say that an OPRF is a partial OPRF (POPRF) if part of the client's input is given to the server, i.e.,  $y = F_k(t, x)$ , where  $t$  is in the public part and  $x$  is the private part hidden from  $S$ . When the client can verify the correctness and consistency of the PRF evaluations, then we consider verifiable OPRF (VOPRF).

**Definition 1 (Partial Oblivious PRF with Preprocessing).** *A partial oblivious PRF (POPRF)  $F$  is a tuple of PPT algorithms*

(Setup, KeyGen, PreProcClient, PreProcServer, Request, BlindEval, Finalize, Eval).

*The Setup and KeyGen algorithms generate public parameters  $pp$  and a public/secret key pair  $(pk, sk)$ , respectively. The preprocessing is a (single-round) protocol between the client  $C$  and server  $S$ , presented as algorithms  $F.PreProcClient$  and  $F.PreProcServer$ , which work as follows (where the random oracles are denoted by  $RO$ ):*

1.  $C$  runs the algorithm  $F.PreProcClient_{pp}^{RO}(pk, T)$  taking as input a public key  $pk$  and an evaluation bound  $T$ . It outputs a local client state  $st_C$  and a preprocessing request message  $preq$ , which is sent to  $S$ .
2.  $S$  runs  $F.PreProcServer_{pp}^{RO}(sk, preq)$  taking as input a secret key  $sk$  and the preprocessing request message  $preq$ . It produces a local server state  $st_S$  and a preprocessing response message  $prep$ , which is sent to  $C$ . (Without loss of generality we assume that the client  $C$  includes the evaluation bound  $T$  and response message  $prep$  in the state  $st_C$ .)

*Oblivious evaluation is carried out as a (single-round) protocol between  $C$  and  $S$ , presented as algorithms  $F.Request$ ,  $F.BlindEval$ ,  $F.Finalize$ , which work as follows:*

1. First,  $C$  runs the algorithm  $F.Request_{pp}^{RO}(pk, t, x, st_C)$  taking as input a public key  $pk$ , a (public) tag  $t$ , a private input  $x$  and previously constructed client state  $st_C$ . It outputs an updated client state  $st_C$  and a request message  $req$  that is sent to  $S$ , or outputs  $\perp$  (a designated error symbol).
2.  $S$  runs  $F.BlindEval_{pp}^{RO}(sk, t, req, st_S)$  taking as input a secret key  $sk$ , a (public) tag  $t$ , the request message  $req$  and previously constructed server state  $st_S$ . It produces a response message  $rep$ , which is sent to  $C$ .
3. Finally,  $C$  runs  $F.Finalize_{pp}^{RO}(rep, st_C)$  taking as input the response message  $rep$  and the previously constructed client state  $st_C$ . It outputs either a PRF evaluation  $y$  or  $\perp$  if  $rep$  is rejected.

*The unblinded evaluation algorithm  $F.Eval$  is deterministic and takes as input a secret key  $sk$ , an input pair  $(t, x)$ , and outputs a PRF evaluation  $y$ .*

*We also define the sets  $\mathcal{SK}, \mathcal{PK}, \mathcal{T}, \mathcal{X}, \mathcal{Z}$ , representing the secret key, public key, (public) tag, private input, and output spaces, respectively. We define the input space as  $\mathcal{T} \times \mathcal{X}$ . We assume there exist efficient algorithms for sampling and membership queries on these sets.*

We remark that fixing  $t = \perp$ , gives us the definition of (plain) OPRF. We extend the correctness definition from [ADDG24], which allows for a small failure probability.

**Definition 2 (Correctness).** *A partial oblivious PRF (POPRF)  $F$  is correct, if for all  $\lambda, T \in \mathbb{N}$ , for all  $pp \leftarrow F.Setup(1^\lambda)$  and  $(pk, sk) \leftarrow F.KeyGen_{pp}^{RO}(1^\lambda)$ , it holds that*

$$\Pr \left[ y = F.Eval_{pp}^{RO}(sk, t, x) \mid \begin{array}{l} (st_C, preq) \leftarrow F.PreProcClient_{pp}^{RO}(pk, T) \\ (st_S, prep) \leftarrow F.PreProcServer_{pp}^{RO}(sk, preq) \\ (st_C, req) \leftarrow F.Request_{pp}^{RO}(pk, t, x, st_C) \\ rep \leftarrow F.BlindEval_{pp}^{RO}(sk, t, req, st_S) \\ y \leftarrow F.Finalize_{pp}^{RO}(rep, st_C) \end{array} \right] = 1 - \text{negl}(\lambda).$$

<p><b>POPFR<math>_{F,\mathcal{A},\mathcal{S},\text{RO}}^b(\lambda)</math></b></p> <hr/> <p>1 : <math>q_p := 0, \mathcal{D} := \perp</math>  2 : <math>\text{st}_F \leftarrow \text{RO}_F.\text{Init}(1^\lambda) \quad // \mathcal{T} \times \mathcal{X} \rightarrow \mathcal{Z}</math>  3 : <math>\text{st}_{\text{RO}} \leftarrow \text{RO}.\text{Init}(1^\lambda)</math>  4 : <math>\text{pp}_1 \leftarrow F.\text{Setup}(1^\lambda)</math>  5 : <math>(\text{st}_S, \text{pk}_0, \text{pp}_0) \leftarrow \mathcal{S}.\text{Init}(\text{pp}_1)</math>  6 : <math>(\text{sk}, \text{pk}_1) \leftarrow F.\text{KeyGen}_{\text{pp}_1}^{\text{RO}}(1^\lambda)</math>  7 : <math>b' \leftarrow \mathcal{A}^{\text{PreProcServer, Eval, BlindEval, Prim}}(\text{pp}_b, \text{pk}_b)</math>  8 : <b>return</b> <math>b'</math></p> <p><b>BlindEval</b>(<math>i, t, \text{req}</math>)</p> <hr/> <p>1 : <b>if</b> <math>i \notin [q_p]</math> <b>then return</b> <math>\perp</math>  2 : <b>parse</b> <math>\mathcal{D}[i] := \text{st}_S</math>  3 : <math>(\text{rep}_0, \text{st}_S) \leftarrow \mathcal{S}.\text{BlindEval}(t, \text{req}, \text{st}_S)</math>  4 : <math>\text{rep}_1 \leftarrow F.\text{BlindEval}_{\text{pp}_1}^{\text{RO}}(\text{sk}, t, \text{req}, \text{st}_S)</math>  5 : <b>return</b> <math>\text{rep}_b</math></p>	<p><b>PreProcServer</b>(<math>\text{preq}</math>)</p> <hr/> <p>1 : <math>q_p := q_p + 1</math>  2 : <math>(\text{prep}_0, \text{st}_S) \leftarrow \mathcal{S}.\text{PreProcServer}(\text{preq}, \text{st}_S)</math>  3 : <math>(\text{st}_S, \text{prep}_1) \leftarrow F.\text{PreProcServer}_{\text{pp}_1}^{\text{RO}}(\text{sk}, \text{preq})</math>  4 : <math>\mathcal{D}[q_p] := \text{st}_S</math>  5 : <b>return</b> <math>\text{prep}_b</math></p> <p><b>Eval</b>(<math>t, x</math>)</p> <hr/> <p>1 : <math>z_0 \leftarrow \text{RO}_F.\text{Eval}((t, x), \text{st}_F)</math>  2 : <math>z_1 \leftarrow F.\text{Eval}_{\text{pp}_1}^{\text{RO}}(\text{sk}, t, x)</math>  3 : <b>return</b> <math>z_b</math></p> <p><b>Prim</b>(<math>x</math>)</p> <hr/> <p>1 : <math>(h_0, \text{st}_S) \leftarrow \mathcal{S}.\text{Eval}(x, \text{st}_S)</math>  2 : <math>h_1 \leftarrow \text{RO}.\text{Eval}(x, \text{st}_{\text{RO}})</math>  3 : <b>return</b> <math>h_b</math></p>
<p><b>POPRIV2<math>_{F,\mathcal{A},\text{RO}}^b(\lambda)</math></b></p> <hr/> <p>1 : <math>q_r := 0, q_p := 0, \mathcal{D}_p := \mathcal{D}_r := \perp</math>  2 : <math>\text{pp} \leftarrow F.\text{Setup}(1^\lambda)</math>  3 : <math>b' \leftarrow \mathcal{A}^{\text{PreProcClient, Request, Finalize, RO}}(\text{pp})</math>  4 : <b>return</b> <math>b'</math></p> <p><b>Request</b>(<math>i, \text{pk}, \text{prep}, t, x_0, x_1</math>)</p> <hr/> <p>1 : <b>if</b> <math>i \notin [q_p]</math> <b>then return</b> <math>\perp</math>  2 : <b>parse</b> <math>\mathcal{D}_p[i] := \text{st}_C</math>  3 : <math>\text{st}_C := \text{st}_C \cup \{\text{prep}\}</math>  4 : <math>(\text{st}_{C,q_r,0}, \text{req}_0) \leftarrow F.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x_0, \text{st}_C)</math>  5 : <math>(\text{st}_{C,q_r,1}, \text{req}_1) \leftarrow F.\text{Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x_1, \text{st}_C)</math>  6 : <b>if</b> <math>\text{req}_0 = \perp \vee \text{req}_1 = \perp</math> <b>then</b>  7 :     <b>return</b> <math>\perp</math>  8 : <math>q_r := q_r + 1</math>  9 : <math>\mathcal{D}_r[i, q_r] := (\text{st}_{C,q_r,0}, \text{st}_{C,q_r,1})</math>  10 : <b>return</b> <math>(\text{req}_b, \text{req}_{1-b})</math></p>	<p><b>PreProcClient</b>(<math>\text{pk}, T</math>)</p> <hr/> <p>1 : <math>q_p := q_p + 1</math>  2 : <math>(\text{st}_C, \text{preq}) \leftarrow F.\text{PreProcClient}_{\text{pp}}^{\text{RO}}(\text{pk}, T)</math>  3 : <math>\mathcal{D}_p[q_p] := \text{st}_C</math>  4 : <b>return</b> <math>\text{preq}</math></p> <p><b>Finalize</b>(<math>i, j, \text{rep}_0, \text{rep}_1</math>)</p> <hr/> <p>1 : <b>if</b> <math>i \notin [q_p] \vee j \notin [q_r]</math> <b>then</b>  2 :     <b>return</b> <math>\perp</math>  3 : <b>parse</b> <math>\mathcal{D}_r[i, j] := (\text{st}_{C,j,0}, \text{st}_{C,j,1})</math>  4 : <math>y_b \leftarrow F.\text{Finalize}_{\text{pp}}^{\text{RO}}(\text{rep}_0, \text{st}_{C,j,b})</math>  5 : <math>y_{1-b} \leftarrow F.\text{Finalize}_{\text{pp}}^{\text{RO}}(\text{rep}_1, \text{st}_{j,1-b})</math>  6 : <b>if</b> <math>y_0 = \perp \vee y_1 = \perp</math> <b>then</b>  7 :     <b>return</b> <math>\perp</math>  8 : <b>return</b> <math>(y_0, y_1)</math></p>

Fig. 2: POPRF and POPRIV2 experiments.

Next, we extend the *pseudorandomness* against malicious clients definition from [TCR<sup>+</sup>22] to account for preprocessing.

**Definition 3 (Pseudorandomness).** *We say that a partial oblivious PRF (POPFR)  $F$  is pseudorandom if for all (stateful) PPT adversaries  $\mathcal{A}$ , there exists a PPT simulator  $\mathcal{S}$  and a negligible function  $\text{negl}(\cdot)$ , such that the following holds,*

$$\text{Adv}_{F,\mathcal{A},\mathcal{S},\text{RO}}^{\text{po-prf}}(\lambda) = \left| \Pr [\text{POPFR}_{F,\mathcal{A},\mathcal{S},\text{RO}}^1(\lambda) = 1] - \Pr [\text{POPFR}_{F,\mathcal{A},\mathcal{S},\text{RO}}^0(\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where the experiment POPRF is defined in Figure 2.

*Remark 1.* The oracle Prim captures access to the random oracle(s) used in the OPRF construction. We also remark that the original pseudorandomness security game given in [TCR<sup>+</sup>22] additionally included a LimitEval oracle, which was given to the simulator  $\mathcal{S}$  inside the BlindEval and Prim oracles. The LimitEval oracle allowed the simulator to obtain a limited number of evaluation outputs, which is quite useful when proving the security of the underlying OPRF construction using a “one-more”-type assumption. However, we do not use any such assumption, and hence, our simulator does not need the LimitEval oracle. Therefore, we opted to remove it from the security model for the sake of clarity.

Lastly, we define the *request privacy*, which provides security against malicious servers in the preprocessing model. We present the stronger POPRIV2 definition satisfied by our proposal, and refer the reader to [DDT25] for the weaker POPRIV1 definition.

**Definition 4 (Request Privacy).** *We say that a POPRF  $\mathcal{F}$  has request privacy against malicious servers if for all PPT adversaries, there exists a negligible function  $\text{negl}(\cdot)$  such that the following holds,*

$$\text{Adv}_{\mathcal{F}, \mathcal{A}, \text{RO}}^{\text{po-priv-2}}(\lambda) = |\Pr[\text{POPRIV2}_{\mathcal{F}, \mathcal{A}, \text{RO}}^1(\lambda) = 1] - \Pr[\text{POPRIV2}_{\mathcal{F}, \mathcal{A}, \text{RO}}^0(\lambda) = 1]| \leq \text{negl}(\lambda),$$

where the experiment POPRIV2 is defined in Figure 2.

## 2.2 Lattice Preliminaries

We start with the definitions of standard problems in lattice-based cryptography.

**Definition 5 (MSIS $_{n,m,\beta}$ ).** *Let  $\mathbf{A} \xleftarrow{\$} R_q^{n \times m}$ , where  $n, m > 0$  and let  $0 < \beta < q$ . We say  $\mathbf{z} \in R_q^m$  is a solution for MSIS $_{n,m,\beta}$  problem if  $\mathbf{A}\mathbf{z} = \mathbf{0}$  over  $R_q$  and  $0 < \|\mathbf{z}\| < \beta$ . If the following inequality holds for an adversary  $\mathcal{A}$*

$$\Pr[0 < \|\mathbf{z}\| < \beta \wedge \mathbf{A}\mathbf{z} = \mathbf{0} \mid \mathbf{A} \xleftarrow{\$} R_q^{n \times m}, \mathbf{z} \leftarrow \mathcal{A}(\mathbf{A})] \geq \epsilon,$$

then we say  $\mathcal{A}$  has advantage  $\epsilon$  in solving MSIS $_{n,m,\beta}$ .

**Definition 6 (MLWE $_{\ell,q,m,r,\chi_e,\chi_k}$ ).** *Let  $\chi_e$  and  $\chi_k$  be error and secret distributions over  $R$  respectively,  $\mathbf{A} \xleftarrow{\$} R_q^{\ell \times m}$ , where  $\ell$  denotes the number of samples,  $m > 0$  denotes the dimension per secret vector, let  $\mathbf{S} \xleftarrow{\$} \chi_k^{m \times r}$  be a matrix of  $r$  secret vectors, and  $\mathbf{E} \xleftarrow{\$} \chi_e^{\ell \times r}$  be a matrix of  $r$  error vectors. The MLWE $_{\ell,q,m,r,\chi_e,\chi_k}$  problem asks an adversary  $\mathcal{A}$  to distinguish between  $(\mathbf{A}, \mathbf{A}\mathbf{S} + \mathbf{E})$  and  $(\mathbf{A}, \mathbf{B})$ , for  $\mathbf{B} \xleftarrow{\$} R_q^{\ell \times r}$ . We say  $\mathcal{A}$  has advantage  $\epsilon$  against the MLWE $_{\ell,q,m,r,\chi_e,\chi_k}$  problem if*

$$\begin{aligned} & \Pr[b = 1 \mid \mathbf{A} \xleftarrow{\$} R_q^{\ell \times m}, \mathbf{S} \xleftarrow{\$} \chi_k^{m \times r}, \mathbf{E} \xleftarrow{\$} \chi_e^{\ell \times r}, \mathbf{B} := \mathbf{A}\mathbf{S} + \mathbf{E}, b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{B})] \\ & - \Pr[b = 1 \mid \mathbf{A} \xleftarrow{\$} R_q^{\ell \times m}, \mathbf{B} \xleftarrow{\$} R_q^{\ell \times r}, b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{B})] \geq \epsilon. \end{aligned}$$

If  $r = 1$  in the above MLWE definition, then we simply write MLWE $_{\ell,q,m,\chi_e,\chi_k}$ . We also sometimes drop the parameter  $\ell$  when the number of samples is not critical for the discussion.

**Definition 7 (knMLWE $_{n,m,h,\chi}$ ).** *Let  $\chi$  be an error distribution over  $R$ ,  $\mathbf{A} \xleftarrow{\$} R_q^{n \times m}$ , where  $n > m > 0$ , and let  $\mathbf{S} \xleftarrow{\$} \chi^{h \times n}$ . The knMLWE $_{n,m,h,\chi}$  problem asks an adversary  $\mathcal{A}$  to distinguish between  $(\mathbf{A}, \mathbf{S}\mathbf{A})$  and  $(\mathbf{A}, \mathbf{U})$ , for  $\mathbf{U} \xleftarrow{\$} R_q^{h \times m}$ . We say  $\mathcal{A}$  has advantage  $\epsilon$  against the knMLWE $_{n,m,h,\chi}$  problem if*

$$\begin{aligned} & \Pr[b = 1 \mid \mathbf{A} \xleftarrow{\$} R_q^{n \times m}, \mathbf{S} \xleftarrow{\$} \chi^{h \times n}, b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{S}\mathbf{A})] \\ & - \Pr[b = 1 \mid \mathbf{A} \xleftarrow{\$} R_q^{n \times m}, \mathbf{U} \xleftarrow{\$} R_q^{h \times m}, b \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{U})] \geq \epsilon. \end{aligned}$$

The duality between the above Knapsack form of MLWE and regular MLWE has been established already. That is, from a practical security perspective,  $\text{knMLWE}_{n,m,h,\chi}$  is as hard as  $\text{MLWE}_{n-m,m,\chi}$ , which we will use to estimate the practical hardness of  $\text{knMLWE}_{n,m,h,\chi}$ . We refer to [EZS<sup>+</sup>19, Appendix C] for a summary discussion (and references therein for more details). Also, note that our  $\text{knMLWE}$  definition is in the “multi-secret” setting where multiple vectors are multiplied by the matrix  $\mathbf{A}$ . A standard hybrid argument of replacing each  $\mathbf{s}_i^\top \mathbf{A}$  with a uniformly random  $\mathbf{u}_i^\top$  reduces its hardness to standard Knapsack form of MLWE.

The second parameter  $m$  in the definitions of MSIS and  $\text{knMLWE}$  does not play a critical role in our practical hardness estimations of these problems against state-of-the-art attacks (i.e., we assume that the adversary has the optimal choice of  $m$ ). Therefore, we sometimes simply omit this parameter.

**Definition 8 (Discrete Gaussian Distribution).** For any  $s > 0$  and dimension  $n \in \mathbb{Z}^+$ , the spherical  $n$ -dimensional Gaussian function with parameter  $s$ <sup>9</sup> is defined as  $\rho_s(\mathbf{x}) := \exp(-\pi \|\mathbf{x}\|^2/s^2)$  for  $\mathbf{x} \in \mathbb{R}^n$ . More generally, given a positive definite symmetric covariance parameter matrix  $\Sigma \in \mathbb{R}^{n \times n}$  and center  $\mathbf{c}$ , the  $n$ -dimensional Gaussian function with covariance parameter  $\Sigma$  and center  $\mathbf{c}$  is defined as  $\rho_{\Sigma, \mathbf{c}}(\mathbf{x}) := \exp(-\pi(\mathbf{x} - \mathbf{c})^\top \Sigma^{-1}(\mathbf{x} - \mathbf{c}))$  for  $\mathbf{x} \in \mathbb{R}^n$ . The discrete Gaussian distribution  $\mathcal{D}_{\Lambda, \Sigma, \mathbf{c}}$  over an  $n$ -dimensional lattice  $\Lambda \subseteq \mathbb{R}^n$  with covariance parameter  $\Sigma$ , centre  $\mathbf{c}$  and support  $\Lambda$  is defined as  $D_{\Lambda, \Sigma, \mathbf{c}}(\mathbf{x}) := \frac{\rho_{\Sigma, \mathbf{c}}(\mathbf{x})}{\sum_{\mathbf{y} \in \Lambda} \rho_{\Sigma, \mathbf{c}}(\mathbf{y})}$  for  $\mathbf{x} \in \Lambda$ . In the spherical case, where  $\Sigma = s^2 \mathbf{I}_n$ , we write  $D_{\Lambda, s, \mathbf{c}}$ , and we omit  $\mathbf{c}$  if it is  $\mathbf{0}$ .

**Lemma 1 ([DPS23], [BLP<sup>+</sup>13, Lemma 2.3]).** There exists a PPT algorithm that, given a basis  $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_n)$  of a full rank  $n$ -dimensional lattice  $\Lambda$ , a positive definite symmetric matrix  $\Sigma$  and a center  $\mathbf{c} \in \mathbb{R}^n$ , returns a sample from  $\mathcal{D}_{\Lambda, \Sigma, \mathbf{c}}$ , assuming the condition  $\sqrt{\ln(2n+4)/\pi} \cdot \max_i \|\Sigma^{-1/2} \mathbf{b}_i\| \leq 1$ .

**Lemma 2 ([KLSS23], [Pei10]).** For  $n \in \mathbb{Z}^+, \epsilon \in \mathbb{R}^+$ , let  $\Sigma_1, \Sigma_2$  be positive definite symmetric matrices such that  $\Sigma_3^{-1} := \Sigma_1^{-1} + \Sigma_2^{-1}$  satisfies  $\sqrt{\Sigma_3} \geq \eta_\epsilon(\mathbb{Z}^n)$  for  $0 < \epsilon < 1/2$ . Then for an arbitrary center  $\mathbf{c} \in \mathbb{Z}^n$ , the distribution

$$\{\mathbf{x}_1 + \mathbf{x}_2 : \mathbf{x}_1 \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^n, \Sigma_1}, \mathbf{x}_2 \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^n, \Sigma_2, \mathbf{c}}\},$$

is within statistical distance  $\leq 2\epsilon$  of  $\mathcal{D}_{\mathbb{Z}^n, \Sigma_1 + \Sigma_2, \mathbf{c}}$ .

**Definition 9 (Smoothing Parameter [MR04, Pei10]).** For an  $n$ -dimensional lattice  $\Lambda$  and  $\epsilon > 0$ , the smoothing parameter  $\eta_\epsilon(\Lambda)$  is the smallest  $s$  such that  $\rho_{1/s}(\Lambda^* \setminus \mathbf{0}) \leq \epsilon$ , where  $\Lambda^*$  denotes the dual lattice of  $\Lambda$ . More generally, for a positive definite symmetric matrix  $\Sigma$ , we say that  $\sqrt{\Sigma} \geq \eta_\epsilon(\Lambda)$  if  $\eta_\epsilon(\sqrt{\Sigma}^{-1} \cdot \Lambda) \leq 1$ .

**Lemma 3 ([MR04]).** For  $n \in \mathbb{Z}^+, \epsilon \in \mathbb{R}^+$ , and  $n$  dimensional lattice  $\Lambda$ , we have

$$\eta_\epsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\epsilon))}{\pi}} \cdot \lambda_n(\Lambda), \quad (1)$$

where  $\lambda_n(\Lambda)$  is the smallest radius of an  $n$ -dimensional ball that contains  $n$  linearly independent vectors in  $\Lambda$ .

**Lemma 4 ([KLSS23]).** For  $n \in \mathbb{Z}^+, \epsilon \in \mathbb{R}^+$ ,  $n$  dimensional lattice  $\Lambda$ , and a positive definite symmetric matrix  $\Sigma$ , we have  $\sqrt{\Sigma} \geq \eta_\epsilon(\Lambda)$  if  $\|\Sigma^{-1}\| \leq \eta_\epsilon(\Lambda)^{-2}$ .

**Lemma 5 (Generalization of [ENP24, Lemma 1] to non-square hint matrices).** Fix  $s_1 \in \mathbb{R}^+$ ,  $n, \ell \in \mathbb{Z}^+$ , a positive definite symmetric matrix  $\Sigma_1$ , and a matrix  $\mathbf{R} \in \mathbb{Z}^{\ell \times n}$ . Furthermore, let  $\Sigma_0 := (\Sigma_1^{-1} + \frac{1}{s_1^2} \mathbf{R}^\top \mathbf{R})^{-1}$ . Then, the following two probability distributions over  $\mathbb{Z}^{n+\ell}$  are equal:

$$\mathcal{D}_1 := \left\{ (\mathbf{e}, \mathbf{h}) : \mathbf{e} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^n, \Sigma_1}, \mathbf{e}' \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^\ell, s_1}, \mathbf{h} = \mathbf{R}\mathbf{e} + \mathbf{e}' \right\}$$

$$\mathcal{D}_2 := \left\{ (\tilde{\mathbf{e}}, \mathbf{h}) : \mathbf{e} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^n, \Sigma_1}, \mathbf{e}' \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^\ell, s_1}, \mathbf{h} = \mathbf{R}\mathbf{e} + \mathbf{e}', \mathbf{c} := \frac{1}{s_1^2} \Sigma_0 \mathbf{R}\mathbf{h}, \tilde{\mathbf{e}} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^\ell, \Sigma_0, \mathbf{c}} \right\}$$

<sup>9</sup> Note that the parameter  $s$  is related to the standard deviation  $\sigma$  by  $s = \sqrt{2\pi} \cdot \sigma$ .

*Proof.* Fix any  $(\mathbf{v}, \mathbf{w}) \in \mathbb{Z}^n \times \mathbb{Z}^\ell$ . For the first distribution we have:

$$\begin{aligned} \mathcal{D}_1(\mathbf{v}, \mathbf{w}) &= \mathcal{D}_{\mathbb{Z}^n, \Sigma_1}(\mathbf{v}) \cdot \mathcal{D}_{\mathbb{Z}^n, s_1}(\mathbf{w} - \mathbf{R}\mathbf{v}) \\ &\propto \exp \left[ -\pi \left( \mathbf{v}^\top \Sigma_1^{-1} \mathbf{v} - \frac{1}{s_1^2} (\mathbf{w} - \mathbf{R}\mathbf{v})^\top (\mathbf{w} - \mathbf{R}\mathbf{v}) \right) \right] \\ &= \exp \left[ -\pi \left( (\mathbf{v} - \mathbf{c})^\top \Sigma_0^{-1} (\mathbf{v} - \mathbf{c}) + \frac{1}{s_1^2} \mathbf{w}^\top \mathbf{w} - \mathbf{c}^\top \Sigma_0^{-1} \mathbf{c} \right) \right], \end{aligned}$$

where  $\mathbf{c} := \frac{1}{s_1^2} \Sigma_0 \mathbf{R}^\top \mathbf{w}$ . Since  $\frac{1}{s_1^2} \mathbf{w}^\top \mathbf{w} - \mathbf{c}^\top \Sigma_0^{-1} \mathbf{c}$  is a constant that does not depend on  $\mathbf{v}$ , it follows that the conditional distribution  $\Pr_{(\mathbf{e}, \mathbf{h}) \leftarrow \mathcal{D}_1}[\mathbf{e} = \mathbf{v} | \mathbf{h} = \mathbf{w}] = \mathcal{D}_{\mathbb{Z}^n, \Sigma_0, \mathbf{c}}(\mathbf{v})$ , which is (by construction of  $\mathcal{D}_2$ ) exactly equal to the conditional distribution  $\Pr_{(\tilde{\mathbf{e}}, \mathbf{h}) \leftarrow \mathcal{D}_2}[\tilde{\mathbf{e}} = \mathbf{v} | \mathbf{h} = \mathbf{w}]$  in  $\mathcal{D}_2$ . Since the marginal distribution of  $\mathbf{h}$  is also exactly the same in  $\mathcal{D}_1$  and  $\mathcal{D}_2$ , we conclude that  $\mathcal{D}_1 = \mathcal{D}_2$ .  $\square$

### 3 MLWE-PRF-RU: New Lattice Problem with a Mix of MLWE and PRF Samples

We introduce a family of new (interactive) assumptions where MLWE and (randomized) PRF samples are mixed up. We call this family *MLWE-PRF with Re-Use*, MLWE-PRF-RU, and instantiate a member, MLWE-PRF-RU<sub>H</sub>, by choosing an appropriate “message mapping”  $H$  for the PRF samples.

We prove a set of reductions to establish strong evidence for the hardness of MLWE-PRF-RU. A high-level (simplified) summary of the problems and the reductions we achieve is depicted in Figure 1. We first introduce an intermediary problem, Hint-MLWE-PRF, and show a reduction from Hint-MLWE-PRF to MLWE-PRF-RU. Here, we interpret the error information leaked in MLWE-PRF-RU as additional hints (hence the name of the problem). Then, we introduce the main problem that forms the basis for the hardness of MLWE-PRF-RU, namely MLWE-PRF. This problem is a mild variant of the MLWE problem, where the adversary is additionally given randomized PRF outputs. We show a reduction from MLWE-PRF to Hint-MLWE-PRF, which overall leads a reduction from MLWE-PRF to MLWE-PRF-RU. As the final piece of the puzzle, we show a reduction from MLWE to MLWE-PRF for natural instantiations of  $H$ , including modelling  $H$  as a random oracle and using the mapping from BP14 PRF [BP14].

In the MLWE-PRF-RU problem, the adversary is allowed to *adaptively* query an adversarially-chosen matrix  $\mathbf{R}$  to an oracle that returns an MLWE sample. After seeing the MLWE samples, the adversary can then (again adaptively) obtain “deficient” randomized PRF samples of the form  $H(x) \cdot \mathbf{k} + \mathbf{e}_p$  with  $\mathbf{e}_p$  dependent on the MLWE error  $\mathbf{e}$  and the matrix  $\mathbf{R}$ , where  $\mathbf{e}_p$  has the form  $\mathbf{e}_p = \mathbf{e}' - \mathbf{R}\mathbf{e}$ , and  $\mathbf{e}'$  is a fresh ‘masking’ error. The assumption states that this real case is indistinguishable from random. Importantly, although the MLWE samples (returned by `SampMLWEcom` oracle in Figure 3) are independent of  $\mathbf{R}$ , the adversary has to *commit* to  $\mathbf{R}$  in advance of seeing the MLWE samples (thus  $\mathbf{R}$  is an input of `SampMLWEcom`) since the MLWE error is reused in the deficient PRF samples. Looking ahead, this matches the (adversarial) behaviour in the OPRF pseudorandomness proof. We formally define the problem below.

**Definition 10** (MLWE-PRF-RU <sub>$H, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}$</sub> ). *Let  $\vec{\chi} = (\chi, \chi_1, \chi_k)$  be discrete distributions over  $R$ ,*

*let  $H$  be a hash family, and let  $Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r \in \mathbb{Z}^+$ . We say that the MLWE-PRF with Re-Use assumption MLWE-PRF-RU<sub>param</sub> holds, for  $\text{param} := (H, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi})$ , if for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$ , such that the following holds,*

$$\text{Adv}_{\text{param}}^{\text{MLWE-PRF-RU}}(\lambda) = |\Pr[\text{MLWE-PRF-RU}^1(\lambda) = 1] - \Pr[\text{MLWE-PRF-RU}^0(\lambda) = 1]| \leq \text{negl}(\lambda),$$

*where the experiment MLWE-PRF-RU<sup>b</sup> is defined in Figure 3,  $Q_x$  denotes the number of `SampPRF-RU` queries for each  $x$ ,  $Q := \sum_x Q_x$  denotes the total number of `SampPRF-RU` queries in the experiment,  $Q_x^\infty$  is the maximum allowed upper bound on  $\max_x \{Q_x\}$ , and  $Q_M$  denotes the total number of `SampMLWEcom` queries.*

MLWE-PRF-RU $_{H, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}}^b(\lambda)$	SampPRF-RU( $\mathbf{A}, x \in \{0, 1\}^L$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{M}[\mathbf{A}] = \perp$ , <b>then return</b> $\perp$
2 : $\mathcal{Y}, \mathcal{M} := \emptyset$ // Initialize empty tables	2 : $(\cdot, \mathbf{e}, \mathbf{R}) := \mathcal{M}[\mathbf{A}]$
3 : $b' \leftarrow \mathcal{A}^{\text{SampMLWEcom}, \text{SampPRF-RU}}(1^\lambda)$	3 : $\mathbf{e}' \xleftarrow{\$} \chi_1^h$
4 : <b>return</b> $b = b'$	4 : $\mathbf{B}_x := \mathbf{H}(x) \in R_q^{h \times m}$
<b>SampMLWEcom</b> ( $\mathbf{R} \in R^{h \times N}$ )	5 : <b>if</b> $b = 0$ <b>then</b>
1 : <b>if</b> $\ \mathbf{R}\ _\infty > \beta_r$ , <b>then return</b> $\perp$	6 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b> $\mathcal{Y}[x] \xleftarrow{\$} R_q^h$
2 : $\mathbf{A} \xleftarrow{\$} R_q^{N \times m}$ , $\mathbf{e} \xleftarrow{\$} \chi^N$	7 : $\mathbf{z} := \mathcal{Y}[x] + \mathbf{e}' - \mathbf{R}\mathbf{e} \in R_q^h$
3 : <b>if</b> $b = 0$ <b>then</b> $\mathbf{c} \xleftarrow{\$} R_q^N$	8 : <b>if</b> $b = 1$ <b>then</b>
4 : <b>if</b> $b = 1$ <b>then</b> $\mathbf{c} := \mathbf{A}\mathbf{k} + \mathbf{e} \in R_q^N$	9 : $\mathbf{z} := \mathbf{B}_x\mathbf{k} + \mathbf{e}' - \mathbf{R}\mathbf{e} \in R_q^h$
5 : $\mathcal{M}[\mathbf{A}] := (\mathbf{c}, \mathbf{e}, \mathbf{R})$	10 : $\mathcal{M}[\mathbf{A}] := \perp$
6 : <b>return</b> $(\mathbf{A}, \mathbf{c})$	11 : <b>return</b> $\mathbf{z}$

Fig. 3: MLWE-PRF-RU security game.

Now, the MLWE-PRF problem is the assumption that the randomized function  $x \mapsto \mathbf{H}(x) \cdot \mathbf{k} + \mathbf{e}_B$  (for “small” Gaussian distributed error  $\mathbf{e}_B$ ) is pseudorandom, even given some additional MLWE samples  $(\mathbf{A}, \mathbf{A}\mathbf{k} + \mathbf{e})$  with  $\mathbf{A}$  having uniformly random entries in  $R_q$  and  $\mathbf{e}$  are “small” Gaussian errors. Thus, MLWE-PRF is a mild variant of the standard PRF pseudorandomness assumption under the function  $\mathbf{H}$ , i.e.  $x \mapsto \lfloor \mathbf{H}(x) \cdot \mathbf{k} \rfloor_p$ , which has a security reduction from MLWR (and MLWE) when  $\mathbf{H}$  is instantiated as in [BLMR13] or [BP14].

**Definition 11** (MLWE-PRF $_{H, Q_P, Q_M, q, m, h, L, \vec{\chi}}$ ). *Let  $\vec{\chi} = (\chi_0, \chi_{0,B}, \chi_k)$  be discrete distributions over  $R$ , let  $\mathbf{H}$  be a hash family, and let  $Q_P, Q_M, q, m, h, L \in \mathbb{Z}^+$ . We say that the MLWE-PRF $_{\text{param}}$  assumption holds, for  $\text{param} := (H, Q_P, Q_M, q, m, h, L, \vec{\chi})$ , if for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$ , such that the following holds,*

$$\text{Adv}_{\text{param}}^{\text{MLWE-PRF}}(\lambda) = |\Pr[\text{MLWE-PRF}^1(\lambda) = 1] - \Pr[\text{MLWE-PRF}^0(\lambda) = 1]| \leq \text{negl}(\lambda),$$

where the experiment MLWE-PRF $^b$  is defined in Figure 4,  $Q_P$  denotes the number of queries to the SampPRF oracle, while  $Q_M$  denotes the number of queries to the SampMLWE oracle.

We present our reduction from MLWE-PRF to MLWE-PRF-RU via an intermediate problem, which is a “hint” variant of the MLWE-PRF-RU problem. Here, we provide the adversary with (i) MLWE hints of the form  $\mathbf{h}_i := \mathbf{R}_i \mathbf{e}_i + \mathbf{e}'_i$  for the MLWE error  $\mathbf{e}_i$ , and (ii) PRF hints of the form  $\mathbf{h}_{B,i} := \mathbf{e}_B + \mathbf{e}'_{B,i}$  for the PRF error  $\mathbf{e}_B$ . We call the latter variant as the Hint-MLWE-PRF problem. Our Hint-MLWE-PRF problem can be viewed as an extension of the “Matrix-Hint-MLWE” problem introduced in [ENP24] (with a reduction from MLWE), which in turn generalizes the Hint-MLWE problem introduced in [KLSS23]. The “Matrix” variant in [ENP24] allows for general hint matrices, rather than block diagonal hint matrices used in [KLSS23]. Extending on these, in our Hint-MLWE-PRF problem, we allow for *adaptively-chosen* general hint matrices for MLWE samples as well as hints for PRF samples that may have *non-random* matrices.

**Definition 12** (Hint-MLWE-PRF $_{H, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}}$ ). *Let  $\vec{\chi} = (\chi, \chi_B, \bar{\chi}_1, \chi_{1,B}, \chi_k)$  be discrete distributions over  $R$ , let  $\mathbf{H}$  be a hash family, and let  $Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r \in \mathbb{Z}^+$ . We say that the Hint-MLWE-PRF $_{\text{param}}$  holds, for  $\text{param} := (H, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi})$ , if for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\cdot)$ , such that the following holds,*

$$|\Pr[\text{Hint-MLWE-PRF}^1(\lambda) = 1] - \Pr[\text{Hint-MLWE-PRF}^0(\lambda) = 1]| \leq \text{negl}(\lambda),$$

MLWE-PRF $_{\mathbf{H}, Q_P, Q_M, q, m, h, L, \vec{\chi}}^b(\lambda)$	SampPRF( $x$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b>
2 : $\mathcal{Y} := \emptyset$ // Initialize empty table	2 : <b>if</b> $b = 0$ <b>then</b>
3 : $b' \leftarrow \mathcal{A}^{\text{SampPRF, SampMLWE}}(1^\lambda)$	3 : $\mathcal{Y}[x] \xleftarrow{\$} R_q^h$
4 : <b>return</b> $b = b'$	4 : <b>if</b> $b = 1$ <b>then</b>
<b>SampMLWE</b> ( $\cdot$ )	5 : $\mathbf{B}_x := \mathbf{H}(x) \in R_q^{h \times m}$
1 : $\mathbf{a} \xleftarrow{\$} R_q^m, e \xleftarrow{\$} \chi_0$	6 : $\mathbf{e}_B \xleftarrow{\$} \chi_{0,B}^h$
2 : <b>if</b> $b = 0$ , <b>then</b> $c \xleftarrow{\$} R_q$	7 : $\mathcal{Y}[x] := \mathbf{B}_x \mathbf{k} + \mathbf{e}_B \in R_q^h$
3 : <b>if</b> $b = 1$ , <b>then</b> $c := \mathbf{a}^\top \mathbf{k} + e \in R_q$	8 : <b>return</b> $\mathcal{Y}[x]$
4 : <b>return</b> $(\mathbf{a}, c)$	

Fig. 4: MLWE-PRF security game.

where the experiment  $\text{Hint-MLWE-PRF}^b$  is defined in Figure 5,  $Q_x$  denotes the number of  $\text{SampHintPRF}$  queries for each  $x$ ,  $Q := \sum_x Q_x$  denotes the total number of  $\text{SampHintPRF}$  queries in the experiment,  $Q_x^\infty$  is the maximum allowed upper bound on  $\max_x \{Q_x\}$ , and  $Q_M$  denotes the total number of  $\text{SampHintMLWE}$  queries.

Observe that in the  $\text{Hint-MLWE-PRF}$  experiment in Figure 5, once a particular  $x$  value is queried to  $\text{SampHintPRF}$  for the first time, the entries  $\mathcal{Y}[x]$  and  $\mathcal{E}[x]$  are fixed. So for any subsequent queries with the same  $x$ , the only elements that change are the masking error  $\mathbf{e}'_B$  and the corresponding hint  $\mathbf{h}_B$ . This behaviour matches with the computation of the same  $\mathbf{B}_x \mathbf{k}$  term (or  $\mathcal{Y}[x]$ ) in the  $\text{MLWE-PRF-RU}$  experiment in Figure 3.

### 3.1 Hardness Reduction from $\text{Hint-MLWE-PRF}$ to $\text{MLWE-PRF-RU}$

We start with the first step of our reduction series. The idea here is that in the  $\text{Hint-MLWE-PRF}$  problem, in addition to the  $\text{MLWE-PRF}$  samples  $(\mathbf{B}_x, \mathbf{y}_B := \mathbf{B}_x \mathbf{k} + \mathbf{e}_B)$  and  $(\mathbf{A}, \mathbf{c} := \mathbf{A} \mathbf{k} + \mathbf{e})$ , the adversary is also provided with appropriate hint vector  $\mathbf{h}_j$  about error  $\mathbf{e}$  with respect to the  $\mathbf{R}$  matrix provided by the adversary, i.e.  $\mathbf{h} := \bar{\mathbf{e}}' - \mathbf{R} \mathbf{e}$ , and also another hint  $\mathbf{h}_B := \mathbf{e}'_B - \mathbf{e}_B$ . The hints  $\mathbf{h}$  and  $\mathbf{h}_B$  then allow the adversary to simulate the “mixed”  $\text{MLWE-PRF-RU}$  samples of the form  $(\mathbf{c} := \mathbf{A} \mathbf{k} + \mathbf{e}, \mathbf{z} := \mathbf{B}_x \mathbf{k} + \hat{\mathbf{e}})$  from the given  $\text{MLWE-PRF}$  samples, since  $\mathbf{y}_B + \mathbf{h} + \mathbf{h}_B = \mathbf{B}_x \mathbf{k} + \mathbf{e}'_B + \bar{\mathbf{e}}' - \mathbf{R} \mathbf{e}$ , where  $\hat{\mathbf{e}} := \mathbf{e}'_B + \bar{\mathbf{e}}' - \mathbf{R} \mathbf{e}$  is the simulated error term for  $\mathbf{z}$ . By the discrete Gaussian convolution Lemma 2, the term  $\mathbf{e}'_B + \bar{\mathbf{e}}'$  is statistically close to a discrete Gaussian with parameter  $s_1 := \sqrt{s_{1,B}^2 + \bar{s}_1^2}$  when the independent errors  $\mathbf{e}'_B$  and  $\bar{\mathbf{e}}'$  are discrete Gaussians with parameters  $s_{1,B}$  and  $\bar{s}_1$ , respectively.

**Theorem 1.** Let  $\vec{\chi} := (\chi, \chi_B, \chi_1, \chi_k)$  and  $\vec{\chi}_h := (\chi, \chi_B, \bar{\chi}_1, \chi_{1,B}, \chi_k)$  be discrete distributions over  $R$ , let  $\mathbf{H}$  be a hash family, and let  $Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r \in \mathbb{Z}^+, \epsilon \in \mathbb{R}^+$ . The  $\text{MLWE-PRF-RU}_{\mathbf{H}, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}}$  assumption holds if the  $\text{Hint-MLWE-PRF}_{\mathbf{H}, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}_h}$  assumption holds. More precisely, for any PPT adversary  $\mathcal{A}$  against the  $\text{MLWE-PRF-RU}$  problem, there exists a PPT adversary  $\mathcal{B}$  against the  $\text{Hint-MLWE-PRF}$  problem, such that

$$\text{Adv}_{\mathcal{A}}^{\text{MLWE-PRF-RU}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{Hint-MLWE-PRF}}(\lambda) + \bar{\epsilon},$$

where  $\bar{\epsilon} = 0$  if  $\chi_1 := \chi_{1,B} + \bar{\chi}_1$  is the convolution of the distributions  $\chi_{1,B}$  and  $\bar{\chi}_1$ . Moreover, if  $\chi_1 := \mathcal{D}_{\mathbb{Z}^d, s_1}$ ,  $\chi_{1,B} := \mathcal{D}_{\mathbb{Z}^d, s_{1,B}}$  and  $\bar{\chi}_1 := \mathcal{D}_{\mathbb{Z}^d, \bar{s}_1}$  with  $s_1 := \sqrt{s_{1,B}^2 + \bar{s}_1^2}$ , and  $s_1 \geq \sqrt{\ln(2hd(1+1/\epsilon))/\pi}$ , then the above advantage bound holds with  $\bar{\epsilon} = 2Q \cdot \epsilon$ .

$\text{Hint-MLWE-PRF}_{H,Q,Q_x^\infty,Q_M,q,m,N,h,L,\beta_r,\vec{\chi}}^b(\lambda)$ <hr/> 1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$ 2 : $\mathcal{Y}, \mathcal{E}, \mathcal{M} := \emptyset$ // Initialize empty tables 3 : $b' \leftarrow \mathcal{A}^{\text{SampHintMLWE}, \text{SampHintPRF}}(1^\lambda)$ 4 : <b>return</b> $b = b'$ <hr/> $\text{SampHintMLWE}(\mathbf{R})$ // $\mathbf{R} \in R^{h \times N}$ <hr/> 1 : <b>if</b> $\ \mathbf{R}\ _\infty > \beta_r$ , <b>then return</b> $\perp$ 2 : $\mathbf{A} \xleftarrow{\$} R_q^{N \times m}$ , $\mathbf{e} \xleftarrow{\$} \chi^N$ // MLWE matrix/error 3 : $\bar{\mathbf{e}}' \xleftarrow{\$} \chi_1^h$ // MLWE hint masking term 4 : $\mathbf{h} := \bar{\mathbf{e}}' - \mathbf{R}\mathbf{e} \in R^h$ 5 : <b>if</b> $b = 0$ , <b>then</b> $\mathbf{c} \xleftarrow{\$} R_q^N$ 6 : <b>if</b> $b = 1$ , <b>then</b> $\mathbf{c} := \mathbf{A}\mathbf{k} + \mathbf{e} \in R_q^N$ 7 : <b>return</b> $(\mathbf{A}, \mathbf{c}, \mathbf{h})$	$\text{SampHintPRF}(x)$ // $x \in \{0, 1\}^L$ <hr/> 1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b> 2 : $\mathcal{E}[x] \xleftarrow{\$} \chi_B^h$ 3 : <b>if</b> $b = 0$ <b>then</b> 4 : $\mathcal{Y}[x] \xleftarrow{\$} R_q^h$ 5 : <b>if</b> $b = 1$ <b>then</b> 6 : $\mathbf{B}_x := \mathbf{H}(x) \in R_q^{h \times m}$ 7 : $\mathcal{Y}[x] := \mathbf{B}_x \mathbf{k} + \mathcal{E}[x] \in R_q^h$ 8 : $\mathbf{e}'_B \xleftarrow{\$} \chi_{1,B}^h$ // PRF hint masking 9 : $\mathbf{h}_B := \mathbf{e}'_B - \mathcal{E}[x]$ 10 : <b>return</b> $(\mathbf{B}_x, \mathcal{Y}[x], \mathbf{h}_B)$
--	---

Fig. 5: Hint-MLWE-PRF security game.

*Proof.* Given an adversary  $\mathcal{A}$  against MLWE-PRF-RU, we construct an adversary  $\mathcal{B}$  against Hint-MLWE-PRF that runs as shown in Figure 6. Note that to be able to get a meaningful response from SimPRF-RU, the adversary  $\mathcal{A}$  has to query SimMLWEcom first and use a matrix  $\mathbf{A}$  returned from SimMLWEcom to query SimPRF-RU. This is exactly the same as in the original game MLWE-PRF-RU (with the corresponding oracles SampMLWEcom and SampPRF-RU). Also, the adversary  $\mathcal{A}$  is allowed to query SimPRF-RU once per  $\mathbf{A}$ , again matching the original MLWE-PRF-RU game. Therefore, without loss of generality, we assume that  $\mathcal{A}$  queries SimPRF-RU exactly once for a matrix  $\mathbf{A}$  and always uses a matrix  $\mathbf{A}$  returned by SimMLWEcom.

$\mathcal{B}_{H,Q,Q_x^\infty,q,m,N,h,L,\beta_r,\vec{\chi}_h}^{\text{SampHintMLWE}, \text{SampHintPRF}}(1^\lambda)$ <hr/> 1 : $\mathcal{H}_a := \emptyset$ 2 : $b' \leftarrow \mathcal{A}^{\text{SimMLWEcom}, \text{SimPRF-RU}}(1^\lambda)$ 3 : <b>return</b> $b = b'$ <hr/> $\text{SimMLWEcom}(\mathbf{R})$ <hr/> 1 : $(\mathbf{A}, \mathbf{c}, \mathbf{h}) \leftarrow \text{SampHintMLWE}(\mathbf{R})$ 2 : $\mathcal{H}_a[\mathbf{A}] := \mathbf{h}$ // store MLWE hint 3 : <b>return</b> $(\mathbf{A}, \mathbf{c})$	$\text{SimPRF-RU}(\mathbf{A}, x)$ <hr/> 1 : <b>if</b> $\mathcal{H}_a[\mathbf{A}] = \perp$ , <b>then return</b> $\perp$ 2 : $(\mathbf{B}_x, \mathbf{y}_B, \mathbf{h}_B) \leftarrow \text{SampHintPRF}(x)$ 3 : $\mathbf{h} := \mathcal{H}_a[\mathbf{A}]$ // retrieve MLWE hint 4 : $\mathbf{z} := \mathbf{y}_B + \mathbf{h}_B + \mathbf{h}$ 5 : $\mathcal{H}_a[\mathbf{A}] = \perp$ 6 : <b>return</b> $\mathbf{z}$
--	--

Fig. 6: Hint-MLWE-PRF solver algorithm  $\mathcal{B}$  given an MLWE-PRF-RU solver algorithm  $\mathcal{A}$ .

Next, to analyze the advantage of  $\mathcal{B}$  against Hint-MLWE-PRF, we consider first the case that  $b = 1$ . In this case, for a query  $\mathbf{R}$  to SimMLWEcom, followed by a query  $(\mathbf{A}, x)$  to SimPRF-RU, the overall response  $(\mathbf{A}, \mathbf{c}, \mathbf{z})$  returned has  $\mathbf{c} = \mathbf{A}\mathbf{k} + \mathbf{e}$ , and  $\mathbf{z} = \mathbf{y}_B + \mathbf{h}_B + \mathbf{h}$  with  $\mathbf{y}_B = \mathbf{B}_x \mathbf{k} + \mathbf{e}_B$ ,  $\mathbf{B}_x = \mathbf{H}(x)$ ,  $\mathbf{h}_B = \mathbf{e}'_B - \mathbf{e}_B$ , and  $\mathbf{h} = \bar{\mathbf{e}}' - \mathbf{R}\mathbf{e}$ , where  $\mathbf{e}_B$  and  $\mathbf{e}'_B$  are sampled from the distributions  $\chi_B^h$  and  $\chi_{1,B}^h$ , respectively. Hence, we

have  $\mathbf{z} = \mathbf{B}_x \mathbf{k} + \hat{\mathbf{e}}$  for  $\hat{\mathbf{e}} := \mathbf{e}'_B + \bar{\mathbf{e}}' - \mathbf{R}\mathbf{e}$ . Therefore, in the Hint-MLWE-PRF<sup>1</sup> game, the view of  $\mathcal{A}$  is simulated with exactly the same distribution as in the real MLWE-PRF-RU<sup>1</sup> game in the case  $\chi_1 = \chi_{1,B} + \bar{\chi}_1$ .

In the other case that  $b = 0$ , for a similar query pair  $(\mathbf{R}, x)$  of  $\mathcal{A}$ , the overall response  $(\mathbf{A}, \mathbf{c}, \mathbf{z})$  returned has uniformly random  $\mathbf{c}$  in  $R_q^N$ , and  $\mathbf{z} = \mathbf{y}_B + \mathbf{h}_B + \mathbf{h} = (\mathcal{Y}[x] - \mathcal{E}[x]) + (\mathbf{e}'_B + \bar{\mathbf{e}}') - \mathbf{R}\mathbf{e}$ . On the other hand, in the MLWE-PRF-RU<sup>0</sup> game,  $\mathbf{c}$  is also uniform in  $R_q^N$  and  $\mathbf{z} = \mathcal{Y}^{RU}[x] + \mathbf{e}' - \mathbf{R}\mathbf{e}$ , where  $\mathcal{Y}^{RU}[x]$  denotes the  $\mathcal{Y}$  oracle in the MLWE-PRF-RU<sup>0</sup> game. Since  $\mathcal{Y}^{RU}[x]$  and  $(\mathcal{Y}[x] - \mathcal{E}[x])$  are both uniformly random in  $R_q^h$  and independent for distinct  $x$  in both games,  $(\mathbf{e}'_B + \bar{\mathbf{e}}')$  and  $\mathbf{e}'$  are both distributed as  $\chi_1^h$  in both games if  $\chi_{1,B} + \bar{\chi}_1 = \chi_1$ . Therefore, in the Hint-MLWE-PRF<sup>0</sup> game, the view of  $\mathcal{A}$  is simulated with exactly the same distribution as in the MLWE-PRF-RU<sup>0</sup> game in the case  $\chi_1 = \chi_{1,B} + \bar{\chi}_1$ .

We conclude that the advantage of  $\mathcal{B}$  against Hint-MLWE-PRF is equal to the advantage of  $\mathcal{A}$  against MLWE-PRF-RU, as claimed, in the case  $\chi_1 = \chi_{1,B} + \bar{\chi}_1$ .

The last part of the Theorem follows from the convolution Lemma 2 that implies that  $\mathcal{D}_{\mathbb{Z}^{hd}, s_{1,B}} + \mathcal{D}_{\mathbb{Z}^{hd}, \bar{s}_1}$  is within statistical distance  $2\epsilon$  of  $\mathcal{D}_{\mathbb{Z}^{hd}, s_1}$  with  $s_1 := \sqrt{s_{1,B}^2 + \bar{s}_1^2}$ , assuming the smoothing condition  $s_1 \geq \sqrt{\ln(2hd(1+1/\epsilon))/\pi}$  is satisfied. So for  $Q$  queries the statistical distance is at most  $2Q \cdot \epsilon$ , as claimed.  $\square$

### 3.2 Hardness Reduction from MLWE-PRF to Hint-MLWE-PRF

Our reduction from MLWE-PRF to Hint-MLWE-PRF can be viewed as an extension of the reduction for the Matrix-Hint-MLWE problem in [ENP24] from the MLWE problem. However, unlike the Matrix-Hint-MLWE problem in [ENP24], which only involves MLWE samples with respect to *random* matrices, our Hint-MLWE-PRF problem also provides the adversary with MLWE-like samples with respect to possibly *non-random* matrices  $\mathbf{B}_x := \mathbf{H}(x)$  for the chosen instantiation of  $\mathbf{H}$  (which is not necessarily a random oracle). That is the reason why we need the additional PRF oracle in the MLWE-PRF problem. Therefore, our reduction to Hint-MLWE-PRF is from MLWE-PRF, rather than a reduction from MLWE as in [ENP24]. Our reduction also handles *adaptive* hint-matrix queries unlike prior work [KLSS23, ENP24].

**Theorem 2.** Let  $\vec{\chi}_h := (\chi, \chi_B, \bar{\chi}_1, \chi_{1,B}, \chi_k)$  be discrete distributions over  $R$ , where  $\chi := \mathcal{D}_{\mathbb{Z}^d, s}$ ,  $\chi_B := \mathcal{D}_{\mathbb{Z}^d, s_B}$ ,  $\bar{\chi}_1 := \mathcal{D}_{\mathbb{Z}^d, \bar{s}_1}$  and  $\chi_{1,B} := \mathcal{D}_{\mathbb{Z}^d, s_{1,B}}$  for some  $s, s_B, \bar{s}_1, s_{1,B}, \eta, \delta, \epsilon, \epsilon_p \in \mathbb{R}^+$ . Let  $\mathbf{H}$  be a hash family, and let  $Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r \in \mathbb{Z}^+$ . Moreover, let  $\vec{\chi} := (\chi_0, \chi_{0,B}, \chi_k)$ , where  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  and  $\chi_{0,B} := \mathcal{D}_{\mathbb{Z}^d, \hat{\Sigma}_{0,B}}$ , where  $s_{0,B} := \sqrt{\|\hat{\Sigma}_{0,B}\|}$  and  $\hat{\rho}_{0,B} := \sigma_{\max}(\hat{\Sigma}_{0,B})/\sigma_{\min}(\hat{\Sigma}_{0,B}) = s_{0,B}^2/\sigma_{\min}(\hat{\Sigma}_{0,B})$ . The Hint-MLWE-PRF $_{\mathbf{H}, Q, Q_x^\infty, Q_M, q, m, N, h, L, \beta_r, \vec{\chi}_h}$  assumption holds if the MLWE-PRF $_{\mathbf{H}, Q, N, Q_M, q, m, h, L, \vec{\chi}}$  assumption holds and the following are satisfied:

$$\bar{s}_1/s \geq \eta \cdot \beta_r d \sqrt{hN}, \quad s_{1,B}/s_B \geq \eta \cdot \sqrt{Q_x^\infty} \quad (2)$$

$$\max(s_{0,B}/s_B, s_0/s) \leq \sqrt{\frac{1 - 1/(1 + \eta^2) - \epsilon_p}{(1 + \delta)}}, \quad (3)$$

$$s_0 \geq \sqrt{\frac{(1 + 1/\delta) \cdot \ln(2Nd(1 + 1/\epsilon))}{\pi}}, \quad s_{0,B} \geq \sqrt{\frac{(\hat{\rho}_{0,B} + 1/\delta) \cdot \ln(2hd(1 + 1/\epsilon))}{\pi}}, \quad (4)$$

$$s \geq \sqrt{\frac{\ln(2Nd + 4)}{\pi \epsilon_p}}, \quad s_B \geq \sqrt{\frac{\ln(2hd + 4)}{\pi \epsilon_p}} \quad (5)$$

More precisely, for any PPT adversary  $\mathcal{A}$  against the Hint-MLWE-PRF problem, there exists a PPT adversary  $\mathcal{B}$  against the MLWE-PRF problem, such that

$$\text{Adv}_{\mathcal{A}}^{\text{Hint-MLWE-PRF}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{MLWE-PRF}}(\lambda) + 2(Q + Q_M) \cdot \epsilon,$$

where  $Q$  and  $Q_M$  denote the total number of SampHintPRF and SampHintMLWE queries, respectively, that the adversary  $\mathcal{A}$  makes, and  $Q_x^\infty$  is the maximum number of hints requested per  $x$  query.

In the proof of the above theorem, we extend on the approach of the MLWE to Hint-MLWE reduction in [KLSS23] and in particular its generalization in [ENP24] to hints that contain linear combinations of error coordinates, as is the case in our hint matrices. Namely, the security reduction simulates the hints  $\mathbf{h}_j$  and  $\mathbf{h}_{B,j}$  with their correct marginal distribution and then simulates samples from the conditional distribution of the error vectors  $\mathbf{e}_j$  and  $\mathbf{e}_B$  respectively given the hints. Here, to deal with adaptive hint matrix  $\mathbf{R}_j$  oracle queries, we use the fact that hint  $\mathbf{h}_j$  depends only on  $\mathbf{R}_j$  and  $\mathbf{e}_j$  but not on  $\mathbf{e}_i$  for  $i \neq j$  and  $\mathbf{e}_B$ , and the independence of  $\mathbf{e}_j$ 's and  $\mathbf{e}_B$ . Hence the conditional distribution of  $\mathbf{e}_j$  given the hints depends only on  $\mathbf{h}_j$  and  $\mathbf{R}_j$  and not on  $\mathbf{h}_i$  and  $\mathbf{R}_i$  for  $i > j$ . This allows our reduction to sample the simulated conditional distribution for  $\mathbf{e}_j$  adaptively for each queried  $\mathbf{R}_j$  as a discrete Gaussian with parameters that depends only on  $\mathbf{R}_j$  and  $\mathbf{h}_j$ , as in the standard Hint-MLWE reductions [KLSS23, ENP24]. For the hints  $\mathbf{h}_{B,j}$  on the PRF error  $\mathbf{e}_B$ , the conditional distribution of  $\mathbf{e}_B$  conditioned on the hints does depend on all the  $\mathbf{e}_B$  hints  $\mathbf{h}_{B,j}$  for that  $x$  value over  $j \in [Q_x^\infty]$ . However, the latter dependence does not cause a problem with adaptive queries, because the  $\mathbf{h}_{B,j}$  hints are all with respect to a fixed identity hint matrix, and does not depend on the adaptively queried  $\mathbf{R}_j$ 's; therefore our security reduction can sample all the simulated  $\mathbf{e}_B$  hints  $\{\mathbf{h}_{B,j}\}_{j \in [Q_x^\infty]}$  for a given  $x$  value in advance when the  $x$  value is first queried to the PRF oracle. We remark that it is to handle this adaptive query issue that we reduce to a hint problem with separate hints  $\mathbf{h}_j$  and  $\mathbf{h}_{B,j}$  rather than a combined hint  $\mathbf{h}_j + \mathbf{h}_{B,j}$ , as mentioned in the technical overview. Our reduction results in a lower bound on the ratio  $s_{1,B}/s \approx \sqrt{Q_x^\infty}$  that is nearly optimal for large  $Q_x$ , as it approximately matches the minimum value of  $s_{1,B}/s$  required to thwart the known ‘‘averaging’’ attack (see Appendix B).

We are now ready to prove our security reduction from MLWE-PRF to Hint-MLWE-PRF.

*Proof (of Theorem 2).* Given an adversary  $\mathcal{A}$  against Hint-MLWE-PRF, we construct an adversary  $\mathcal{B}$  against MLWE-PRF that runs as shown in Figure 7.

**Proof of sampleability conditions.** Note that for the discrete Gaussian sampling in Line 9 of SimHintMLWE to be PPT sampleable, it is sufficient by Lemma 1 that  $\Sigma$  is positive definite, i.e.  $\sigma_{\min}(\Sigma) > 0$ , where  $\Sigma := \Sigma_0 - s_0^2 \mathbf{I}$ , and that<sup>10</sup>  $\delta_p \cdot M_\Sigma \leq 1$ , where  $\delta_p := \sqrt{\frac{\ln(2Nd+4)}{\pi}}$  and  $M_\Sigma$  denotes the max column norm of  $\sqrt{\Sigma^{-1}}$ , which is bounded as  $M_\Sigma \leq \sqrt{\sigma_{\max}(\Sigma^{-1})} \leq \frac{1}{\sqrt{\sigma_{\min}(\Sigma)}} = \frac{1}{\sqrt{\sigma_{\min}(\Sigma_0) - s_0^2}}$ . So the sufficient sampleability condition is

$$\sigma_{\min}(\Sigma_0) \geq (1 + \delta) \cdot s_0^2 + \delta_p^2, \quad (6)$$

for some  $\delta \geq 0$ . On the other hand, we have:  $\sigma_{\min}(\Sigma_0) = \frac{1}{\|\Sigma_0^{-1}\|}$  and

$$\begin{aligned} \|\Sigma_0^{-1}\| &= \|\Sigma_1^{-1} + \frac{1}{\bar{s}_1^2} \mathbf{R}^\top \mathbf{R}\| && \text{(by definition)} \\ &\leq \|\Sigma_1^{-1}\| + \frac{1}{\bar{s}_1^2} \|\mathbf{R}^\top \mathbf{R}\| && \text{(by triangle inequality)} \\ &\leq \frac{1}{s^2} + \frac{1}{\bar{s}_1^2} \|\mathbf{R}\|^2 && \text{(using } \|\mathbf{R}^\top \mathbf{R}\| = \|\mathbf{R}\|^2\text{)} \\ &\leq \frac{1}{s^2} + \frac{1}{\bar{s}_1^2} ((\beta_r d)^2 h N) && \text{(using } \|\mathbf{R}\|_\infty \leq \beta_r\text{)} \\ &\leq \frac{1}{s^2} \cdot (1 + 1/\eta^2). && (7) \end{aligned}$$

The last inequality uses the assumed lower bound (2) on  $\bar{s}_1/s$ . Therefore, using (7), we see that the sampleability condition (6) is implied by the desired lower bound  $\frac{s^2}{1+1/\eta^2} \geq (1 + \delta)s_0^2 + \delta_p^2$ , whereas the assumed condition (3) on  $s_0/s$  is equivalent to the lower bound  $\frac{s^2}{1+1/\eta^2} \geq (1 + \delta)s_0^2 + \epsilon_p s^2$ . Comparing these assumed

<sup>10</sup> We remark that there is a minor omission in the Hint-MLWE security reduction analysis in [KLSS23, Thm 1] which omitted this extra PPT sampleability condition and only required positive definiteness; it leads to the extra term  $\epsilon_p$  in the bound.

$\mathcal{B}^{\text{SampMLWE, SampPRF}}$ $H, Q, Q_x^\infty, q, m, N, h, L, \beta_r, \vec{\chi}(1^\lambda)$	$\text{SimHintPRF}(x)$
1 : $\mathcal{Y}_{\text{sim}}, \mathcal{H}_B, \mathcal{J}, \mathcal{M} := \emptyset$ // initialize empty tables 2 : $b' \leftarrow \mathcal{A}^{\text{SimHintMLWE, SimHintPRF}}(1^\lambda)$ 3 : <b>return</b> $b'$	1 : <b>if</b> $\mathcal{J}[x] = \perp$ , <b>then</b> $\mathcal{J}[x] = 0$ 2 : $\mathcal{J}[x] := \mathcal{J}[x] + 1, j := \mathcal{J}[x]$ 3 : $\mathbf{B}_x := \mathbf{H}(x)$ 4 : <b>if</b> $\mathcal{Y}_{\text{sim}}[x] = \perp$ <b>then</b> 5 : $\tilde{\mathbf{e}}_B \stackrel{\$}{\leftarrow} \chi_B^h$ // simulated error 6 : <b>for</b> $i \in [Q_x^\infty]$ <b>do</b> // set all hints for given $x$ 7 : $\mathbf{e}'_{B,i} \stackrel{\$}{\leftarrow} \chi_{1,B}^h$ 8 : $\mathbf{h}_{B,i} := \mathbf{e}'_{B,i} - \tilde{\mathbf{e}}_B$ 9 : $\mathcal{H}_B[x, i] := \mathbf{h}_{B,i}$ 10 : <b>endfor</b> 11 : $\mathbf{R}_B^\top := -[\mathbf{I}, \dots, \mathbf{I}] \in \mathbb{Z}^{hd \times Q_x^\infty \cdot hd}$ 12 : $\mathbf{h}_B^\top := [\mathbf{h}_{B,1}^\top, \dots, \mathbf{h}_{B,Q_x^\infty}^\top]$ 13 : $\mathbf{y}'_B \leftarrow \text{SampPRF}(x)$ // get PRF sample // next steps: correct the error in $\mathbf{y}'_B$ 14 : $\Sigma_{1,B} := s_B^2 \mathbf{I} \in \mathbb{Z}^{hd \times hd}$ 15 : $\Sigma_{0,B} := \left( \Sigma_{1,B}^{-1} + \frac{1}{s_{1,B}^2} \mathbf{R}_B^\top \mathbf{R}_B \right)^{-1}$ 16 : $\mathbf{d}_B := \frac{1}{s_{1,B}^2} \Sigma_{0,B} \mathbf{R}_B^\top \mathbf{h}_B$ 17 : $\mathbf{t}_B \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^{hd}, \Sigma_{0,B} - \hat{\Sigma}_{0,B}, \mathbf{d}_B}$ // correction term 18 : $\mathcal{Y}_{\text{sim}}[x] := \mathbf{y}'_B + \mathbf{t}_B \in R_q^h$ 19 : <b>return</b> $(\mathbf{B}_x, \mathcal{Y}_{\text{sim}}[x], \mathcal{H}_B[x, j])$
<b>SimHintMLWE</b> ( $\mathbf{R}$ )	
1 : <b>if</b> $\ \mathbf{R}\ _\infty > \beta_r$ , <b>then return</b> $\perp$ 2 : $\tilde{\mathbf{e}}' \stackrel{\$}{\leftarrow} \tilde{\chi}_1^h, \tilde{\mathbf{e}} \stackrel{\$}{\leftarrow} \chi^N$ // simulated errors in hint 3 : $\mathbf{h} := \tilde{\mathbf{e}}' - \mathbf{R}\tilde{\mathbf{e}}$ 4 : $(\mathbf{A}, \mathbf{c}') \leftarrow \text{SampMLWE}^N()$ // get $N$ MLWE samples // next steps: correct the error in $\mathbf{c}'$ 5 : $\Sigma_1 := s^2 \mathbf{I} \in \mathbb{Z}^{Nd \times Nd}$ 6 : $\Sigma_0 := \left( \Sigma_1^{-1} + \frac{1}{s_1^2} \mathbf{R}^\top \mathbf{R} \right)^{-1}$ 7 : $\mathbf{d}_j := \frac{1}{s_1^2} \Sigma_0 \mathbf{R}^\top \mathbf{h}$ 8 : $\Sigma := \Sigma_0 - s_0^2 \mathbf{I}$ 9 : $\mathbf{t} \stackrel{\$}{\leftarrow} \mathcal{D}_{\mathbb{Z}^{Nd}, \Sigma, \mathbf{d}}$ // correction term 10 : $\mathbf{c} := \mathbf{c}' + \mathbf{t} \in R_q^N$ 11 : <b>return</b> $(\mathbf{A}, \mathbf{c}, \mathbf{h})$	

Fig. 7: MLWE-PRF solver algorithm  $\mathcal{B}$  given a Hint-MLWE-PRF solver algorithm  $\mathcal{A}$ .

and desired lower bounds on  $\frac{s^2}{1+1/\eta^2}$ , we see that the former implies the latter if  $\epsilon_p s^2 \geq \delta_p^2$ , which in turn is equivalent to the assumed lower bound (5) on  $s$ , as required.

Similarly, for the discrete Gaussian sampling in Line 17 of  $\text{SimHintPRF}$  to be PPT sampleable, it is sufficient by Lemma 1 that  $\Sigma_B$  is positive definite, i.e.  $\sigma_{\min}(\Sigma_B) > 0$ , where  $\Sigma_B := \Sigma_{0,B} - \hat{\Sigma}_{0,B}$  and  $\delta_{pB} \cdot M_{\Sigma_B} \leq 1$ , where  $\delta_{pB} := \sqrt{\frac{\ln(2hd+4)}{\pi}}$  and  $M_{\Sigma_B}$  denotes the max column norm of  $\sqrt{\Sigma_B^{-1}}$ , which is bounded as  $M_{\Sigma_B} \leq \sqrt{\sigma_{\max}(\Sigma_B^{-1})} \leq \frac{1}{\sqrt{\sigma_{\min}(\Sigma_B)}} \leq \frac{1}{\sqrt{\sigma_{\min}(\Sigma_{0,B}) - \sigma_{\max}(\hat{\Sigma}_{0,B})}} = \frac{1}{\sqrt{\sigma_{\min}(\Sigma_{0,B}) - s_{0,B}^2}}$ . So the sufficient sampleability condition is

$$\sigma_{\min}(\Sigma_{0,B}) \geq (1 + \delta) \cdot s_{0,B}^2 + \delta_{pB}^2, \quad (8)$$

for some  $\delta \geq 0$ . On the other hand, we have:  $\sigma_{\min}(\Sigma_{0,B}) = \frac{1}{\|\Sigma_{0,B}^{-1}\|}$  and

$$\begin{aligned} \|\Sigma_{0,B}^{-1}\| &= \left\| \Sigma_{1,B}^{-1} + \frac{1}{s_{1,B}^2} \mathbf{R}_B^\top \mathbf{R}_B \right\| && \text{(by definition)} \\ &\leq \|\Sigma_{1,B}^{-1}\| + \frac{1}{s_{1,B}^2} \|\mathbf{R}_B^\top \mathbf{R}_B\| && \text{(by triangle inequality)} \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{s_B^2} + \frac{Q_x^\infty}{s_{1,B}^2} && \text{(by the fact } \mathbf{R}_B^\top \mathbf{R}_B = Q_x^\infty \mathbf{I} \text{)} \\
&\leq \frac{1}{s_B^2} \cdot (1 + 1/\eta^2). && (9)
\end{aligned}$$

The last inequality uses the assumed lower bound (2) on  $s_{1,B}/s_B$ . Therefore, using (9), we see that the sampleability condition (8) is implied by the desired lower bound  $\frac{s_B^2}{1+1/\eta^2} \geq (1+\delta)s_{0,B}^2 + \delta_{pB}^2$ , whereas the assumed condition (3) on  $s_{0,B}/s_B$  is equivalent to the lower bound  $\frac{s_B^2}{1+1/\eta^2} \geq (1+\delta)s_{0,B}^2 + \epsilon_p s_B^2$ . Comparing these assumed and desired lower bounds on  $\frac{s_B^2}{1+1/\eta^2}$ , we see that the former implies the latter if  $\epsilon_p s_B^2 \geq \delta_{pB}^2$ , which in turn is equivalent to the assumed lower bound (5) on  $s_B$ , as required.

### Analysis of $\mathcal{B}$ 's advantage

**Case 1:  $b = 1$ ; Part 1: analysis of SimHintMLWE oracle.** Next, to analyze the advantage of  $\mathcal{B}$  against MLWE-PRF, we consider first the case that  $b = 1$ . In this case, consider a query  $\mathbf{R}_j$  by  $\mathcal{A}$  to SimHintMLWE. The vector  $\mathbf{c}'_j$  returned by the SampMLWE has the form

$$\mathbf{c}'_j = \mathbf{A}_j \mathbf{k} + \mathbf{e}'_j, \quad (10)$$

where  $\mathbf{A}_j$  is uniformly random in  $R_q^{h \times N}$  and  $\mathbf{e}'_j$  is sampled from the distribution  $\mathcal{D}_{\mathbb{Z}^{hd}, s_0^2 \mathbf{I}}$ . Hence, we have

$$\mathbf{c}_j = \mathbf{A}_j \mathbf{k} + \mathbf{e}_j, \quad (11)$$

where  $\mathbf{e}_j := \mathbf{e}'_j + \mathbf{t}_j$  is sampled from the distribution  $\mathcal{D}_{\text{conv}} := \mathcal{D}_{\mathbb{Z}^{Nd}, s_0^2 \mathbf{I}} + \mathcal{D}_{\mathbb{Z}^{Nd}, \Sigma_{0,j} - s_0^2 \mathbf{I}, \mathbf{d}_j}$ . By the convolution Lemma 2, the distribution  $\mathcal{D}_{\text{conv}}$  is within statistical distance  $\leq 2\epsilon$  of  $\mathcal{D}_{\mathbb{Z}^{Nd}, \Sigma_{0,j}, \mathbf{d}_j}$ , assuming that the following smoothing condition holds:

$$\sqrt{\Sigma_{3,j}} \geq \eta_\epsilon(\mathbb{Z}^{Nd}), \text{ where } \Sigma_{3,j} := \frac{1}{s_0^2} \mathbf{I} + (\Sigma_{0,j} - s_0^2 \mathbf{I})^{-1}. \quad (12)$$

We show in Claim 1 below that (12) is satisfied. Therefore, by the ‘‘one-matrix hint’’ conditional distribution Lemma 5, for each  $j \in [Q_M]$ , the pair  $(\mathbf{c}_j, \mathbf{h}_j)$  returned by the SimHintMLWE simulator is within statistical distance  $\leq 2\epsilon$  of the distribution of the output returned by the SampHintMLWE oracle in the real Hint-MLWE-PRF<sup>1</sup> game. Note that the  $j$ 'th response hint  $\mathbf{h}_j$  provides a hint about the  $j$ 'th secret error  $\mathbf{e}_j$  only, and does not depend on any other secret errors  $\mathbf{e}_i$  and the secret errors  $\mathbf{e}_i$ 's are independent. So we can apply the ‘‘one-matrix hint’’ conditional distribution Lemma 5 inductively for each query  $j \in [Q_M]$ . Hence, we conclude that the joint distribution of the tuples  $((\mathbf{A}_j, \mathbf{c}_j, \mathbf{h}_j)_{j \in [Q_M]})$  in the view of  $\mathcal{A}$  simulated by  $\mathcal{B}$  in the MLWE-PRF<sup>1</sup> game is within statistical distance  $\leq 2Q_M \epsilon$  of the distribution of  $\mathcal{A}$ 's view in the real Hint-MLWE-PRF<sup>1</sup> game.

**Case 1:  $b = 1$ ; Part 2: analysis of SimHintPRF oracle.** Similarly, in SimHintPRF, the PRF output  $\mathbf{y}'_B$  returned by the SampPRF oracle for a query  $x$  has the form

$$\mathbf{y}'_B = \mathbf{B}_x \mathbf{k} + \mathbf{e}'_B, \quad (13)$$

where  $\mathbf{B}_x = \mathbf{H}(x)$  and  $\mathbf{e}'_B$  is sampled from the distribution  $\mathcal{D}_{\mathbb{Z}^{hd}, \hat{\Sigma}_{0,B}}$ . Hence, we have

$$\mathbf{y}_B = \mathbf{B}_x \mathbf{k} + \mathbf{e}_B, \quad (14)$$

where  $\mathbf{e}_B := \mathbf{e}'_B + \mathbf{t}_B$  is sampled from the distribution  $\mathcal{D}_{\text{convB}} := \mathcal{D}_{\mathbb{Z}^{hd}, \hat{\Sigma}_{0,B}} + \mathcal{D}_{\mathbb{Z}^{hd}, \Sigma_{0,B} - \hat{\Sigma}_{0,B}, \mathbf{d}_B}$ . By the convolution Lemma 2, the distribution  $\mathcal{D}_{\text{convB}}$  is within statistical distance  $\leq 2\epsilon$  of  $\mathcal{D}_{\mathbb{Z}^{hd}, \Sigma_{0,B}, \mathbf{d}_B}$ , assuming that the following smoothing condition holds:

$$\sqrt{\Sigma_{3,B}} \geq \eta_\epsilon(\mathbb{Z}^{hd}), \text{ where } \Sigma_{3,B}^{-1} := \hat{\Sigma}_{0,B}^{-1} + (\Sigma_{0,B} - \hat{\Sigma}_{0,B})^{-1}. \quad (15)$$

We show below in Claim 1 that (15) is satisfied. Therefore, by the matrix hint conditional distribution Lemma 5 with respect to the  $Q_x^\infty$ -hint matrix  $\mathbf{R}_B := -[\mathbf{I}, \dots, \mathbf{I}]^\top \in \mathbb{Z}^{Q_x^\infty \times hd \times hd}$  (since all  $Q_x^\infty$  hints  $\mathbf{h}_{B,j}$  for the same  $x$  are with respect to the same error  $\mathbf{e}_B$ ), we conclude that the collection of triples  $(\mathbf{B}_x, \mathbf{y}_B, \mathbf{h}_{B,j})_{j \in [Q_x]}$  returned by the SimHintPRF simulator for each queried  $x$  is within statistical distance  $\leq 2\epsilon$  of the distribution of the output returned by the SampHintPRF oracle in the real Hint-MLWE-PRF<sup>1</sup> game. Adding up over all  $\leq Q$   $x$  queries, we conclude that the joint distribution of the  $(\mathbf{B}_x, \mathbf{y}_B, \mathbf{h}_{B,j})_{j \in [Q_x]}$  tuples in the view of  $\mathcal{A}$  simulated by  $\mathcal{B}$  in the MLWE-PRF<sup>1</sup> game is within statistical distance  $\leq 2Q\epsilon$  of the distribution of  $\mathcal{A}$ 's view in the real Hint-MLWE-PRF<sup>1</sup> game.

Combining Part 1 and Part 2, the overall view of  $\mathcal{A}$  simulated by  $\mathcal{B}$  in the MLWE-PRF<sup>1</sup> game is within statistical distance  $\leq 2(Q + Q_M)\epsilon$  of the distribution of  $\mathcal{A}$ 's view in the real Hint-MLWE-PRF<sup>1</sup> game.

**Case 2:**  $b = 0$ . We now consider the second case that  $b = 0$ . In this case, the response vectors  $\mathbf{y}'_B$  and  $\mathbf{c}'_j$  are sampled from the uniform distribution on  $R_q^h$  and  $R_q^N$ , respectively. Hence,  $\mathbf{y}_B = \mathbf{y}'_B + \mathbf{t}_B$  and  $\mathbf{c}_j = \mathbf{c}'_j + \mathbf{t}_j$  are also uniformly distributed on  $R_q^h$  and  $R_q^N$ , respectively. They are independent of the hint vectors  $\mathbf{h}_j$  and  $\mathbf{h}_{B,j}$ , while the marginal distribution of the hint vectors remain the same as in the  $b = 1$  game. Therefore, the distribution of the view of  $\mathcal{B}$  simulated by  $\mathcal{A}$  in the MLWE-PRF<sup>0</sup> game is exactly equal to the distribution of  $\mathcal{B}$ 's view in the real Hint-MLWE-PRF<sup>0</sup> game.

From the above analysis for Case 1 and Case 2, we conclude that  $\text{Adv}_{\mathcal{B}}^{\text{MLWE-PRF}}(\lambda) \geq \text{Adv}_{\mathcal{A}}^{\text{Hint-MLWE-PRF}}(\lambda) - 2(Q + Q_M) \cdot \epsilon$ , which is the claimed advantage bound. What remains to prove is the smoothing conditions.

*Claim 1.* The smoothing conditions (12) and (15) are satisfied based on the assumptions in the Theorem statement.

*Proof of Claim 1:* By Lemma 4, the condition (12) holds if

$$\|\Sigma_{3,j}^{-1}\| \leq \eta_\epsilon(\mathbb{Z}^{Nd})^{-2}. \quad (16)$$

From (6), we have  $\sigma_{\min}(\Sigma_{0,j} - s_0^2 \mathbf{I}) \geq \delta \cdot s_0^2$ . Therefore, the triangle inequality gives

$$\|\Sigma_{3,j}^{-1}\| \leq \frac{1}{s_0^2} + \|(\Sigma_{0,j} - s_0^2 \mathbf{I})^{-1}\| = \frac{1}{s_0^2} + \frac{1}{\sigma_{\min}(\Sigma_{0,j} - s_0^2 \mathbf{I})} \leq \frac{1 + 1/\delta}{s_0^2}. \quad (17)$$

Combining (17) and (16), we conclude that (12) is satisfied if  $s_0 \geq \sqrt{1 + 1/\delta} \cdot \eta_\epsilon(\mathbb{Z}^{Nd})$ , and the latter condition is implied by the assumed lower bound (4) on  $s_0$ , using Lemma 3, as required.

Similarly, by Lemma 4, the condition (15) holds if

$$\|\Sigma_{3,B}^{-1}\| \leq \eta_\epsilon(\mathbb{Z}^{hd})^{-2}. \quad (18)$$

From (8), we have  $\sigma_{\min}(\Sigma_{0,B} - \hat{\Sigma}_{0,B}) \geq \delta \cdot s_{0,B}^2$ . Therefore, the triangle inequality gives

$$\|\Sigma_{3,B}^{-1}\| \leq \|\hat{\Sigma}_{0,B}^{-1}\| + \|(\Sigma_{0,B} - \hat{\Sigma}_{0,B})^{-1}\| \quad (19)$$

$$= \frac{1}{\sigma_{\min}(\hat{\Sigma}_{0,B})} + \frac{1}{\sigma_{\min}(\Sigma_{0,B} - \hat{\Sigma}_{0,B})} \quad (20)$$

$$\leq \frac{\hat{\rho}_{0,B} + 1/\delta}{s_{0,B}^2}. \quad (21)$$

Combining (21) and (18), we conclude that (15) is satisfied if  $s_{0,B} \geq \sqrt{\hat{\rho}_{0,B} + 1/\delta} \cdot \eta_\epsilon(\mathbb{Z}^{hd})$ , and the latter condition is implied by the assumed lower bound (4) on  $s_{0,B}$ , using Lemma 3, as required.  $\square$

### 3.3 Hardness Reduction from MLWE to MLWE-PRF

The final piece of our analyses on the security of MLWE-PRF-RU is to reduce from MLWE to MLWE-PRF. This reduction is contingent upon the concrete choice of the mapping  $\mathbf{H}$  and we discuss two natural options in this section.

**H is a random oracle.** It is easy to see from the description of MLWE-PRF in Figure 4 that if H is a random oracle, then SampPRF simply outputs MLWE samples with error distribution  $\chi_{0,B}$ . Assuming  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  and  $\chi_{0,B} := \mathcal{D}_{\mathbb{Z}^d, s_{0,B}}$  with  $s_0 \leq s_{0,B}$  (without loss of generality), we can construct samples of SampPRF using regular MLWE samples with error distribution  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  by adding an error term with Gaussian width  $\sqrt{s_{0,B}^2 - s_0^2}$  as similarly done before. If  $s_0 \geq s_{0,B}$ , then we can similarly construct MLWE samples with error distribution  $\chi_0$  using MLWE samples with error distribution  $\chi_{0,B}$ . We therefore get the following theorem.

**Theorem 3 (MLWE  $\rightarrow$  MLWE-PRF for random oracle H).** *Let H be a random oracle, and  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  and  $\chi_{0,B} := \mathcal{D}_{\mathbb{Z}^d, s_{0,B}}$ . Let  $\hat{s} := \min\{s_0, s_{0,B}\}$ . The MLWE-PRF $_{H, Q_P, Q_M, q, m, h, L, \chi_0, \chi_{0,B}, \chi_k}$  assumption holds if MLWE $_{Q_P + Q_M, q, m, 1, \hat{\chi}, \chi_k}$  assumption holds for  $\hat{\chi} := \mathcal{D}_{\mathbb{Z}^d, \hat{s}}$ . More precisely, for any PPT adversary  $\mathcal{A}$  against the MLWE-PRF $_{H, Q_P, Q_M, q, m, h, L, \chi_0, \chi_{0,B}, \chi_k}$  problem, there exists a PPT adversary  $\mathcal{B}$  against the MLWE $_{Q_H + Q_P + Q_M, q, m, 1, \hat{\chi}, \chi_k}$  problem with  $Q_H$  being the number of random oracle H queries, such that*

$$\text{Adv}_{\mathcal{A}}^{\text{MLWE-PRF}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda).$$

**H is defined based on BP14 PRF.** As discussed in [BP14], the BMLR13 PRF [BLMR13] can be seen as a special case of BP14 PRF. Therefore, we focus here on BP14 PRF instantiation; the results extend to BMLR13 PRF as a special case.

**Definition 13 (Notations for BP14 PRF).** *Let  $\ell := \lceil \log_2 q \rceil$ . The public parameters are two random matrices  $\mathbf{A}_0, \mathbf{A}_1 \xleftarrow{\$} R_q^{m \times m\ell}$ . The BP14 PRF computation involves a gadget matrix  $\mathbf{G}$  and its (non-linear) inverse computation  $\mathbf{G}^{-1}: R_q^{m \times m\ell} \rightarrow R_q^{m\ell \times m\ell}$  (i.e., digit decomposition).*

*Let  $x = (x_0, x_1, \dots, x_{L-1}) \in \{0, 1\}^L$ . For  $i = 1, \dots, L-1$ , define  $\mathbf{B}_{x/i} := \mathbf{G}^{-1}(\mathbf{A}_{x_i} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_{i+1}} \dots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}})))$  and  $x \setminus i := x_0 \dots x_{i-1}$  the  $i$ -bit prefix of  $x$ . Therefore,  $\mathbf{B}_{x/i} = \mathbf{G}^{-1}(\mathbf{A}_{x_i} \cdot \mathbf{B}_{x/i+1})$  and  $\mathbf{G} \cdot \mathbf{B}_{x/i} = \mathbf{A}_{x_i} \cdot \mathbf{B}_{x/i+1}$ .*

*For given  $\mathbf{A}_0, \mathbf{A}_1$  and  $x$ , let  $\chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$  be the distribution that samples  $\mathbf{e}_{x,0}, \dots, \mathbf{e}_{x,L-1} \xleftarrow{\$} \chi_0^{m\ell}$  and outputs  $\mathbf{e}_x^\top := \sum_{j=0}^{L-2} \mathbf{e}_{x,j}^\top \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$ . We write  $\mathbf{e}_{x/i}^\top := \sum_{j=i}^{L-2} \mathbf{e}_{x,j}^\top \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$ . Note that the same  $\mathbf{e}_{x,j-1} = \mathbf{e}_{x',j-1}$  is used in sampling  $\mathbf{e}_x, \mathbf{e}_{x'}$  for  $x, x'$  with the same  $j$ -bit prefix, i.e. if  $x \setminus j = x' \setminus j$ .*

*The message mapping  $H^{\text{bp}}$  for BP14 PRF is defined as follows.*

$$H_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}: x \mapsto \mathbf{B}_x^{\text{bp}}, \text{ s.t. } (\mathbf{B}_x^{\text{bp}})^\top := \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \dots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-2}} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))). \quad (22)$$

Observe that  $(H_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x) \cdot \mathbf{k})^\top = \mathbf{k}^\top \cdot \mathbf{A}_{x_0} \mathbf{B}_{x/1}$ .

We now recall the pseudorandomness result from [BP14] most relevant to us.

**Lemma 6 (MLWE  $\rightarrow$  BP14-PRF [BP14, Theorem 3.8]).** *The function family  $F_{\mathbf{A}_0, \mathbf{A}_1, \mathbf{k}}(x) := H_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x) \cdot \mathbf{k} + \mathbf{e}_x$  is pseudorandom for  $\mathbf{A}_0, \mathbf{A}_1 \xleftarrow{\$} R_q^{m \times m\ell}$ ,  $\mathbf{k} \xleftarrow{\$} R_q^m$  and  $\mathbf{e}_x \xleftarrow{\$} \chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$  if multi-secret MLWE problem (i.e., MLWE with secret matrix  $\mathbf{S}$ ) is hard where each secret key vector is sampled uniformly random in  $R_q^m$  and each error vector is sampled independently from  $\chi_0^{m\ell}$ .*

The high-level idea for the proof of Lemma 6 in [BP14] is to “peel off” each layer of computation in  $H(x)$  one by one using the observation that  $\mathbf{k}^\top \cdot \mathbf{A}_{x_i} \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/i}^\top = (\mathbf{k}^\top \cdot \mathbf{A}_{x_i} + \mathbf{e}_{x,i}^\top) \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top$ . Assuming MLWE hardness, we can replace the term inside the parentheses by a uniformly random vector, say  $\mathbf{u}_{x \setminus (i+1)}^\top$ , over  $R_q$ , that is independently sampled for each  $(i+1)$ -bit prefix  $x \setminus (i+1)$ . We can then decompose  $\mathbf{u}_{x \setminus (i+1)}^\top = \mathbf{k}_{x \setminus (i+1)}^\top \cdot \mathbf{G} + \mathbf{v}_{i+1}^\top$  for some  $\mathbf{v}_{i+1}^\top$  uniform in the quotient group  $R_q^{m\ell} / (\mathbf{G} R_q^m)$  and independent of  $\mathbf{k}_{x \setminus (i+1)}^\top$  (these  $\mathbf{k}_{x \setminus (i+1)}^\top$ , sampled independently for each  $(i+1)$ -bit prefix  $x \setminus (i+1)$ , make up the (multi-secret) MLWE secret matrix  $\mathbf{S}$ ). Plugging in  $\mathbf{u}_{x \setminus (i+1)}^\top = \mathbf{k}_{x \setminus (i+1)}^\top \cdot \mathbf{G} + \mathbf{v}_{i+1}^\top$  back to the above expression, we will get a similarly formed expression for the next index  $i+1$ . We can therefore continue to inductively peel off all layers until at the end we are left with a uniformly random vector  $\mathbf{u}_{x \setminus L}^\top = \mathbf{u}_x^\top$  which is independently sampled for each  $x \in \{0, 1\}^L$ , i.e. a uniformly random function.

From this blueprint, the MLWE  $\rightarrow$  BP14-PRF reduction naturally extends to a MLWE  $\rightarrow$  MLWE-PRF reduction as follows. In the first hybrid, while peeling off the first layer ( $i = 0$ ) of the PRF samples as

discussed above, with  $\mathbf{k}^\top \cdot \mathbf{A}_{x_0} + \mathbf{e}_{x_0}^\top$  simulated from the MLWE oracle samples, we additionally simulate the SampMLWE outputs in MLWE-PRF by requesting more samples from the MLWE oracle with respect to the same secret  $\mathbf{k}$ . So at this first layer, hardness of MLWE implies we can replace both PRF samples and SampMLWE outputs by independent uniformly random outputs. Once the additional SampMLWE outputs are replaced by uniform, they remain uniform for the remaining hybrids, and the remaining hybrids can proceed similarly as in [BP14]. This high-level idea leads to the following theorem, whose formal proof is provided in Appendix D.1.

**Theorem 4 (MLWE  $\rightarrow$  MLWE-PRF for BP14 H).**

Let  $H^{\text{bp}}$  be the BP14 function family with  $L$ -bit input and error distribution  $\chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$  as in Def. 13. Let  $\chi_k := \mathcal{U}(R_q)$ ,  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  and  $\chi_{0,B} := \chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$ . The MLWE-PRF $_{H^{\text{bp}}, Q_P, Q_M, q, m, h, L, \chi_0, \chi_0, B, \chi_k}$  assumption holds if the MLWE $_{2m\ell + Q_M, q, m, 1, \chi_0, \chi_k}$  and MLWE $_{2m\ell, q, m, Q_P, \chi_0, \chi_k}$  assumptions hold.

More precisely, for any PPT adversary  $\mathcal{A}$  against the MLWE-PRF $_{H^{\text{bp}}, Q_P, Q_M, q, m, h, L, \chi_0, \chi_0, B, \chi_k}$  problem, there exist PPT adversaries  $\mathcal{B}, \bar{\mathcal{B}}$  against the MLWE $_{2m\ell + Q_M, q, m, 1, \chi_0, \chi_k}$  and MLWE $_{2m\ell, q, m, Q_P, \chi_0, \chi_k}$  problems respectively, such that

$$\text{Adv}_{\mathcal{A}}^{\text{MLWE-PRF}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda) + (L - 1) \cdot \text{Adv}_{\bar{\mathcal{B}}}^{\text{MLWE}}(\lambda).$$

Note that for each fixed  $x$ , the PRF error distribution  $\chi_{0,B} := \chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$  in Theorem 4 marginalized for a single PRF input  $x$  is a convolution of Gaussian distributions scaled by short matrices, and hence is statistically close to a non-spherical Gaussian distribution, by standard discrete Gaussian convolution results. Our MLWE-PRF  $\rightarrow$  Hint-MLWE-PRF reduction Theorem 2 supports such non-spherical Gaussian distributions for  $\chi_{0,B}$ . However, Theorem 2 is currently not fully compatible with Theorem 4. This is because in the joint distribution in Theorem 4 of two PRF errors ( $\mathbf{e}_x^\top, \mathbf{e}_{x'}^\top$ ) for distinct inputs  $x, x'$  with a common prefix, there is a correlation between the errors (due to the error vectors sharing a common additive error term corresponding to the common prefix), while such two errors are assumed independently sampled in Theorem 2. Extending Theorem 4 to support independent error samples or Theorem 2 to support correlated error samples for distinct  $x$  is currently an open question for  $H^{\text{bp}}$  in the standard model. Still, we view Theorem 4 as a bonus qualitative evidence for the hardness of MLWE-PRF with  $H^{\text{bp}}$  in the standard model, in addition to our hardness result Theorem 3 for MLWE-PRF with a random oracle  $H$  and independent error samples. We also remark that both MLWE  $\rightarrow$  MLWE-PRF reductions in Theorems 3 and 4 are tight.

## 4 LeOPaRd : Our Lattice-Based VOPRF Proposal

At a very high-level, our goal is to realize the 2HashDH OPRF idea [JKK14] in the lattice setting as first done by Albrecht et al. [ADDS21]. Our construction ensures that a client with an input/tag pair  $(x, t) \in \{0, 1\}^L$  can recover the PRF value  $[\mathbf{B}_{x,t}\mathbf{k}]_p$  after interacting with the server, where  $\mathbf{B}_{x,t} := H(x, t)$ , for some message mapping  $H$ , and  $\mathbf{k}$  is the secret key of the server. We discuss potential instantiations of the function  $H$  and its relation to the existing lattice-based PRF schemes in Appendix C.2. As discussed before, our OPRF construction, LeOPaRd, is round-optimal (see Remark 3), and supports partial obliviousness and verifiability.

We would like the OPRF evaluation to begin with the client linearly encrypting the matrix  $\mathbf{B}_{x,t}$  such that  $\mathbf{C}_x := \mathbf{R}\mathbf{A}_r + \mathbf{B}_{x,t}$ , followed by the server computing  $\mathbf{u}_x := \mathbf{C}_x\mathbf{k} + \mathbf{e}'_s$  for some error vector  $\mathbf{e}'_s$  (so that  $\mathbf{u}_x$  contains the term  $\mathbf{B}_{x,t}\mathbf{k}$ ). However, to minimize the client's control over  $\mathbf{C}_x$  so that  $\mathbf{u}_x$  does not leak server's secret key, we want the client to commit to the pair  $\mathbf{R}$ , and generate the matrix  $\mathbf{A}_r$  via a random oracle using the resulting commitment  $\mathbf{c}_r$ .

This way, the client is forced to pick their randomness  $\mathbf{R}$  before seeing the random matrix  $\mathbf{A}_r$ . We show that the adversary can later decide adaptively on the input  $x$  without harming security. In fact, this not only suffices for reducing the pseudorandomness of LeOPaRd to our new MLWE-PRF-RU assumption (Definition 10), but it also allows us to utilize preprocessing (Figure 8) for improving the efficiency during the online phase of our protocol (Figure 9).

During the preprocessing phase, the client can sample multiple randomness  $\mathbf{R}_i$  and commit to them as  $\mathbf{c}_{r,i}$ , which the server can use to derive the corresponding  $\mathbf{A}_{r,i}$  and  $\mathbf{v}_{k,i}$  values before the online phase

<b>F.Setup(<math>1^\lambda</math>)</b> <hr/> 1 : $\forall i \in [3]: \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)$ 2 : $\forall i \in [2]: \text{crs}_i \leftarrow \text{NIZK}_i.\text{Setup}(1^\lambda)$ 3 : <b>return</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$	<b>F.KeyGen(pp)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{k} \xleftarrow{\$} \chi_{\mathbf{k}}^m$ 3 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)$ 4 : <b>return</b> $(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))$
<b>F.PreProcClient(pp, pk, T)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\$} \chi_r^{h \times (\ell+m)}$ 3 : $\forall i \in [T]: \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{R}_i; \rho_{r,i})$ 4 : $\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}$ 6 : $\text{st}_{\text{C}} := (\text{st}_{\text{pre}}, T, j := 1, \perp)$ 7 : <b>return</b> $(\text{st}_{\text{C}}, \text{prep} := \{\mathbf{c}_{r,i}\}_{i \in [T]})$ // Client includes $\text{prep}$ (received from Server) in $\text{st}_{\text{pre}}$	<b>F.PreProcServer(pp, sk, preq)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}$ 4 : $\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\forall i \in [T]: \mathbf{e}_{s,i} \xleftarrow{\$} \chi^{\ell+m}$ 6 : $\forall i \in [T]: \mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i} \in R_q^{\ell+m}$ 7 : $\text{st}_{\text{S}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 8 : <b>return</b> $(\text{st}_{\text{S}}, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})$
<b>RO<sub>r</sub>(<math>\mathbf{c}_{r,i}</math>)</b> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$	<b>RO<sub>z</sub>(<math>t, z, \mathbf{z}</math>)</b> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$

Fig. 8: Setup, preprocessing and random oracles of LeOPaRd.

starts (we note that these  $\mathbf{v}_{k,i}$  values are given to the client as well). Then, during the online phase, the client consumes the  $\mathbf{R}_j$  values one by one (where  $j$  denotes the counter that the client keeps in Figure 9). Concretely, the client uses  $\mathbf{R}_j$  and  $\mathbf{A}_{r,j}$  to derive the masked input  $\mathbf{C}_x := \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{B}_{x,t}$ , and additionally, computes a commitment to the input  $x$ , that we denote as  $\mathbf{d}_x$ <sup>11</sup>, and shares both  $\mathbf{C}_x$  and  $\mathbf{d}_x$  with the server.

Upon receiving  $\mathbf{C}_x$  and  $\mathbf{d}_x$  from the client, the server responds to the client with  $\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s$ , where  $\mathbf{e}'_s$  denote some appropriately distributed errors. At this point, the client can use its secret randomness  $\mathbf{R}_j$  to recover the PRF output  $[\mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j}]_p = [\mathbf{B}_{x,t} \mathbf{k} + \mathbf{e}_f]_p = [\mathbf{B}_{x,t} \mathbf{k}]_p$ , where  $\mathbf{e}_f := \mathbf{e}'_s - \mathbf{R}_j \mathbf{e}_s$  is some final error term with small coefficients (relative to  $q$ ).

To additionally achieve verifiability and protect against malicious adversaries, we require the client and server to prove the well-formedness of their computations with NIZK proofs.

Our full POPRF construction is given in Figures 8 and 9, and the NIZK relations of client and server proofs, i.e.,  $\pi_c$  and  $\pi_s$ , are given in Figure 10. In the protocol description,  $\text{NIZK}_1$  and  $\text{NIZK}_2$  denote the NIZK argument system of the client and server, respectively. We also summarize the main notations/parameters in Table 2. We instantiate the commitment scheme COM using Regev-style encryption [Reg05], which can also be seen as an extractable version of the BDLOP commitment [BDL<sup>+</sup>18]. Given that such an encryption is quite standard by now, we defer the details to Appendix F.

*Remark 2.* In our LeOPaRd description in Figure 9, the client's first move F.Request has the server's public key  $\text{pk}$  as an input as we follow the formal model from [TCR<sup>+</sup>22]. However, the client does not actually make use of  $\text{pk}$  in F.Request, in fact, only ever uses it in verifying  $\pi_s$ . There are important applications of

<sup>11</sup> We note that we require the commitment  $\mathbf{d}_x$  just to make the security proof of Theorem 5 go through, since during the proof we need to extract the input  $x$ .

<p><b>F.Request</b>(pp, pk, t, x, st<sub>C</sub>)</p> <hr/> <p>1 : <b>parse</b> pp := (H, {ck<sub>i</sub>}<sub>i∈[3]</sub>, {crs<sub>i</sub>}<sub>i∈[2]</sub>)</p> <p>2 : <b>parse</b> st<sub>C</sub> := (st<sub>pre</sub>, T, j, ·)</p> <p>3 : st<sub>pre</sub> := {c<sub>r,i</sub>, A<sub>r,i</sub>, R<sub>i</sub>, v<sub>k,i</sub>, ρ<sub>r,i</sub>}<sub>i∈[T]</sub></p> <p>4 : <b>if</b> j &gt; T <b>then return</b> ⊥</p> <p>5 : B<sub>x,t</sub> := H(x, t) ∈ R<sub>q</sub><sup>h×m</sup></p> <p>6 : C<sub>x</sub> := R<sub>j</sub>A<sub>r,j</sub> + B<sub>x,t</sub> ∈ R<sub>q</sub><sup>h×m</sup></p> <p>7 : d<sub>x</sub> ← COM.Commit(ck<sub>1</sub>, x; ρ<sub>x</sub>)</p> <p>8 : stmt<sub>1</sub> := (c<sub>r,j</sub>, C<sub>x</sub>, d<sub>x</sub>, ck<sub>1</sub>, ck<sub>3</sub>, A<sub>r,j</sub>, H, t, j)</p> <p>9 : wit<sub>1</sub> := (R<sub>j</sub>, x, ρ<sub>r,j</sub>, ρ<sub>x</sub>)</p> <p>10 : π<sub>c</sub> ← NIZK<sub>1</sub>.P(crs<sub>1</sub>, stmt<sub>1</sub>, wit<sub>1</sub>)</p> <p>11 : st<sub>j</sub> := (t, x, C<sub>x</sub>, pk)</p> <p>12 : st<sub>C</sub> := (st<sub>pre</sub>, T, j + 1, st<sub>j</sub>)</p> <p>13 : req := (d<sub>x</sub>, C<sub>x</sub>, π<sub>c</sub>, j)</p> <p>14 : <b>return</b> (st<sub>C</sub>, req)</p> <hr/> <p><b>F.Finalize</b>(pp, rep, st<sub>C</sub>)</p> <hr/> <p>1 : <b>parse</b> pp := (H, {ck<sub>i</sub>}<sub>i∈[3]</sub>, {crs<sub>i</sub>}<sub>i∈[2]</sub>)</p> <p>2 : <b>parse</b> st<sub>C</sub> := (st<sub>pre</sub>, T, j + 1, st<sub>j</sub>)</p> <p>3 : <b>parse</b> st<sub>pre</sub> := {c<sub>r,i</sub>, A<sub>r,i</sub>, R<sub>i</sub>, v<sub>k,i</sub>, ρ<sub>r,i</sub>}<sub>i∈[T]</sub></p> <p>4 : <b>parse</b> st<sub>j</sub> := (t, x, C<sub>x</sub>, c<sub>k</sub>)</p> <p>5 : <b>parse</b> rep := (u<sub>x</sub>, π<sub>s</sub>)</p> <p>6 : stmt<sub>2</sub> := (u<sub>x</sub>, C<sub>x</sub>, c<sub>k</sub>, ck<sub>2</sub>, v<sub>k,j</sub>, A<sub>r,j</sub>)</p> <p>7 : <b>if</b> NIZK<sub>2</sub>.V(crs<sub>2</sub>, π<sub>s</sub>, stmt<sub>2</sub>) ≠ 1 <b>then</b></p> <p>8 :   <b>return</b> ⊥</p> <p>9 : z := [u<sub>x</sub> - R<sub>j</sub>v<sub>k,j</sub>]<sub>p</sub> ∈ R<sub>p</sub><sup>h</sup>   // = [B<sub>x,t</sub>k]<sub>p</sub></p> <p>10 : y := RO<sub>z</sub>(t, x, z)</p> <p>11 : <b>return</b> y</p>	<p><b>F.BlindEval</b>(pp, pk, sk, t, req, st<sub>S</sub>)</p> <hr/> <p>1 : <b>parse</b> pp := (H, {ck<sub>i</sub>}<sub>i∈[3]</sub>, {crs<sub>i</sub>}<sub>i∈[2]</sub>)</p> <p>2 : <b>parse</b> pk := c<sub>k</sub>, sk := (k, ρ<sub>k</sub>)</p> <p>3 : <b>parse</b> st<sub>S</sub> := {c<sub>r,i</sub>, A<sub>r,i</sub>, e<sub>s,i</sub>, v<sub>k,i</sub>}<sub>i∈[T]</sub></p> <p>4 : <b>parse</b> req := (d<sub>x</sub>, C<sub>x</sub>, π<sub>c</sub>, j)</p> <p>5 : stmt<sub>1</sub> := (c<sub>r,j</sub>, C<sub>x</sub>, d<sub>x</sub>, ck<sub>1</sub>, ck<sub>3</sub>, A<sub>r,j</sub>, H, t, j)</p> <p>6 : <b>if</b> NIZK<sub>1</sub>.V(crs<sub>1</sub>, π<sub>c</sub>, stmt<sub>1</sub>) ≠ 1 <b>then</b></p> <p>7 :   <b>return</b> ⊥</p> <p>8 : e'<sub>s</sub> ←<sup>\$</sup> χ<sub>1</sub><sup>h</sup></p> <p>9 : u<sub>x</sub> := C<sub>x</sub>k + e'<sub>s</sub> ∈ R<sub>q</sub><sup>h</sup></p> <p>10 : stmt<sub>2</sub> := (u<sub>x</sub>, C<sub>x</sub>, c<sub>k</sub>, ck<sub>2</sub>, v<sub>k,j</sub>, A<sub>r,j</sub>)</p> <p>11 : wit<sub>2</sub> := (k, e<sub>s,j</sub>, e'<sub>s</sub>, ρ<sub>k</sub>)</p> <p>12 : π<sub>s</sub> ← NIZK<sub>2</sub>.P(crs<sub>2</sub>, stmt<sub>2</sub>, wit<sub>2</sub>)</p> <p>13 : <b>return</b> rep := (u<sub>x</sub>, π<sub>s</sub>)</p> <hr/> <p><b>F.Eval</b>(sk, t, x)</p> <hr/> <p>1 : <b>parse</b> sk := (k, ρ<sub>k</sub>)</p> <p>2 : B<sub>x,t</sub> := H(x, t) ∈ R<sub>q</sub><sup>h×m</sup></p> <p>3 : z := [B<sub>x,t</sub>k]<sub>p</sub> ∈ R<sub>p</sub><sup>h</sup></p> <p>4 : y := RO<sub>z</sub>(t, x, z)</p> <p>5 : <b>return</b> y</p>
--	--

Fig. 9: LeOPaRd : Our verifiable POPRF construction.

OPRFs where the client is required to only remember a password (and some common reference string, which may be embedded in application software) such as password-authenticated key exchange (PAKE) [JKX18] and password-protected secret sharing (PPSS) [JKK14, JKKX16]. LeOPaRd can support such applications in the password-only model using ideas as in [JKK14], where the authenticity of the server's public key is *not* assumed during the online OPRF evaluation (i.e., no PKI is required).

*Remark 3 (Preprocessing).* We note that unlike prior works [APRR24, HKL<sup>+</sup>25, BDFH25, DDT25] that require preprocessing in order to reduce their online rounds, our protocol given in Figure 9 is round-optimal by default (i.e., we can simply bring computation of c<sub>r</sub> and v<sub>k</sub> from preprocessing to inside F.Request and F.BlindEval, respectively). Hence, the sole reason for us to employ preprocessing is to achieve smaller online communication size (not smaller number of communication rounds).

*Remark 4 (Threshold VOPRF).* Although we presented (for the sake of simplicity) our construction from Figures 8 and 9 in a single server setting, we note that it easily extends to *n*-out-of-*n* setting, as

The server proof  $\pi_s$  proves the following,

$$\pi_s := \left\{ \begin{array}{l} \mathcal{D}(\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \mathbf{ck}_2, \mathbf{v}_k, \mathbf{A}_r, \beta_k, \beta, \beta_1); \mathcal{N}(\mathbf{k}, \mathbf{e}_s, \mathbf{e}'_s, \rho_k) \mid \|\mathbf{k}\|_\infty \leq \beta_k \wedge \\ \|\mathbf{e}_s\|_\infty \leq \beta \wedge \|\mathbf{e}'_s\|_\infty \leq \beta_1 \wedge \mathbf{c}_k = \text{COM.Commit}(\mathbf{ck}_2, \mathbf{k}; \rho_k) \wedge \\ \mathbf{v}_k = \mathbf{A}_r \mathbf{k} + \mathbf{e}_s \bmod q \wedge \mathbf{u}_x = \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s \bmod q \end{array} \right\}.$$

The client proof  $\pi_c$  proves the following,

$$\pi_c := \left\{ \begin{array}{l} \mathcal{D}(\mathbf{c}_r, \mathbf{C}_x, \mathbf{d}_x, \mathbf{ck}_1, \mathbf{ck}_3, \mathbf{A}_r, \mathbf{H}, t, j, \beta_r); \mathcal{N}(\mathbf{R}, x, \rho_r, \rho_x) \mid \|\mathbf{R}\|_\infty \leq \beta_r \wedge \\ \mathbf{B}_{x,t} := \mathbf{H}(x, t) \wedge \mathbf{c}_r = \text{COM.Commit}(\mathbf{ck}_3, \mathbf{R}; \rho_r) \wedge \\ \mathbf{d}_x = \text{COM.Commit}(\mathbf{ck}_1, x; \rho_x) \wedge \mathbf{C}_x = \mathbf{R} \mathbf{A}_r + \mathbf{B}_{x,t} \bmod q \end{array} \right\}.$$

Fig. 10: Relations for NIZK proofs performed in LeOPaRd.

Notation	Description
$\lambda$	security parameter
$\kappa$	correctness parameter
$q$	system modulus
$p$	rounding modulus
$d$	ring dimension of $R = \mathbb{Z}[X]/(X^d + 1)$
$\chi, \chi_1$	server's error distributions
$\beta, \beta_1$	$\ell_\infty$ -norm bounds on server's errors $\mathbf{e}_s, \mathbf{e}'_s$
$\chi_k$	server's key distribution
$\chi_r$	client's randomness distribution
$\beta_r$	$\ell_\infty$ -norm bound on client's randomness $\mathbf{R}$
$s_0$	error width in server's MLWE security
$m$	dimension of server key $\mathbf{k}$ (over $R_q$ )
$h$	# of rows of $\mathbf{B}_x$
$\ell$	client MLWE dimension parameter
$n_c, n_s$	dimensions of trapdoor keys used by client, server

Table 2: Summary of main notations/parameters used for LeOPaRd.

in [AG24], due to the key-homomorphic properties of our VPOPRF. Concretely, we can consider  $n$  servers with secret and public key pairs  $(\mathbf{k}_i, \mathbf{c}_{k,i})$  generated identically as in lines 2-3 of F.KeyGen in Figure 8. Analogously, during the server preprocessing, each server computes shares of  $\mathbf{v}_{k,i}$  (line 6 of F.PreProcServer in Figure 8) using their secret key shares  $\mathbf{k}_i$ , and in blind signing, each server computes the pair  $(\mathbf{u}_{x,i}, \pi_{s,i})$  identically as in lines 9-12 of F.BlindEval in Figure 9. Finally, the user replaces line 9 of F.Finalize from Figure 9 with the computation  $\mathbf{z} := \lfloor \sum_{i=1}^n \mathbf{u}_{x,i} - \mathbf{R} \sum_{i=1}^n \mathbf{v}_{k,i} \rfloor_p$  in order to obtain the final OPRF value.

*Remark 5.* We note that certain NIZK proof systems such as LNP22 [LNP22] already constructs a commitment to its witness inside the NIZK proof. In fact, as discussed in Appendix F, the trapdoor commitment scheme we use is the trapdoor version of the BDLOP commitment [BDL<sup>+</sup>18] used in [LNP22]. Therefore, by unboxing the NIZK proofs, it may be possible (in certain cases) to optimize our LeOPaRd design. However, we forego such optimizations in this work as they would make the overall protocol significantly more complex.

*Remark 6.* In our MLWE-PRF-RU security reduction from Section 3, it is important to know the maximum number of OPRF evaluation queries the adversary makes with the *same*  $\mathbf{B}_{x,t} = \mathbf{H}(x, t)$ . In the partially oblivious setting, since the server already knows  $t$ , we can let the server keep track of queried tags; and either never allow the same tag be queried twice or restrict the number of queries under the same tag, e.g., to  $2^{16}$ . This way, the adversary would be restricted to seeing a limited number of OPRF evaluation results for a particular  $\mathbf{B}_{x,t}$ . In this case, the impact of the term  $Q_x^\infty$  in Theorem 2 would effectively diminish.

Some applications such as Privacy Pass [DGS<sup>+</sup>18] can benefit from batching multiple queries in one go. It is not difficult to see that LeOPaRd can support this via the client/server consuming more pre-processing outputs at once. In the online phase, both the client and the server can compute a single NIZK proof attesting

to the validity of the whole batched query. When using a succinct argument system like LaBRADOR [BS23], the proof size will increase by a very small factor. For example, as discussed in [ADDG24, App. A.2], their LaBRADOR proof size increases from 45KB (for one single query) to just 79KB for a batch of 64 queries ( $< 1.8\times$  increase). Therefore, the per-query communication cost of the NIZK proof drops down to a very small value. In Appendix C, we discuss options for concretely instantiating  $\mathsf{H}$  (including BP14 PRF [BP14], BLMR13 PRF [BLMR13] and random oracle) and the NIZK proofs.

#### 4.1 Correctness and Security Analyses

**Correctness analysis.** We upper bound the correctness error probability of our protocol in Lemma 7, in terms of the protocol parameters. Due to space limitations, we only sketch the proofs and refer the reader to Appendix D.2 for full proofs and details.

**Lemma 7.** *Fix  $\kappa, h, d$ . Let  $B_f(\kappa, d)$  denote an upper bound on a fixed coordinate of  $\mathbf{e}_f = \mathbf{e}'_s - \mathbf{R}\mathbf{e}_s$  (as in (49)) that holds except with probability  $p_e \leq 2^{-(\kappa+2+\log(hd))}$ . Also, assume that the function family  $\mathsf{H}$  satisfies  $\epsilon_u$ -uniformity (as per Definition 19) with*

$$\epsilon_u \leq 2^{-(\kappa+2+\log(hd))}. \quad (23)$$

Then, for any fixed  $x \in \{0, 1\}^L$ ,  $2^{-\kappa}$ -correctness holds (as per Definition 2) if

$$q/p \geq 2^{\kappa+2} \cdot hd \cdot (2B_f(\kappa, d) + 1). \quad (24)$$

**Remarks:**

- **(worst-case bound)** Note that  $B_f(\kappa, d)$  is also upper bounded by the worst-case bound ( $p_e = 0$ ):

$$B_f(\kappa, d) \leq (m + \ell)d\beta\beta_r + \beta_1 \quad (25)$$

- **(bound for Gaussian distributions)** Assume that  $\chi_r = \mathcal{U}(\mathbb{S}_{\beta_r})$ ,  $\chi = \mathcal{D}_s$ , and  $\chi_1 = \mathcal{D}_{s_1}$ . Then, using the Gaussian convolution lemma (Lemma 2) for the term  $\mathbf{R}\mathbf{e}_s$  and Gaussian tail bound, we get

$$B_f(\kappa, d) \leq \sqrt{\ln(2)(\kappa + 2 + \log(hd))/\pi} \cdot \left( s_1 + \beta_r s \sqrt{(\ell + m)d} \right). \quad (26)$$

- **(bound for uniform distributions)** In practice, one may want to avoid using Gaussian distribution in the implementation. We provide the bounds for using only uniform distribution as well. Assume that  $\chi_r = \mathcal{U}(\mathbb{S}_{\beta_r})$ ,  $\chi = \mathcal{U}(\mathbb{S}_{\beta})$ , and  $\chi_1 = \mathcal{U}(\mathbb{S}_{\beta_1})$ . Using the central limit theorem Gaussian approximation for the distribution of the coordinates of the term  $\mathbf{R}\mathbf{e}_s$  in (49), their standard deviation is given by  $\sigma\sigma_r\sqrt{(m + \ell)d}$ , where  $\sigma := \sqrt{\frac{1}{12}((2\beta + 1)^2 - 1)}$  and  $\sigma_r := \sqrt{\frac{1}{12}((2\beta_r + 1)^2 - 1)}$  are the standard deviations of the  $\mathbf{R}$  and  $\mathbf{e}_s$  coordinates, respectively. Using a Gaussian tail bound and the upper bound  $\beta_1$  for the coordinates of  $\mathbf{e}'_s$ , we have the approximate upper bound

$$B_f(\kappa, d) \leq \beta_1 + \sigma\sigma_r\sqrt{(m + \ell)d \cdot 2\ln(2)(\kappa + 2 + \log(hd))}. \quad (27)$$

- The existing correctness definition in Definition 2 assumes only 1 key generation run and 1 OPRF protocol evaluation. For a modified correctness definition with  $Q_{\text{eval}}$  total key generation and protocol evaluation pairs, the bound in (24) will be multiplied by  $Q_{\text{eval}}$ .

In our evaluation protocol, the client obtains  $\mathbf{u}_x - \mathbf{R}\mathbf{v}_k = \mathbf{B}_{x,t}\mathbf{k} + \mathbf{e}_f \in R_q^h$ , where  $\mathbf{e}_f = \mathbf{e}'_s - \mathbf{R}\mathbf{e}_s$  is an error term due to the server and client’s randomness, and then rounds the result to the nearest multiple of  $q/p$ . The correctness proof bounds the probability of a PRF rounding error  $\Pr[\lfloor \mathbf{B}_{x,t}\mathbf{k} + \mathbf{e}_f \rfloor_p \neq \lfloor \mathbf{B}_{x,t}\mathbf{k} \rfloor_p]$ , which occurs only if  $\mathbf{B}_{x,t}\mathbf{k}$  falls “close” to a rounding interval boundary (midpoints between consecutive multiples of  $q/p$ ), i.e. within distance  $B_f$  of an upper bound on the coordinates of error term  $\mathbf{e}_f$ . For this, we exploit the uniformity, up to negligible statistical distance  $\epsilon_u$ , of coordinates of  $\mathbf{B}_{x,t}\mathbf{k}$  (a property of  $\mathsf{H}$  which we call  $\epsilon_u$ -uniformity) and choose  $q/p$  sufficiently large by a factor  $\approx 2^\kappa$  compared to  $B_f$  to achieve  $2^{-\kappa}$  correctness error. We upper bound the  $\epsilon_u$ -uniformity of the BLMR13 and BP14 PRF instantiations of  $\mathsf{H}$  for this purpose.

**Security analysis.** We first prove pseudorandomness against malicious clients with Theorem 5 and then request privacy against malicious servers (POPRIV2) with Theorem 6. Due to space constraints, we give only the sketch of the main proof steps, and refer the reader to Appendix D.3 and Appendix D.4 for the full proofs, respectively. In Appendix E.2, we discuss an issue in a proof of [TCR<sup>+</sup>22], where they prove that correctness and POPRIV2 together implies uniqueness (a stronger form of verifiability). We show that this implication also requires a *key binding* property, and for completeness, we provide a full proof of the statement that correctness, POPRIV2 and key binding implies uniqueness. Our LeOPaRd proposal satisfies all these properties. Although our correctness error is not necessarily less than  $2^{-128}$ , the verifiability property is not affected by the correctness error. We may only have a case where a different PRF output is computed in Finalize with probability  $2^{-\kappa}$ .

**Theorem 5.** *The POPRF construction F from Figures 8 and 9 satisfies pseudorandomness given in Definition 3, with random oracles  $\text{RO}_r$  and  $\text{RO}_z$ , if:*

- The client argument system  $\text{NIZK}_1$  is computationally sound and the server argument system  $\text{NIZK}_2$  is computationally zero-knowledge (Definition 18),
- The commitment scheme COM is computationally hiding and extractable (Definitions 15 and 17), and
- The MLWE-PRF-RU<sub>param</sub> assumption (Definition 10) holds for  $\text{param} := (\text{H}, Q, Q_{x,t}^\infty, Q_M, q, m, \ell + m, h, L, \beta_r, \vec{\chi})$  and  $\vec{\chi} := (\chi, \chi_1, \chi_k)$ , where  $Q$  corresponds to the number of BlindEval queries,  $Q_{x,t}^\infty$  to the maximum number of identical  $(x||t)$  queries to BlindEval, and  $Q_M$  to the number of  $\text{RO}_r$  queries, respectively.

More precisely, for any PPT adversary  $\mathcal{A}$ , there exist PPT adversaries  $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$  and  $\mathcal{B}_6$  against the computational zero-knowledge of  $\text{NIZK}_2$ , computational soundness of  $\text{NIZK}_1$ , hiding and extractability of COM, and MLWE-PRF-RU<sub>param</sub> assumption, respectively, such that

$$\begin{aligned} \text{Adv}_{\mathcal{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^{\text{po-prf}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_2, \mathcal{B}_1}^{\text{CZK}}(\lambda) + Q_B \cdot \text{Adv}_{\text{NIZK}_1, \mathcal{B}_2}^{\text{CS}}(\lambda) + \text{Adv}_{\text{COM}, \mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + Q_{\text{RO}_r} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_4}^{\text{Ext}}(\lambda) + Q_B \cdot \text{Adv}_{\text{COM}, \mathcal{B}_5}^{\text{Ext}}(\lambda) + \text{Adv}_{\mathcal{B}_6}^{\text{iMLWE}_{\text{param}}}(\lambda) \\ &\quad + Q_{\text{RO}_z} \left( \frac{1}{p} + \frac{1}{q} \right)^{dh}, \end{aligned}$$

where  $Q_B$ ,  $Q_{\text{RO}_r}$  and  $Q_{\text{RO}_z}$  denote the number of BlindEval,  $\text{RO}_r$  and  $\text{RO}_z$  queries, respectively, that the adversary  $\mathcal{A}$  makes.

Our proof essentially revolves around the idea of removing traces of the server’s PRF key  $\mathbf{k}$  from the computation of the POPRF output, such that the simulator  $\mathcal{S}$  of the pseudorandomness game does not need any secret key dependent information. In order to do so, we first simulate the server proof  $\pi_s$ , such that from that point onward we do not need  $\mathbf{k}$  for the generation of the proof. Next, we commit to an all zero vector  $\mathbf{0}_m$  during the computation of the public key  $\text{pk}$ , which allows us to remove  $\mathbf{k}$  from the computation of  $\text{pk}$ . Lastly, we rely on our MLWE-PRF-RU assumption to remove all traces of  $\mathbf{k}$  from the computation of  $\mathbf{u}_x$  and  $\mathbf{v}_k$ , and yet maintain the correctness of the scheme. Before relying on the MLWE-PRF-RU assumption though, we use the extractability property of the commitment scheme to extract the client’s input  $x$  and randomness  $\mathbf{R}$ , which we need for making the appropriate calls to the oracles of the MLWE-PRF-RU assumption. After applying all these changes, we end up in a position where we do not require the secret key  $\mathbf{k}$  anymore, and hence, the simulator  $\mathcal{S}$  can be constructed easily. We note that to make these argument go through, we also need to carefully program the random oracles  $\text{RO}_r$  and  $\text{RO}_z$  during the proof. Furthermore, our security proof takes into account the possibility of the adversary making many queries to various oracles, which are captured via the terms like  $Q_B$ ,  $Q_{\text{RO}_r}$ , etc., and we ensure the hardness of the underlying problems under (polynomially-bounded) many queries.

**Theorem 6.** *The POPRF construction F from Figures 8 and 9 satisfies the request privacy against malicious servers (POPRIV2), given in Definition 4, with random oracles  $\text{RO}_r$  and  $\text{RO}_z$ , if:*

- The client argument system  $\text{NIZK}_1$  is computationally zero-knowledge and the server argument system  $\text{NIZK}_2$  is computationally sound (Definition 18),

- The commitment scheme COM is computationally hiding and extractable (Definitions 15 and 17), and
- The  $\text{knMLWE}_{\ell+m,m,h,\chi}$  assumption holds (Definition 7).

More precisely, for any PPT adversary  $\mathcal{A}$ , there exist PPT adversaries  $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$  and  $\mathcal{B}_6$  against the computational zero-knowledge of  $\text{NIZK}_1$ , computational hiding of COM, computational soundness of  $\text{NIZK}_2$ , extractability of COM and  $\text{knMLWE}_{m,\ell+m,h,\chi}$  assumption, respectively, such that

$$\begin{aligned} \text{Adv}_{\mathcal{F}, \mathcal{A}, \text{RO}_r, \text{RO}_z}^{\text{po-priv-2}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_1, \mathcal{B}_1}^{\text{CZK}}(\lambda) + T \cdot Q_P \cdot \text{Adv}_{\text{COM}, \mathcal{B}_2}^{\text{CH}}(\lambda) + 2Q_R \cdot \text{Adv}_{\text{COM}, \mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + 2Q_F \cdot \text{Adv}_{\text{NIZK}_2, \mathcal{B}_4}^{\text{CS}}(\lambda) + Q_R^{\text{pk}} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_5}^{\text{Ext}}(\lambda) + 2Q_R \cdot \text{Adv}_{\mathcal{B}_6}^{\text{knMLWE}}(\lambda), \end{aligned}$$

where  $Q_P, Q_R$  and  $Q_F$  denote the number of PreProcClient, Request and Finalize oracle queries, respectively, and  $Q_R^{\text{pk}}$  denotes the number of queries to the Request oracle with distinct  $\text{pk}$  inputs that the adversary  $\mathcal{A}$  makes.

In order to prove request privacy against malicious servers, we aim to remove traces of the client’s input  $x$ , such that at the end the transcript observed by the adversary is independent of the challenge bit  $b$ . To this end, we first simulate the client proof  $\pi_c$ , such that we do not require anymore the input  $x$  for computing the proof. Next, we change the client commitments  $\mathbf{c}_r$  and  $\mathbf{d}_x$  to be commitments of all zero values  $\mathbf{0}_{h \cdot (\ell+m)}$  and  $0^{|x|}$  instead of  $\mathbf{R} x$ , respectively, which removes another occurrence of the input  $x$ . Finally, we rely on the Knapsack MLWE (Definition 7) assumption to replace computation of  $\mathbf{C}_x$  with a uniformly random matrix. We note that in order to maintain the correctness of the scheme, during this last change we rely on the extractability of the commitment scheme. More precisely, we extract the secret key  $\mathbf{k}$  from the public key  $\text{pk}$ , and use  $\mathbf{k}$  to correct the POPRF evaluation inside Finalize. Importantly, we can use Regev-style [Reg05] *lightweight* commitments, as opposed to requiring *heavy* full trapdoors to extract from a commitment of the form  $\mathbf{A}\mathbf{k} + \mathbf{e}$  used in [ADDS21].

At this point we have removed all occurrences of the input  $x$ , and hence, the adversary’s view is independent of the input bit  $b$ .

## 5 Practical Performance Analysis

Our performance analysis focuses on estimating the sizes of various components of LeOPaRd. Once protocol components are small enough, a practically-acceptable computational performance is often also achieved given that lattice-based schemes involve quite simple operations like matrix-vector multiplications (see examples in [ESLL19, EZS<sup>+</sup>19, LNS20, NS24]).

One of the most important factors impacting the performance of LeOPaRd is the size of the modulus  $q$ . Since the dimension parameters for suitable MLWE security grow linearly with  $\log q$ , the overall communication is in fact *quadratic* in  $\log q$ . Particularly, the sizes of commitments/encryptions (i.e.,  $\mathbf{c}_r, \mathbf{C}_x$ ) on the client side and  $\mathbf{v}_k$  on the server side for LeOPaRd) quickly become large for very large  $q$ . This is not just specific to our scheme, but true in general for lattice-based commitment/encryption schemes (as also used e.g. in [ADDS21]). Therefore, our first goal is to minimize  $q$  as much as possible.

As shown in the correctness analysis (Lemma 7), the size of  $q$  itself is heavily dominated by the  $2^\kappa$  term. Therefore, given this bound is *statistical* and for correctness (rather than a security property), we consider a range of  $\kappa$  values between 16 and 64. One may wonder what if a (malicious) server wants to stop the client from getting the correct OPRF result. There are two parts to consider here. First, as discussed in Remark 8, our correctness analysis easily extends to cover such a malicious server setting as the NIZK proof  $\pi_s$  proves shortness of the server’s error terms, and also we can simply increase  $\kappa$ . Second, OPRFs are often used in the server-client setting, and if clients do not receive the correct PRF output and therefore end up not receiving the intended service from the system, then the server would harm its own reputation and lose clients (while still not able to break client’s privacy). Hence, we believe a rational server would not intentionally try to stop the client from getting the correct OPRF result.

For performance analysis, we focus on the instantiation of the message mapping obtained by truncating  $\text{H}^{\text{bp}}$  of the BP14 PRF (see Appendix C.2) since this choice leads to a better performance overall. As discussed

in Appendix C, one may consider a general base  $\gamma$  for the gadget matrix  $\mathbf{G}$ , i.e.,  $\mathbf{g} = (1, \gamma, \gamma^2, \dots)$ , and  $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}$ .

We set  $\chi = \mathcal{D}_s$ ,  $\chi_1 = \mathcal{D}_{s_1}$  (matching our MLWE-PRF-RU analyses in Section 3) for width parameters  $s, s_1$ . Since  $\chi_r$  (client’s randomness distribution) is not involved in MLWE-PRF-RU analyses at all and the client’s randomness  $\mathbf{R}$  is one-time, we simply set  $\chi_r = \mathcal{U}(\mathbb{S}_{\beta_r})$  for  $\beta_r = 1$ . Since BP14 PRF allows for a small secret key, we sample the server’s key  $\mathbf{k}$  from the same  $\chi$  distribution, i.e.,  $\chi_k = \chi = \mathcal{D}_s$ . To estimate the practical hardness of MLWE-PRF-RU problem, we rely on the reductions in Section 3, particularly the conditions in Theorem 2. Based on this theorem, we require the hardness of MLWE with the same secret dimension  $m$  as in MLWE-PRF-RU and error Gaussian parameter  $s_0$ .

Our parameter setting then proceeds as follows. Fix a value for  $\kappa \in \{16, 32, 64\}$ . We iterate over different values of  $d \in \{64, 128, 256\}$  that optimize the communication efficiency for both the server and the client. Recall that for BP14 PRF, we have  $h = 1$ . Therefore, for a fixed  $d$ , we also set  $p$  as the smallest integer satisfying  $(1/p+1/q)^{dh} \leq 2^{-128}$  (to satisfy a requirement from the pseudorandomness analysis in Theorem 5). Here, the term  $1/q$  is much smaller (since  $q$  is always larger than  $2^{32}$ ) and therefore does not really play an important role. Then, based on Theorem 2, we first fix  $s_B = s$ ,  $s_{0,B} = s_0$  (choosing  $\chi_{0,B}$  as spherical Gaussian),  $\delta = \eta = 1$ ,  $Q = 2^{64}$  and  $\epsilon = 2^{-128}$ , and compute the smallest  $(s_0, s, \bar{s}_1, s_{1,B})$  values. The parameter  $s_1 = \sqrt{s_{1,B}^2 + \bar{s}_1^2}$  is set as in Theorem 1. Here we set  $Q_x^\infty = 2^\kappa$  for the fully oblivious setting, and  $Q_x^\infty = 2^{16}$  for the partially oblivious setting. We explain the reasoning behind this towards the end of this section. In this procedure, we have  $N = m + \ell$  and we estimate the values for  $m$  and  $\ell$  initially, and correct the estimate iteratively for accuracy.

Now we consider the size of  $q$  based on the correctness analysis (Lemma 7) and set  $\log q$  as small as possible. Then, we set the smallest value for  $m$  so that MLWE security with error width  $s_0$  at 128 bits is achieved (server security). As discussed above, this comes from the security reductions in Section 3. Similarly, we set the smallest value for  $\ell$  so that  $\text{MLWE}_{\ell, \chi_r}$  security at 128 bits is achieved (client security). We measure the practical hardness of MLWE using the lattice estimator [APS15] and aim for a “root Hermite factor” (RHF) of around 1.0045 as in earlier works [ESLR23, ESZ22b, LNP22, ESLL19]. RHF is a commonly used measure to estimate the quality of lattice reduction to solve a particular lattice problem. Earlier works such as [ESLR23, ESZ22b, LNP22, ESLL19] considered the same RHF value for 128-bit security level. The same MLWE dimension parameters  $w_c = \ell$  (by client) and  $w_s = m$  (by server) can be used to establish the hiding property of commitment scheme used by the client and the server as they have the same  $\log q$  and their randomnesses are sampled from the same distributions as the earlier MLWE error terms. We also note that given many MLWE and MLWE-PRF-RU samples may be leaked due to many OPRF evaluation queries, we also need to take into account combinatorial attacks (e.g., using Gröbner basis) against MLWE. We verified using the lattice estimator [APS15] that the combinatorial attack complexity is always above  $2^{128}$  for our parameter settings.

Once the majority of the parameters are selected as above, what remains is to consider (i) the binding property of the commitment scheme, (ii) the indistinguishability of a commitment key with a trapdoor from a regular commitment key, and (iii) decryption/extraction correctness for the commitments. First we note that the  $\mathbf{d}_x$  commitment simply commits to a small input  $x$  with entropy at most 256 bits (probably even less). Therefore, we can simply use a regular lattice-based encryption for  $\mathbf{d}_x$ , which we assume to cost 1 KB. Now for the other commitments, given that we are dealing with (relatively) large moduli, we see that the third requirement above is easily met. Again due to (relatively) large moduli and that the MSIS solution norm bounds are quite small (due to tight norm-bound NIZK proofs), the second requirement turns out to be more dominant than the first one in setting the values for  $(n_s, n_c)$  (the  $n$  parameters from Appendix F used by the server and client, respectively). As discussed in Appendix F, we require the same MLWE assumptions over  $R_q$  with error width  $s_0$  (for server’s commitment) and errors sampled from  $\chi_r$  (for client’s commitment). As a result, we get that  $n_s = m$  and  $n_c = \ell$ . Note that the dimensions of the randomnesses used inside the commitments by the server and the client do not have a significant impact in our parameter setting and efficiency estimates since the well-formedness of the commitments are proven by LaBRADOR, and therefore, this proof cost gets amortized along with everything else (see below for more discussion on the proof costs). For the MSIS hardness of our parameter settings, it turns out that module ranks (from both the client and

the server’s side) as small as 1-2 is already sufficient. Therefore, we can comfortably use the standard low-order bit-dropping technique from [DKL<sup>+</sup>18]. For simplicity, we assume  $D = 12$  bits can be dropped in the ‘binding’ parts of the commitments  $\mathbf{c}_r$  and  $\mathbf{c}_k$  (i.e., top  $n_c$  and  $n_s$  rows of  $\mathbf{c}_r$  and  $\mathbf{c}_k$ , respectively). A similar  $D$  parameter has been used in earlier works with smaller moduli such as [ESLR23, LNP22]. We believe that a larger  $D$  can be used, but the saving is not very significant, and therefore, we opt to not pursue further analysis here so as not to over-complicate the discussions. This concludes the setting of all parameters for LeOPaRd except the underlying NIZK proofs.

In Table 3, we present some example parameter settings and communication sizes aiming at 128-bit security level. The communication of the server’s public key is a *one-time* offline cost. Here, we consider two cases: (i) fully oblivious (i.e. no tag) setting where  $Q_x^\infty = 2^\kappa$ ; and (ii) partially oblivious setting where  $Q_{x,t}^\infty = 2^{16}$  (both corresponding to  $Q_x^\infty$  in Theorem 2). In the former case, we are restricting the total number of OPRF queries to  $2^\kappa$ , matching also the correctness error. In the latter case, we consider that the server has the ability to see a part of the PRF input, i.e. tag  $t$ , and therefore, can limit the number of OPRF queries under the same tag to be at most  $2^{16}$ . If the tag corresponds to a user identifier as in [ECS<sup>+</sup>15], then this would mean  $2^{16}$  queries per identifier, which we believe is quite reasonable (for honest users) in practice. We also note that our parameter setting aiming at 90-95 bits of security (to compare with [AG24] in Table 1) is provided in Table 4.

Scheme	$\kappa$	$m = n_s$	$\ell = n_c$	$s_0$	$s$	$s_1$	$\log q$	Server PK size	Server Offline (per query)	Server Online (w/o NIZK)	Client Offline (per query)	Client Online (w/o NIZK)
LeOPaRd	16	24	27	9.90	21.5	11262	42	13.50	16.73	0.33	23.39	8.88
Partial	32	34	37	9.90	21.6	12866	59	28.16	32.73	0.46	46.77	16.67
OPRF	64	54	56	9.93	21.6	15535	92	72.56	79.06	0.72	114.78	39.81
LeOPaRd	16	24	27	9.90	21.5	11262	42	13.50	16.73	0.33	23.39	8.88
(plain)	32	38	41	10.25	23.5	$\approx 2^{21}$	66	35.62	40.73	0.52	58.55	20.59
OPRF	64	67	70	10.93	28.5	$\approx 2^{37}$	114	113.06	122.02	0.89	178.69	60.67

Table 3: Summary of our parameters aiming around 128 bit security and communication cost results. All communication and PK sizes are in KB. For all settings, we have  $d = 64$ ,  $p = 4$ ,  $h = 1$ , and  $\beta_r = 1$ .

Now we get to estimating the sizes of LaBRADOR-based [BS23] NIZKs (in combination with LNP22 [LNP22] to achieve zero-knowledge). One can observe that its proof size is barely impacted by the witness size (see, e.g., [BS23, Fig. 1]) thanks to its recursive nature. In [ADDG24], the authors report that their LaBRADOR-based well-formedness NIZK proof (including the LNP22 part) under a 75-bit modulus for FHE ciphertexts is about 45 KB for a single query. We believe the proofs needed for LeOPaRd more naturally fit LaBRADOR because, unlike [ADDG24], we do not need to use a different *proof* modulus than the one used in the core OPRF protocol. Hence, we estimate our proof cost to be about 45 KB (for a single query). Our estimate is also inline with similar LaBRADOR-based proof estimates in [AG24]. In fact, our estimates may even be more conservative than [AG24] because their modulus  $q$  is much larger around  $2^{143}$ - $2^{169}$ . In Table 3, we exclude the costs of NIZKs to provide a clearer presentation of the advantage of our core techniques (that are outside of the NIZKs). From here, we can see that reducing the costs of underlying NIZKs is a promising step to improve the efficiency of post-quantum OPRFs, which we leave as a future work.

## Acknowledgements

We want to thank the anonymous reviewers for their very helpful comments and suggestions. We want to thank Martin Albrecht for helpful discussions and for answering our questions regarding the prior OPRF works. This work was (partially) supported by an Amazon Research Award Fall 2023, the France 2030 ANR Project ANR-22-PECY-003 SecureCompute, by the Austrian Science Fund (FWF) Project J4879-N, by the Bakar Funds and Peder Sather Funds, and Australian Research Council Discovery Grants DP180102199 and

Scheme	$\kappa$	$m = n_s$	$\ell = n_c$	$s_0$	$s$	$s_1$	$\log q$	Server PK size	Server Offline (per query)	Server Online (w/o NIZK)	Client Offline (per query)	Client Online (w/o NIZK)
LeOPaRd	16	18	20	9.88	21.4	10088	42	10.12	12.47	0.33	17.48	6.91
Partial	32	25	27	9.90	21.5	11347	59	20.70	23.97	0.46	34.34	12.52
OPRF	64	40	42	9.93	21.6	13670	91	53.12	58.30	0.71	84.93	29.44
LeOPaRd	16	18	20	9.88	21.4	10088	42	10.12	12.47	0.33	17.48	6.91
(plain)	32	28	30	10.25	23.5	$\approx 2^{21}$	66	26.25	29.91	0.52	43.08	15.44
OPRF	64	50	52	10.93	28.4	$\approx 2^{37}$	114	84.38	90.84	0.89	133.17	45.53
[AG24]		$Q = 2^{16}$ ;	95-bit security				143	108.30	N/A	1.80	N/A	71.50
(plain)		$Q = 2^{32}$ ;	90-bit security				151	113.50	N/A	1.80	N/A	75.50
OPRF		$Q = 2^{64}$ ;	167-bit security				169	221.50	N/A	1.80	N/A	169.00

Table 4: Summary of our parameters aiming around 90-95 bit security and communication cost results. All communication and PK sizes are in KB. For all settings, we have  $d = 64$ ,  $p = 4$ ,  $h = 1$ , and  $\beta_r = 1$ . We also provide the results for [AG24] as the most closely-related work.  $Q$  denotes a global upperbound on the number of evaluation queries (as in Table 1).

DP220101234. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of Amazon or other funding agencies.

## References

- ADDG24. Martin R. Albrecht, Alex Davidson, Amit Deo, and Daniel Gardham. Crypto dark matter on the torus - oblivious PRFs from shallow PRFs and TFHE. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 447–476. Springer, Cham, May 2024. [2](#), [3](#), [8](#), [26](#), [30](#), [65](#)
- ADDS21. Martin R. Albrecht, Alex Davidson, Amit Deo, and Nigel P. Smart. Round-optimal verifiable oblivious pseudorandom functions from ideal lattices. In Juan Garay, editor, *PKC 2021, Part II*, volume 12711 of *LNCS*, pages 261–289. Springer, Cham, May 2021. [2](#), [3](#), [4](#), [5](#), [7](#), [22](#), [28](#), [37](#), [39](#), [65](#), [66](#)
- AG24. Martin R. Albrecht and Kamil Doruk Gür. Verifiable oblivious pseudorandom functions from lattices: Practical-ish and thresholdisable. In Kai-Min Chung and Yu Sasaki, editors, *ASIACRYPT 2024, Part IV*, volume 15487 of *LNCS*, pages 205–237. Springer, Singapore, December 2024. [2](#), [3](#), [4](#), [7](#), [25](#), [30](#), [31](#), [37](#), [38](#)
- ALS20. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In *CRYPTO (2)*, *LNCS*, pages 470–499. Springer, 2020. [37](#)
- APRR24. Navid Alapati, Guru-Vamsi Policharla, Srinivasan Raghuraman, and Peter Rindal. Improved alternating-moduli PRFs and post-quantum signatures. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VIII*, volume 14927 of *LNCS*, pages 274–308. Springer, Cham, August 2024. [24](#)
- APS15. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Math. Cryptol.*, 9(3):169–203, 2015. [29](#)
- BDFH25. Ward Beullens, Lucas Dodgson, Sebastian H. Faller, and Julia Hesse. The 2Hash OPRF framework and efficient post-quantum instantiations. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VIII*, volume 15608 of *LNCS*, pages 332–362. Springer, Cham, May 2025. [2](#), [3](#), [24](#)
- BDL<sup>+</sup>18. Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In *SCN*, volume 11035 of *Lecture Notes in Computer Science*, pages 368–385. Springer, 2018. [23](#), [25](#), [69](#)
- BFM88. Manuel Blum, Paul Feldman, and Silvio Micali. Non-interactive zero-knowledge and its applications (extended abstract). In *20th ACM STOC*, pages 103–112. ACM Press, May 1988. [36](#)
- BIP<sup>+</sup>18. Dan Boneh, Yuval Ishai, Alain Passelègue, Amit Sahai, and David J. Wu. Exploring crypto dark matter: New simple PRF candidates and their applications. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 699–729. Springer, Cham, November 2018. [3](#)

- BLMR13. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *CRYPTO (1)*, volume 8042 of *LNCS*, pages 410–428. Springer, 2013. [6](#), [13](#), [21](#), [26](#), [38](#), [68](#)
- BLNS23. Ward Beullens, Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Lattice-based blind signatures: Short, efficient, and round-optimal. In *CCS*, pages 16–29. ACM, 2023. [39](#)
- BLP<sup>+</sup>13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. ACM, 2013. [11](#)
- Bol03. Alexandra Boldyreva. Threshold signatures, multisignatures and blind signatures based on the gap-Diffie-Hellman-group signature scheme. In Yvo Desmedt, editor, *PKC 2003*, volume 2567 of *LNCS*, pages 31–46. Springer, Berlin, Heidelberg, January 2003. [66](#)
- BP14. Abhishek Banerjee and Chris Peikert. New and improved key-homomorphic pseudorandom functions. In *CRYPTO (1)*, volume 8616 of *LNCS*, pages 353–370. Springer, 2014. [6](#), [12](#), [13](#), [21](#), [22](#), [26](#), [38](#), [41](#)
- BS23. Ward Beullens and Gregor Seiler. Labrador: Compact proofs for R1CS from module-sis. In *CRYPTO (5)*, volume 14085 of *LNCS*, pages 518–548. Springer, 2023. [2](#), [4](#), [26](#), [30](#), [37](#), [40](#)
- CHL22. Sílvia Casacuberta, Julia Hesse, and Anja Lehmann. SoK: Oblivious pseudorandom functions. In *2022 IEEE European Symposium on Security and Privacy*, pages 625–646. IEEE Computer Society Press, June 2022. [1](#)
- CJ25. Jung Hee Cheon and Daehyun Jang. Cryptanalysis on lightweight verifiable homomorphic encryption. In *ASIACRYPT (7)*, volume 16251 of *Lecture Notes in Computer Science*, pages 366–397. Springer, 2025. [3](#)
- COS20. Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. In *EUROCRYPT (1)*, volume 12105 of *Lecture Notes in Computer Science*, pages 769–793. Springer, 2020. [39](#)
- DDT25. Alex Davidson, Amit Deo, and Louis Tremblay Thibault. Pool: A practical OT-based OPRF from learning with rounding. In Chun-Ying Huang, Jyh-Cheng Chen, Shiuh-Pyng Shieh, David Lie, and Véronique Cortier, editors, *ACM CCS 2025*, pages 1038–1052. ACM Press, October 2025. [3](#), [4](#), [8](#), [10](#), [24](#)
- DGS<sup>+</sup>18. Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. Privacy pass: Bypassing internet challenges anonymously. *PoPETs*, 2018(3):164–180, July 2018. [1](#), [25](#)
- Di 03. Giovanni Di Crescenzo. Equivocable and extractable commitment schemes. In Stelvio Cimato, Clemente Galdi, and Giuseppe Persiano, editors, *SCN 02*, volume 2576 of *LNCS*, pages 74–87. Springer, Berlin, Heidelberg, September 2003. [35](#)
- DKL<sup>+</sup>18. Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018. [30](#)
- DPS23. Julien Devevey, Alain Passelègue, and Damien Stehlé. G+G: A fiat-shamir lattice signature based on convolved gaussians. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part VII*, volume 14444 of *LNCS*, pages 37–64. Springer, Singapore, December 2023. [11](#)
- ECS<sup>+</sup>15. Adam Everspaugh, Rahul Chatterjee, Samuel Scott, Ari Juels, and Thomas Ristenpart. The Pythia PRF service. In *USENIX Security Symposium*, pages 547–562. USENIX Association, 2015. [30](#)
- ENP24. Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerge: Distributed key generation and robust threshold signature from lattices. In *CRYPTO (7)*, volume 14926 of *Lecture Notes in Computer Science*, pages 425–458. Springer, 2024. [6](#), [7](#), [11](#), [13](#), [16](#), [17](#)
- ENS20. Muhammed F. Esgin, Ngoc Khanh Nguyen, and Gregor Seiler. Practical exact proofs from lattices: New techniques to exploit fully-splitting rings. In *ASIACRYPT (2)*, volume 12492 of *LNCS*, pages 259–288. Springer, 2020. [37](#)
- ESLL19. Muhammed F. Esgin, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. Lattice-based zero-knowledge proofs: New techniques for shorter and faster constructions and applications. In *CRYPTO (1)*, volume 11692 of *Lecture Notes in Computer Science*, pages 115–146. Springer, 2019. (Full version at [ia.cr/2019/445](#)). [28](#), [29](#)
- ESLR23. Muhammed F. Esgin, Ron Steinfeld, Dongxi Liu, and Sushmita Ruj. Efficient hybrid exact/relaxed lattice proofs and applications to rounding and VRFs. In *CRYPTO (5)*, volume 14085 of *LNCS*, pages 484–517. Springer, 2023. (Full version at [ia.cr/2022/141](#)). [29](#), [30](#), [37](#)
- ESZ22a. Muhammed F. Esgin, Ron Steinfeld, and Raymond K. Zhao. Efficient verifiable partially-decryptable commitments from lattices and applications. In *Public Key Cryptography (PKC) (1)*, volume 13177 of *Lecture Notes in Computer Science*, pages 317–348. Springer, 2022. (Full version at [ia.cr/2022/142](#)). [69](#)
- ESZ22b. Muhammed F. Esgin, Ron Steinfeld, and Raymond K. Zhao. MatRiCT<sup>+</sup>: More efficient post-quantum private blockchain payments. In *IEEE Symposium on Security and Privacy (S&P)*, pages 1281–1298. IEEE, 2022. (Full version at [ia.cr/2021/545](#)). [29](#)

- EZS<sup>+</sup>19. Muhammed F. Esgin, Raymond K. Zhao, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. MatRiCT: Efficient, scalable and post-quantum blockchain confidential transactions protocol. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pages 567–584. ACM, 2019. (Full version at [ia.cr/2019/1287](https://ia.cr/2019/1287)). 11, 28
- FIPR05. Michael J. Freedman, Yuval Ishai, Benny Pinkas, and Omer Reingold. Keyword search and oblivious pseudorandom functions. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 303–324. Springer, Berlin, Heidelberg, February 2005. 2
- GdKQ<sup>+</sup>24. Phillip Gajland, Bor de Kock, Miguel Quaresma, Giulio Malavolta, and Peter Schwabe. SWOOSH: Efficient lattice-based non-interactive key exchange. In Davide Balzarotti and Wenyuan Xu, editors, *USENIX Security 2024*. USENIX Association, August 2024. 4
- GKR<sup>+</sup>21. Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. In *USENIX Security Symposium*, pages 519–535. USENIX Association, 2021. 39
- GSW13. Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 75–92. Springer, 2013. 68
- Hei25. Lena Heimberger. Private communication, 2025. 4
- HKL<sup>+</sup>25. Lena Heimberger, Daniel Kales, Riccardo Lolato, Omid Mir, Sebastian Ramacher, and Christian Rechberger. Leap: A fast, lattice-based OPRF with application to private set intersection. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 254–283. Springer, Cham, May 2025. 3, 4, 24
- JKK14. Stanislaw Jarecki, Aggelos Kiayias, and Hugo Krawczyk. Round-optimal password-protected secret sharing and T-PAKE in the password-only model. In *ASIACRYPT (2)*, volume 8874 of *Lecture Notes in Computer Science*, pages 233–253. Springer, 2014. 1, 2, 3, 4, 5, 22, 24
- JKKX16. Stanislaw Jarecki, Aggelos Kiayias, Hugo Krawczyk, and Jiayu Xu. Highly-efficient and composable password-protected secret sharing (or: How to protect your Bitcoin wallet online). In *2016 IEEE European Symposium on Security and Privacy*, pages 276–291. IEEE Computer Society Press, March 2016. 1, 24
- JKX18. Stanislaw Jarecki, Hugo Krawczyk, and Jiayu Xu. OPAQUE: An asymmetric PAKE protocol secure against pre-computation attacks. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 456–486. Springer, Cham, April / May 2018. 1, 4, 24
- JL09. Stanislaw Jarecki and Xiaomin Liu. Efficient oblivious pseudorandom function with applications to adaptive OT and secure computation of set intersection. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 577–594. Springer, Berlin, Heidelberg, March 2009. 2
- Kat21. Shuichi Katsumata. A new simple technique to bootstrap various lattice zero-knowledge proofs to QROM secure NIZKs. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 580–610, Virtual Event, August 2021. Springer, Cham. 4, 65
- KBR13. Sriram Keelveedhi, Mihir Bellare, and Thomas Ristenpart. DupLESS: Server-aided encryption for deduplicated storage. In Samuel T. King, editor, *USENIX Security 2013*, pages 179–194. USENIX Association, August 2013. 1
- KLSS23. Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-mlwe. In *CRYPTO (5)*, volume 14085 of *Lecture Notes in Computer Science*, pages 549–580. Springer, 2023. 6, 7, 11, 13, 16, 17
- LNP22. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In *CRYPTO (2)*, volume 13508 of *LNCS*, pages 71–101. Springer, 2022. Full version at (<https://eprint.iacr.org/archive/2022/284/20220814:160144>). 4, 25, 29, 30, 37, 69
- LNPS21. Vadim Lyubashevsky, Ngoc Khanh Nguyen, Maxime Plançon, and Gregor Seiler. Shorter lattice-based group signatures via “almost free” encryption and other optimizations. In *ASIACRYPT (4)*, volume 13093 of *LNCS*, pages 218–248. Springer, 2021. 68, 69
- LNS20. Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Gregor Seiler. Practical lattice-based zero-knowledge proofs for integer relations. In *ACM CCS*, pages 1051–1070. ACM, 2020. 28, 37
- LPA<sup>+</sup>19. Lucy Li, Bijeeta Pal, Junade Ali, Nick Sullivan, Rahul Chatterjee, and Thomas Ristenpart. Protocols for checking compromised credentials. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 1387–1403. ACM Press, November 2019. 1
- MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. In *FOCS*, pages 372–381. IEEE Computer Society, 2004. 11

- NO25. Ngoc Khanh Nguyen and George O'Rourke. More efficient lattice-based zero-knowledge proofs with straight-line extractability. In *APKC@AsiaCCS*, pages 34–43. ACM, 2025. [4](#)
- NS24. Ngoc Khanh Nguyen and Gregor Seiler. Greyhound: Fast polynomial commitments from lattices. In *CRYPTO (10)*, volume 14929 of *LNCS*, pages 243–275. Springer, 2024. [28](#), [40](#)
- Pei10. Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97. Springer, 2010. [11](#)
- Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM, 2005. [23](#), [28](#), [68](#)
- SSS23. Shifeng Sun, Ron Steinfeld, and Amin Sakzad. Incremental symmetric puncturable encryption with support for unbounded number of punctures. *Des. Codes Cryptogr.*, 91(4):1401–1426, 2023. [38](#), [48](#)
- TCR<sup>+</sup>22. Nirvan Tyagi, Sofía Celi, Thomas Ristenpart, Nick Sullivan, Stefano Tessaro, and Christopher A. Wood. A fast and simple partially oblivious PRF, with applications. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part II*, volume 13276 of *LNCS*, pages 674–705. Springer, Cham, May / June 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [8](#), [9](#), [10](#), [23](#), [27](#), [65](#), [66](#)
- TPY<sup>+</sup>19. Kurt Thomas, Jennifer Pullman, Kevin Yeo, Ananth Raghunathan, Patrick Gage Kelley, Luca Invernizzi, Borbala Benko, Tadek Pietraszek, Sarvar Patel, Dan Boneh, and Elie Bursztein. Protecting accounts from credential stuffing with password breach alerting. In Nadia Heninger and Patrick Traynor, editors, *USENIX Security 2019*, pages 1556–1571. USENIX Association, August 2019. [1](#)
- YAZ<sup>+</sup>19. Rupeng Yang, Man Ho Au, Zhenfei Zhang, Qiuliang Xu, Zuoxia Yu, and William Whyte. Efficient lattice-based zero-knowledge arguments with standard soundness: Construction and applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 147–175. Springer, Cham, August 2019. [65](#)
- YBH<sup>+</sup>25. Yibin Yang, Fabrice Benhamouda, Shai Halevi, Hugo Krawczyk, and Tal Rabin. Gold OPRF: Post-quantum oblivious power-residue PRF. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *2025 IEEE Symposium on Security and Privacy*, pages 259–278. IEEE Computer Society Press, May 2025. [2](#), [3](#), [4](#)
- ZSE<sup>+</sup>24. Xinyu Zhang, Ron Steinfeld, Muhammed F. Esgin, Joseph K. Liu, Dongxi Liu, and Sushmita Ruj. Loquat: A snark-friendly post-quantum signature based on the legendre PRF with applications in ring and aggregate signatures. In *CRYPTO (1)*, volume 14920 of *LNCS*, pages 3–38. Springer, 2024. [39](#)

## A Additional Preliminaries

### A.1 Extractable Commitment Scheme

In this work, we consider an extractable commitment scheme [Di 03], which is formally defined as follows.

**Definition 14 (Extractable Commitment Scheme).** A (non-interactive) extractable commitment scheme consists of a tuple of algorithms  $\text{COM} = (\text{Setup}, \text{Commit}, \text{Verify})$  defined as follows:

$\text{Setup}(1^\lambda)$ : is a PPT algorithm that on input a (unary encoded) security parameter  $\lambda$ , outputs a commitment key  $\text{ck}$ .

$\text{Commit}(\text{ck}, m)$ : is a PPT algorithm that on input a commitment key  $\text{ck}$  and a message  $m \in \{0, 1\}^\lambda$ , outputs a commitment  $c$  and an opening information  $d$ .

$\text{Verify}(\text{ck}, c, d, m)$ : is a DPT algorithm that on input a commitment key  $\text{ck}$ , commitment  $c$ , an opening information  $d$  and a message  $m \in \{0, 1\}^\lambda$ , outputs a bit  $b \in \{0, 1\}$ .

In the above  $\text{Commit}$  definition, we leave the randomness part implicit and assume that it is generated internally inside the  $\text{Commit}$  function. When proving well-formedness of commitments via a zero-knowledge proof, the randomness will also be part of the prover's witness. When we need to specify the randomness  $r$  in such cases, we will write  $\text{Commit}(\text{ck}, m; r)$  to explicitly refer to the internally generated randomness of the commitment.

We require the standard notion of *correctness*, which says that for every  $\lambda \in \mathbb{N}$ , every  $\text{ck} \leftarrow \text{Setup}(1^\lambda)$  and every message  $m \in \{0, 1\}^\lambda$ , it holds that

$$\Pr [\text{Verify}(\text{ck}, \text{Commit}(\text{ck}, m), m) = 1] = 1.$$

In terms of security, we require the commitment scheme to satisfy *computational hiding* and *computational binding* properties, along with *extractability*.

**Definition 15 (Computational Hiding).** A commitment scheme  $\text{COM}$  is computationally hiding if for every PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$ , such that

$$\Pr [\text{CH}_{\text{COM}, \mathcal{A}}(\lambda) = 1] \leq \frac{1}{2} + \text{negl}(\lambda),$$

where the experiment  $\text{CH}_{\text{COM}, \mathcal{A}}$  is defined as follows

$\text{CH}_{\text{COM}, \mathcal{A}}(\lambda)$
1 : $\text{ck} \leftarrow \text{Setup}(1^\lambda)$
2 : $(m_0, m_1) \leftarrow \mathcal{A}(\text{ck})$
3 : $b \xleftarrow{\$} \{0, 1\}$
4 : $(c, d) \leftarrow \text{Commit}(\text{ck}, m_b)$
5 : $b' \leftarrow \mathcal{A}(c)$
6 : <b>return</b> $b = b'$

**Definition 16 (Computational Binding).** A commitment scheme  $\text{COM}$  is computationally binding if for every PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$ , such that

$$\Pr [\text{CB}_{\text{COM}, \mathcal{A}}(\lambda) = 1] \leq \text{negl}(\lambda),$$

where the experiment  $\text{CB}_{\text{COM}, \mathcal{A}}$  is defined as follows

$\text{CB}_{\text{COM}, \mathcal{A}}(\lambda)$
1 : $\text{ck} \leftarrow \text{Setup}(1^\lambda)$
2 : $(c, d, d', m, m') \leftarrow \mathcal{A}(\text{ck})$
3 : <b>assert</b> $m \neq m'$
4 : <b>return</b> $\text{Verify}(\text{ck}, c, d, m) = 1 \wedge \text{Verify}(\text{ck}, c, d', m') = 1$

**Definition 17 (Extractability).** A commitment scheme  $\text{COM}$  is extractable, if there exists a pair of PPT algorithm  $\mathcal{E} = (\mathcal{E}_1, \mathcal{E}_2)$ , called the extractor, such that for all  $m \in \{0, 1\}^\lambda$  and for every PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$ , such that

$$\left| \Pr [\mathcal{A}(\text{ck}) = 1 \mid \text{ck} \leftarrow \text{Setup}(1^\lambda)] - \Pr [\mathcal{A}(\text{ck}) = 1 \mid (\text{ck}, \text{td}) \leftarrow \mathcal{E}_1(1^\lambda)] \right| \leq \text{negl}(\lambda),$$

and

$$\Pr \left[ \begin{array}{c} \text{Verify}(c, d, m) = 1 \\ \wedge \\ m \neq m' \end{array} \middle| \begin{array}{c} (\text{ck}, \text{td}) \leftarrow \mathcal{E}_1(1^\lambda) \\ (c, d, m) \leftarrow \mathcal{A}(\text{ck}) \\ m' \leftarrow \mathcal{E}_2(\text{ck}, \text{td}, c) \end{array} \right] \leq \text{negl}(\lambda).$$

## A.2 Non-Interactive Zero-Knowledge Arguments

Let  $R$  be a binary relation and  $L$  the language consisting of statements in  $R$ . We formally define a non-interactive zero-knowledge (NIZK) argument system [BFM88] as follows.

**Definition 18 (Non-Interactive Zero-Knowledge Argument System).** A non-interactive zero-knowledge (NIZK) argument system  $\text{NIZK}$  for a language  $L \in \text{NP}$  (with witness relation  $R$ ) is a tuple of algorithms  $\text{NIZK} = (\text{PGen}, \text{P}, \text{V})$ , such that:

$\text{PGen}(1^\lambda)$ : is a PPT algorithm that on input a (unary encoded) security parameter  $\lambda$ , outputs a common reference string  $\text{crs}$ .

$\text{P}(\text{crs}, x, w)$ : is a PPT algorithm that on input a common reference string  $\text{crs}$ , a statement  $x$  and a witness  $w$ , outputs a proof  $\pi$ .

$\text{V}(\text{crs}, x, \pi)$ : is a DPT algorithm that on input a common reference string  $\text{crs}$ , a statement  $x$  and a proof  $\pi$ , outputs a bit  $b$ .

We require  $\text{NIZK}$  to meet the following properties:

**Perfect Completeness.** For every  $(x, w) \in R$  we have that

$$\Pr [\text{V}(\text{crs}, x, \pi) = 1 \mid \text{crs} \leftarrow \text{PGen}(1^\lambda), \pi \leftarrow \text{P}(\text{crs}, x, w)] = 1.$$

**Computational Soundness.** For every  $x \notin L$ , and every PPT adversary  $\mathcal{A}$ , we have that

$$\Pr [\text{V}(\text{crs}, x, \pi) = 1 \mid \text{crs} \leftarrow \text{PGen}(1^\lambda), \pi \leftarrow \mathcal{A}(\text{crs}, x)] \leq \text{negl}(\lambda).$$

**Computational Zero-Knowledge.** There exists a PPT algorithm  $\mathcal{S} = (\mathcal{S}_1, \mathcal{S}_2)$  such that for every PPT adversary  $\mathcal{A}$ ,

$$\left| \Pr [\mathcal{A}^{\text{P}(\text{crs}, \cdot, \cdot)}(\text{crs}) = 1 \mid \text{crs} \leftarrow \text{PGen}(1^\lambda)] - \Pr [\mathcal{A}^{\mathcal{O}(\text{crs}_s, \tau_s, \cdot, \cdot)}(\text{crs}) = 1 \mid (\text{crs}_s, \tau_s) \leftarrow \mathcal{S}_1(1^\lambda)] \right| \leq \text{negl}(\lambda),$$

where  $\mathcal{O}(\text{crs}_s, \tau_s, \cdot, \cdot)$  is an oracle that outputs  $\perp$  on input  $(x, w)$  when  $(x, w) \notin R$  and outputs  $\pi \leftarrow \mathcal{S}_2(\text{crs}_s, \tau_s, x)$  when  $(x, w) \in R$ .

## B Averaging Attack Against MLWE-PRF-RU

In this section we describe a simple error “averaging” attack on MLWE-PRF-RU (and our OPRF protocol as well as those using the same blueprint [ADDS21, AG24]).

In this attack, the adversary uses the  $Q_x^\infty$  samples on same  $x$  queried to the oracle to compute  $\mathbf{z}'_j := \mathbf{u}_j - \mathbf{R}_j \mathbf{c}_j = \mathbf{B}_x \mathbf{k} + \mathbf{e}'_j - \mathbf{R}_j \mathbf{e}_j$  for  $j \in [Q_x^\infty]$ . The attacker then computes the average (over the rationals) of these noisy secrets to get  $\bar{\mathbf{z}} := \frac{1}{Q_x^\infty} \sum_{j \in [Q_x^\infty]} \mathbf{z}'_j = \mathbf{B}_x \mathbf{k} + \bar{\mathbf{e}}$  in which the error  $\bar{\mathbf{e}} := \frac{1}{Q_x^\infty} \sum_{j \in [Q_x^\infty]} (\mathbf{e}'_j - \mathbf{R}_j \mathbf{e}_j)$  has a reduced standard deviation  $\bar{\sigma} \approx \sigma_1 / \sqrt{Q_x^\infty}$ . If  $\bar{\sigma}$  is a constant factor smaller than  $1/2$ , all the  $\bar{\mathbf{e}}$  error coordinates will likely be  $< 1/2$  and rounding  $\bar{\mathbf{z}}$  to the nearest integers will give the attacker the secret  $\mathbf{B}_x \mathbf{k}$  (and hence  $\mathbf{k}$ ) with high probability. Consequently, to prevent the averaging attack, a necessary condition is  $\sigma_1 \geq \sqrt{Q_x^\infty}/c$  for a small constant  $c > 1$ . More generally, a necessary condition to prevent the averaging attack from reducing the error standard deviation (and MLWE security) below the  $\sigma$  standard deviation used to set MLWE security in the samples  $\mathbf{c}_j$ , is that  $\sigma_1/\sigma \geq \sqrt{Q_x^\infty}/c$ . This lower bound matches (up to the small factor  $c$ ) our MLWE-PRF-RU security reduction parameter lower bound discussed above (for large  $Q_x^\infty$ ), showing the optimality of our reduction condition up to a small factor. Note that the attack and our security reduction bounds only depend on the maximum no. of samples queried per  $x$  value. This is in contrast to the suboptimal Rényi-divergence-based security analysis in [AG24], which seems to inherently lose a factor proportional to  $\sqrt{Q}$ , where  $Q$  is the total number of oracle queries over *all*  $x$  values.

## C Instantiating NIZK Proofs and Message Mapping H

In this section, we discuss possible ways to instantiate the message mapping  $\mathbf{H}$  and the underlying non-interactive zero-knowledge (NIZK) proofs performed by the server and the client.

### C.1 Proof by the Server

The NIZK proof,  $\pi_s$ , (see Figure 10) conducted by the server is the most typical proof needed in lattice-based cryptography, and there are various proof systems that can be used to instantiate it. Some notable ones are the LANES [ALS20, ENS20, LNS20], LANES<sup>+</sup> [ESLR23], LNP22 [LNP22] and LaBRADOR [BS23] proof systems. As discussed in Appendix F, we use a lattice-based commitment scheme, and thus, these proof systems can natively support its well-formedness proof. Given we require a relatively large modulus  $q$ , our experiments show that LaBRADOR in combination with LNP22 to achieve zero-knowledge offers the best communication performance among these options. Thanks to its succinctness features, it also scales very efficiently in the batched setting. Note that even though LaBRADOR as described in [BS23] does not directly achieve zero-knowledge, it is easy to achieve it as discussed in “Zero-Knowledge Property” part of [BS23, Page 3]. In particular, we assume LNP22 is used as the shim protocol to achieve zero-knowledge. We report on the performance results in Section 5.

### C.2 Instantiating the Message Mapping H

In this section, we primarily focus on the *secret-dependent* part of the message mapping  $\mathbf{H}$  and discuss our main option on how the message mapping  $\mathbf{H}: x \mapsto \mathbf{B}_x \in R_q^{h \times m}$  can be instantiated. That is, we assume the public input part  $t$  to be empty. When the tag  $t$  is present, the mapping would simply be computed on  $\hat{x} := x || t$  as  $\mathbf{H}(\hat{x})$  and it is easy to verify the computation w.r.t. this input extended with *public* information (which we discuss more in Appendix C.3). There are two main considerations: (i) security of MLWE-PRF-RU w.r.t. to  $\mathbf{H}$  chosen, and (ii) efficiency of our OPRF, particularly the underlying NIZK proofs especially by the client. The latter is the primary consideration as we believe MLWE-PRF-RU remains secure for all instantiations discussed in this work given the reductions in Section 3.

In line with the MLWE-PRF-RU security reductions in Section 3, we believe that a natural way to instantiate  $\mathbf{H}$  is to use the mappings in existing lattice-based PRFs so that the **SampPRF** oracle in MLWE-PRF really serves as an oracle outputting (randomized) PRF samples (in the case of  $b = 1$ ). Therefore, we discuss options based on existing lattice-based PRFs.

**Using the mapping from BP14 PRF [BP14] for H.** The option we see as the most suitable one is using the message mapping employed in BP14 PRF [BP14]. Let  $\mathfrak{l} := \lceil \log q \rceil$ . Here, the PRF computation involves a gadget matrix  $\mathbf{G}$  and its (non-linear) inverse computation  $\mathbf{G}^{-1}: R_q^{m \times m\mathfrak{l}} \rightarrow R_q^{m\mathfrak{l} \times m\mathfrak{l}}$ . The  $\mathbf{G}^{-1}$  mapping is the standard bit decomposition operation. In this case, we have two matrices  $\mathbf{A}_0, \mathbf{A}_1 \xleftarrow{\$} R_q^{m \times m\mathfrak{l}}$  published as public parameters. Then, for  $\mathbf{k} \leftarrow \chi_k^m$  for some distribution  $\chi_k$ <sup>12</sup> on  $R_q$ , we compute<sup>13</sup>

$$F_{\mathbf{k}}^{\text{bp}}(x) = \lfloor \mathbf{B}_x^{\text{bp}} \cdot \mathbf{k} \rfloor_p, \quad (28)$$

where the message mapping  $H^{\text{bp}}$  is defined as

$$H^{\text{bp}}: x \mapsto (\mathbf{B}_x^{\text{bp}})^\top := \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-2}} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))). \quad (29)$$

If there is a tag  $t$  present, we treat it as the last bits of the input to  $H^{\text{bp}}$ . For better efficiency of the OPRF, we can truncate  $\mathbf{B}_x^{\text{bp}}$  and just use its first row  $\mathbf{B}_x^{\text{bp}}[0]$  for the final PRF computation (as also done in [AG24]). Note that this does not impact security because pseudorandomness of  $\lfloor \mathbf{B}_x \cdot \mathbf{k} \rfloor_p$  implies pseudorandomness of  $\lfloor \mathbf{B}_x[0] \cdot \mathbf{k} \rfloor_p$  in general as the latter is just the first coordinate of the former vector. Overall, for the BP14 option, we take the message mapping used in LeOPaRd as  $H: x \mapsto \mathbf{B}_x := \mathbf{B}_x^{\text{bp}}[0] \in R_q^{1 \times m}$ , resulting in  $h = 1$ . As discussed in [BP14], another well-known lattice-based PRF scheme, BMLR13 PRF, [BLMR13] can be seen as a special case of BP14 PRF.

For the client's NIZK proof in this case (discussed further in Appendix C.3), we will need to treat each  $\mathbf{B}_i := \mathbf{G}^{-1}(\mathbf{A}_{x_i} \cdot \mathbf{B}_{i+1}) \in R_q^{m\mathfrak{l} \times m\mathfrak{l}}$  for  $i = L-1, \dots, 0$  (with  $\mathbf{B}_L := \mathbf{I}_m$ ) as a variable. Therefore, there are  $m^2 \mathfrak{l}^2 L$  variable polynomials.

We may also consider a generalized version of this PRF proposal where the gadget matrix works with base  $\gamma \geq 2$  (instead of  $\gamma = 2$  as before), i.e.,  $\mathbf{g} = (1, \gamma, \gamma^2, \dots)$ , and  $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}$ . In this case, we would have that  $\mathfrak{l} = \lceil \log_\gamma q \rceil$ .

**Using the mapping from BLMR13 PRF [BLMR13] for H.** An other option is the BLMR13 PRF [BLMR13] (particularly its variant over module lattices as described in [SSS23]). In this setting, we have two square matrices  $\mathbf{A}_0, \mathbf{A}_1 \in R_q^{m \times m}$  (with binary entries) published as public parameters. To compute the PRF for an input  $x = (x_0, \dots, x_{L-1})$  in binary and  $\mathbf{k} \xleftarrow{\$} R_q^m$ , we compute

$$F_{\mathbf{k}}^{\text{blmr}}(x) = \left\lfloor \prod_{i=0}^{L-1} \mathbf{A}_{x_i} \cdot \mathbf{k} \right\rfloor_p. \quad (30)$$

As a result, the message mapping  $H^{\text{blmr}}$  is defined as

$$H^{\text{blmr}}: x \mapsto \mathbf{B}_x^{\text{blmr}} = \prod_{i=0}^{L-1} \mathbf{A}_{x_i}. \quad (31)$$

If there is a tag  $t$  present, we treat  $t$  as the first part of the message in this case. That is, if  $t = (t_0, \dots)$  and  $x = (x_0, \dots)$ , then  $\mathbf{B}_x^{\text{blmr}} = \prod_i \mathbf{A}_{t_i} \prod_i \mathbf{A}_{x_i}$ . As in Appendix C.2, we can truncate the final PRF output and take  $\mathbf{B}_x = \mathbf{B}_x^{\text{blmr}}[0] \in R_q^{1 \times m}$  as the first row of  $\mathbf{B}_x^{\text{blmr}}$  in LeOPaRd (i.e.,  $h = 1$ ).

For the client's NIZK proof  $\pi_c$  (discussed in Appendix C.3), we will need to treat each  $\mathbf{A}_{x_i}$  as a set of different variables. It is easy to observe that there are  $m^2 L$  variable polynomials in that case. As discussed in [SSS23], we require  $m \geq 2$  as otherwise this module variant of the PRF scheme is insecure. In fact,  $m = n \lceil \log q \rceil$  is required in the provable security result in [BLMR13] to base pseudorandomness on MLWE hardness with rank  $n$ .

<sup>12</sup> Note that in this case we may have small secret key coefficients.

<sup>13</sup> We note that our presentation here is for multiplying the PRF key from the *right* unlike [BP14] where it is from the *left*.

**Using a (symmetric) PRG for H.** Another option one may consider is to use a (symmetric) PRG to instantiate H. In this case, we may simply set  $h = 1$  to minimize a dimension of certain matrices/vectors involved in communication and NIZK proofs for improved performance. In this case, the `SampPRF` oracle in MLWE-PRF would output samples (computationally) indistinguishable from MLWE samples (based on the PRG pseudorandomness property). One may then employ a suitable zk-SNARK system to instantiate the client’s proof  $\pi_c$ . We do not see this option as competitive (in terms of efficiency) as the prior ones where fully lattice-based tools can be employed for improved compatibility and performance. Therefore, we do not discuss this option in further detail. However, as further progress is made in the development of more efficient zk-SNARK solutions, this option (just like the others) may become more efficient over time.

**Modelling H as a random oracle.** We also note the option of modeling H as a random oracle. In this case, we may again simply set  $h = 1$  to minimize a dimension of certain matrices/vectors involved in communication and NIZK proofs for improved performance. In this case, the `SampPRF` oracle in MLWE-PRF would simply output MLWE samples, and the hardness of MLWE-PRF-RU problem is more easily established based on MLWE as discussed in Section 3. Assuming H to be a random oracle would also reduce the control of an OPRF pseudorandomness adversary on how  $\mathbf{C}_x$  is computed.

There are two (related) caveats with this choice. Firstly, we would need a NIZK proof that proves pre-image knowledge of a hash function (modelled as a random oracle). There are generic zk-SNARK systems like [COS20] that can achieve this for standard hash function such as SHA3, but they are currently not as efficient as we would like them to be. Alternatively, we may use a more algebraic hash function like Poseidon [GKR+21] to help reduce the cost of the NIZK well-formedness proof. Overall, the first caveat is that the associated message-mapping well-formedness proof is not going to be as efficient.

The second caveat is that we would assume H to be a random oracle, but then, as we need to prove message-mapping well-formedness via a NIZK proof, we would also need to assume that H has a compact circuit representation. This is the same issue that has already been encountered in earlier works such as the recursive proof in [COS20], the blind signature in [BLNS23] and the aggregate signature in [ZSE+24]. Similar to earlier works, we can (heuristically) assume that the construction remains secure when the random oracle is instantiated with a hash function and then instantiate the NIZK accordingly.

Overall, while it does not currently appear as a very competitive option in terms of (NIZK) efficiency, we believe that the main advantage of modelling H as a random oracle arises in helping with the security reductions (as in Section 3) and establishing a tighter (theoretical) bound on an adversary’s winning chance in OPRF pseudorandomness.

### C.3 Proof by the Client

In this section, we discuss how the client’s NIZK proof can be instantiated. One can observe from the NIZK relation proven by the client (given in Figure 10) that the majority of the proof components are the most common relations proven in lattice-based cryptography (when the commitment scheme is instantiated using a suitable lattice-based extractable commitment scheme). There is perhaps one exception to this: (i) proving  $\mathbf{B}_{x,t} := \mathbf{H}(x,t)$ . Accordingly, we next discuss the proof of  $\mathbf{B}_{x,t} := \mathbf{H}(x,t)$ . Depending on the instantiation of the message mapping H, the client’s proof may vary significantly. We look more closely at the instantiation of H discussed in Appendix C.2.

**NIZK for BP14 PRF mapping.** Recall we have  $h = 1$  and  $\mathsf{l} := \lceil \log_\gamma q \rceil$  for some  $\gamma \geq 2$ . We first discuss the case when no tag  $t$  is present, i.e.,  $\mathbf{B}_x := \mathbf{H}(x)$ . The case of  $\mathbf{B}_{x,t} := \mathbf{H}(x,t)$  is discussed at the end of this section. For a known  $\mathbf{A}_r$  and  $\mathbf{C}_x$ , the prover wants to prove knowledge of short matrix  $\mathbf{R}$  and a bit-string  $x = (x_0, x_1, \dots, x_{L-1}) \in \{0, 1\}^L$  such that  $\mathbf{C}_x = \mathbf{R}\mathbf{A}_r + \mathbf{B}_x \bmod q$  where  $\mathbf{B}_x := \mathbf{B}_x^{\text{bp}}[0]$  and  $\mathbf{B}_x^{\text{bp}}$  is defined in (29).

Using a similar strategy as in [ADDS21] (see “**Proof system 1: Proofs of Masked Partial PRF Computation**”), define variables  $\mathbf{B}_i \in R_q^{m\mathsf{l} \times m\mathsf{l}}$  for  $i = L-1, \dots, 0$  as  $\mathbf{B}_{L-1} := \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}})$  and  $\mathbf{B}_i := \mathbf{G}^{-1}(\mathbf{A}_{x_i} \cdot \mathbf{B}_{i+1})$  for  $i = L-2, \dots, 0$ . Using this, we have  $\mathbf{B}_x = (\mathbf{G} \cdot \mathbf{B}_0)^\top[0]$ . Then the statement being

proven now can be transformed to prove the following system with  $\mathbf{B}_L := \mathbf{I}$ :

$$\begin{aligned}\mathbf{G} \cdot \mathbf{B}_i &= \mathbf{A}_{x_i} \cdot \mathbf{B}_{i+1}, \quad i = 0, \dots, L-1 \\ \mathbf{C}_x &= \mathbf{R}\mathbf{A}_r + (\mathbf{G} \cdot \mathbf{B}_0)^\top [0].\end{aligned}$$

This system of equations above is not linear due to the  $x_i$  and  $\mathbf{B}_{i+1}$  terms and the fact that  $\mathbf{G}^{-1}$  is not a linear operator. We can manipulate the system of equations above to obtain

$$\begin{aligned}\mathbf{G} \cdot \mathbf{B}_{L-1} &= \mathbf{A}_0 \cdot (1 - x_{L-1}) + \mathbf{A}_1 \cdot x_{L-1}, \\ \mathbf{G} \cdot \mathbf{B}_{L-2} &= \mathbf{A}_0 \cdot (1 - x_{L-2}) \cdot \mathbf{B}_{L-1} + \mathbf{A}_1 \cdot x_{L-2} \cdot \mathbf{B}_{L-1}, \\ &\vdots \\ \mathbf{G} \cdot \mathbf{B}_0 &= \mathbf{A}_0 \cdot (1 - x_0) \cdot \mathbf{B}_1 + \mathbf{A}_1 \cdot x_0 \cdot \mathbf{B}_1, \\ \mathbf{C}_x &= \mathbf{R}\mathbf{A}_r + (\mathbf{G} \cdot \mathbf{B}_0)^\top [0].\end{aligned}$$

The required proof now boils down to proving that a sequence of quadratic equations holds (since  $\mathbf{A}_i$ 's are public), which can be easily handled by LaBRADOR [BS23]. Observe that the witness dimension (over  $R_q$ ) is in the order of  $m^2 l^2 L$ . In our parameter settings (see Section 5), the largest  $m$  is around 64 and the largest  $\log_2 q$  is less than 128; hence  $m^2 l^2 L = m^2 L (\log_2 q / \log_2 \gamma)^2 \leq 2^{32}$  for  $L = 64$ . By choosing larger  $\gamma$  or smaller rank parameter  $m$ , this witness dimension can be significantly reduced. In [NS24], the authors report a 53 KB LaBRADOR-based proof size for a polynomial evaluation proof of degree  $2^{30}$  with all running times around a couple of minutes or less. Therefore, we believe the proof we require will be reasonably efficient.

Note that when a tag  $t$  is used in computing  $\mathbf{B}_{x,t} := \mathbf{H}(x, t)$  (i.e.,  $\mathbf{B}_{\hat{x}} := \mathbf{H}(\hat{x})$  for  $\hat{x} := x||t$ ), the above system of equations will have a minor change in the first expression such that we will have  $\mathbf{G} \cdot \mathbf{B}_{L-1} = \mathbf{A}_0 \cdot (1 - x_{L-1}) \cdot \mathbf{B}_t + \mathbf{A}_1 \cdot x_{L-1} \cdot \mathbf{B}_t$  for some *public* matrix  $\mathbf{B}_t$  dependent on the tag  $t$ . The NIZK proof can equivalently prove this similar system of equations.

**NIZK for BLMR13 PRF mapping.** Here, for a known  $\mathbf{A}_r$  and  $\mathbf{C}_x$ , the prover (client) wants to prove knowledge of short  $\mathbf{R}$  and a bit-string  $x = (x_0, x_1, \dots, x_{L-1}) \in \{0, 1\}^L$  such that  $\mathbf{C}_x = \mathbf{R}\mathbf{A}_r + \mathbf{B}_x \bmod q$  where  $\mathbf{B}_x := \mathbf{B}_x^{\text{blmr}}[0]$  and  $\mathbf{B}_x^{\text{blmr}}$  is defined in (31). To be more specific,  $\mathbf{B}_x^{\text{blmr}} = \prod_{i=0}^{L-1} \mathbf{A}_{x_i}$ , where  $\mathbf{A}_0, \mathbf{A}_1$  are

two public matrices. Now, define variables  $\mathbf{B}_i \in R_q^{m \times m}$  for  $i = L-1, \dots, 0$  as  $\mathbf{B}_i := \prod_{j=0}^i \mathbf{A}_{x_j}$ . Then, the statement being proved can be expressed as follows.

$$\begin{aligned}\mathbf{B}_0 &= \mathbf{A}_{x_0}, \\ \mathbf{B}_{i+1} &= \mathbf{B}_i \cdot \mathbf{A}_{x_{i+1}}, \quad i = 0, \dots, L-2, \\ \mathbf{C}_x &= \mathbf{R}\mathbf{A}_r + \mathbf{B}_{L-1}.\end{aligned}$$

This system of equations above is not linear due to the  $x_i$  and  $\mathbf{B}_{i+1}$  terms. We can represent the terms  $\mathbf{B}_i \cdot \mathbf{A}_{x_{i+1}}$  as  $\mathbf{B}_i \cdot \mathbf{A}_0 \cdot (1 - x_{i+1}) + \mathbf{B}_i \cdot \mathbf{A}_1 \cdot x_{i+1}$ , then we obtain

$$\begin{aligned}\mathbf{B}_0 &= \mathbf{A}_0 \cdot (1 - x_0) + \mathbf{A}_1 \cdot x_0, \\ \mathbf{B}_1 &= \mathbf{B}_0 \cdot \mathbf{A}_0 \cdot (1 - x_1) + \mathbf{B}_0 \cdot \mathbf{A}_1 \cdot x_1, \\ &\vdots \\ \mathbf{B}_{L-1} &= \mathbf{B}_{L-2} \cdot \mathbf{A}_0 \cdot (1 - x_{L-1}) + \mathbf{B}_{L-2} \cdot \mathbf{A}_1 \cdot x_{L-1}, \\ \mathbf{C}_x &= \mathbf{R}\mathbf{A}_r + \mathbf{B}_{L-1}.\end{aligned}$$

Now the required proof boils down to proving that a sequence of quadratic equations holds, which can again be easily handled by LaBRADOR [BS23]. Observe that the witness dimension (over  $R_q$ ) is again in the order of  $m^2 L$ . Note also that when a tag  $t$  is present, we will have a similar situation as described before for the BP14 PRF, where some *public* matrix  $\mathbf{B}_t$  will appear in the first expression describing  $\mathbf{B}_0$ .

## D Deferred Definitions and Proofs

In the security proofs, we use various tables (denoted such as  $\mathcal{T}$ ) to store values. For ease of presentation, we write  $\mathcal{T}[x] = \perp$  to mean that a value,  $x$ , is not stored/defined in the table  $\mathcal{T}$  before. For a newly-initialized table,  $\mathcal{T}[x] = \perp$  for all  $x$ .

### D.1 Proof of Hardness Reduction MLWE $\rightarrow$ MLWE-PRF for BP14 $H^{\text{bp}}$

**Theorem 4 (MLWE  $\rightarrow$  MLWE-PRF for BP14 H).**

Let  $H^{\text{bp}}$  be the BP14 function family with  $L$ -bit input and error distribution  $\chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$  as in Def. 13. Let  $\chi_k := \mathcal{U}(R_q)$ ,  $\chi_0 := \mathcal{D}_{\mathbb{Z}^d, s_0}$  and  $\chi_{0, B} := \chi_{\mathbf{A}_0, \mathbf{A}_1, x, s_0}$ . The MLWE-PRF $_{H^{\text{bp}}, Q_P, Q_M, q, m, h, L, \chi_0, \chi_0, B, \chi_k}$  assumption holds if the MLWE $_{2m\ell + Q_M, q, m, 1, \chi_0, \chi_k}$  and MLWE $_{2m\ell, q, m, Q_P, \chi_0, \chi_k}$  assumptions hold.

More precisely, for any PPT adversary  $\mathcal{A}$  against the MLWE-PRF $_{H^{\text{bp}}, Q_P, Q_M, q, m, h, L, \chi_0, \chi_0, B, \chi_k}$  problem, there exist PPT adversaries  $\mathcal{B}, \bar{\mathcal{B}}$  against the MLWE $_{2m\ell + Q_M, q, m, 1, \chi_0, \chi_k}$  and MLWE $_{2m\ell, q, m, Q_P, \chi_0, \chi_k}$  problems respectively, such that

$$\text{Adv}_{\mathcal{A}}^{\text{MLWE-PRF}}(\lambda) \leq \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda) + (L-1) \cdot \text{Adv}_{\bar{\mathcal{B}}}^{\text{MLWE}}(\lambda).$$

*Proof (of Theorem 4).*

Let  $\mathcal{A}$  denote an adversary against MLWE-PRF. As outlined in the main text, the security proof follows in a series of hybrid games similar to the reduction in [BP14], where we start with the hybrid **Hybrid** $_{0,1}$  in which the view of  $\mathcal{A}$  is the same as in the pseudorandom ( $b = 1$ ) experiment MLWE-PRF $^1$  defined in Figure 4 with  $H = H^{\text{bp}}$ , and we end with the hybrid **Hybrid** $_{L-1,0}$ , in which the view of  $\mathcal{A}$  is the same as in the random ( $b = 0$ ) experiment MLWE-PRF $^0$ . For  $i \in \{0, \dots, L-1\}$  and  $\eta \in \{0, 1\}$ , We say that **Hybrid** $_{i,\eta}$  returns 1 if  $\mathcal{A}$  outputs  $b' = 1$ .

**Hybrid** $_{0,1}$ : This corresponds to the MLWE-PRF $^1$  game (Definition 11), and the interaction follows as in Figure 11.

MLWE-PRF $_{H, Q_P, Q_M, q, m, m\ell, L, \vec{\chi}}^b(\lambda)$	SampPRF( $x$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b>
2 : $\mathcal{Y} := \emptyset$ // Initialize empty table	2 : $(\mathbf{B}_x)^\top := (\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x))^\top \in R_q^{m \times m\ell}$
3 : $\mathcal{E} := \emptyset$ // Initialize empty table	3 : // $= \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$
4 : $b' \leftarrow \mathcal{A}^{\text{SampPRF}, \text{SampMLWE}}(1^\lambda)$	4 : // $= \mathbf{A}_{x_0} \cdot \mathbf{B}_{x/1}$
5 : <b>return</b> $b = b'$	5 : $\forall j \in \{0, \dots, L-1\}$ :
<b>SampMLWE</b> ()	6 : <b>if</b> $\mathcal{E}[x \setminus (j+1)] = \perp$ <b>then</b>
1 : $\mathbf{a} \xleftarrow{\$} R_q^m, e \xleftarrow{\$} \chi_0$	7 : $\mathcal{E}[x \setminus (j+1)] \xleftarrow{\$} \chi_0^{m\ell}$
2 : $c := \mathbf{k}^\top \mathbf{a} + e \in R_q$	8 : $\mathbf{e}_{x,j} := \mathcal{E}[x \setminus (j+1)]$
3 : <b>return</b> $(\mathbf{a}, c)$	9 : $\mathbf{e}_x^\top := \sum_{j=0}^{L-2} \mathbf{e}_{x,j}^\top \cdot \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$
	10 : // $= \mathbf{e}_{x_0}^\top \cdot \mathbf{B}_{x/1} + \mathbf{e}_{x/1}^\top$
	11 : $\mathcal{Y}[x]^\top := \mathbf{k}^\top \mathbf{B}_x + \mathbf{e}_x^\top \in R_q^{m\ell}$
	12 : // $= (\mathbf{k}^\top \mathbf{A}_{x_0} + \mathbf{e}_{x,0}^\top) \cdot \mathbf{B}_{x/1} + \mathbf{e}_{x/1}^\top$
	13 : <b>return</b> $\mathcal{Y}[x]$

Fig. 11: **Hybrid** $_{0,1}$  in which  $\mathcal{A}$ 's view is as in experiment MLWE-PRF $^1$

Since the view of  $\mathcal{A}$  in **Hybrid**<sub>0,1</sub> is the same as in the pseudorandom ( $b = 1$ ) experiment MLWE-PRF<sup>1</sup> defined in Figure 4 with  $\mathbf{H} = \mathbf{H}^{\text{bp}}$ , it follows that

$$\Pr[\mathbf{Hybrid}_{0,1} = 1] = \Pr[\text{MLWE-PRF}^1(\lambda) = 1].$$

**Hybrid**<sub>0,0</sub>: In this hybrid, we make the following two changes:

- In SampPRF, we replace  $(\mathbf{k}^\top \mathbf{A}_{x_0} + \mathbf{e}_{x_0}^\top)$  in the computation of  $\mathcal{Y}[x]$  with a uniformly random  $\mathbf{u}_{x_0}^\top$  in  $R_q^{m_l}$ , where the pair  $\mathbf{u}_{x_0}^\top$  for  $x_0 \in \{0, 1\}$  are independently sampled.
- In SampMLWE, we replace  $\mathbf{k}^\top \mathbf{a} + e$  in the computation of  $c$  with a uniform and independent ring element  $c \xleftarrow{\$} R_q$ .

This hybrid is shown in Figure 12.

MLWE-PRF <sub>H,Q_P,Q_M,q,m,l,L,\vec{\chi}}</sub> <sup>b</sup> ( $\lambda$ )	SampPRF( $x$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b>
2 : $\mathcal{Y} := \emptyset$ // Initialize empty table	2 : $(\mathbf{B}_x)^\top := (\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x))^\top \in R_q^{m \times m_l}$
3 : $\mathcal{E} := \emptyset$ // Initialize empty table	3 : // $= \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$
4 : $\mathcal{U} := \emptyset$ // Initialize empty table	4 : // $= \mathbf{A}_{x_0} \cdot \mathbf{B}_{x/1}$
5 : $b' \leftarrow \mathcal{A}^{\text{SampPRF, SampMLWE}}(1^\lambda)$	5 : $\forall j \in \{0, \dots, L-1\}$ :
6 : <b>return</b> $b = b'$	6 : <b>if</b> $\mathcal{E}[x \setminus (j+1)] = \perp$ <b>then</b>
<b>SampMLWE()</b>	7 : $\mathcal{E}[x \setminus (j+1)] \xleftarrow{\$} \chi_0^{m_l}$
1 : $\mathbf{a} \xleftarrow{\$} R_q^m, e \xleftarrow{\$} \chi_0$	8 : $\mathbf{e}_{x,j} := \mathcal{E}[x \setminus (j+1)]$
2 : $c \xleftarrow{\$} R_q$	9 : $\mathbf{e}_x^\top := \sum_{j=0}^{L-2} \mathbf{e}_{x,j}^\top \cdot \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$
3 : <b>return</b> $(\mathbf{a}, c)$	10 : // $= \mathbf{e}_{x_0}^\top \cdot \mathbf{B}_{x/1} + \mathbf{e}_{x/1}^\top$
	11 : <b>if</b> $\mathcal{U}[x_0] = \perp$ <b>then</b>
	12 : $\mathcal{U}[x_0] \xleftarrow{\$} R_q^{m_l}$
	13 : $\mathbf{u}_{x_0} := \mathcal{U}[x_0]$
	14 : $\mathcal{Y}[x]^\top := \mathbf{u}_{x_0}^\top \cdot \mathbf{B}_{x/1} + \mathbf{e}_{x/1}^\top \in R_q^{m_l}$
	15 : // $= \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$ .
	16 : <b>return</b> $\mathcal{Y}[x]$

Fig. 12: **Hybrid**<sub>0,0</sub> of the MLWE to MLWE-PRF reduction proof

We now construct an adversary  $\mathcal{B}$  against the MLWE<sub>2m\_l+Q\_M,q,m,1,\chi\_0,\chi\_k</sub> problem, such that

$$|\Pr[\mathbf{Hybrid}_{0,0} = 1] - \Pr[\mathbf{Hybrid}_{0,1} = 1]| = \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda). \quad (32)$$

On input MLWE instance  $(\mathbf{A}, \mathbf{b}) \in R_q^{(2m_l+Q_M) \times m} \times R_q^{2m_l+Q_M}$ , adversary  $\mathcal{B}$  parses them as  $\mathbf{A}^T := (\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_M) \in R_q^{m \times m_l} \times R_q^{m \times m_l} \times R_q^{m \times Q_M}$  and  $\mathbf{b}^T := (\mathbf{b}_0^\top, \mathbf{b}_1^\top, \mathbf{b}_M^\top) \in R_q^{m_l} \times R_q^{m_l} \times R_q^{Q_M}$ , and runs  $\mathcal{A}$  simulating its oracles as in **Hybrid**<sub>0,0</sub> but with the following modifications:

- In SampPRF, we replace assignment  $\mathcal{U}[x_0] \xleftarrow{\$} R_q^{m_l}$  in line 12 with the assignment  $\mathcal{U}[x_0] := \mathbf{b}_{x_0} \in R_q^{m_l}$ .
- In SampMLWE, for the  $i$ 'th oracle call of  $\mathcal{A}$  (with  $i \in [Q_M]$ ), we replace assignment  $c \xleftarrow{\$} R_q$  in line 2 with the assignment  $c := b_{M,i}$ , where  $b_{M,i} \in R_q$  denotes the  $i$ th  $R_q$ -coordinate of  $\mathbf{b}_M^\top \in R_q^{Q_M}$ .

When  $\mathcal{A}$  terminates and outputs  $b'$ ,  $\mathcal{B}$  terminates and outputs  $b'$ .

Observe that if the input MLWE instance comes from the ‘Real’ MLWE distribution ( $b = 1$ ) with secret vector  $\mathbf{s}$  sampled from  $\chi_k^m$  and error vector  $\mathbf{e}^\top = (\mathbf{e}_0^\top, \mathbf{e}_1^\top, \mathbf{e}_M^\top) \in R_q^{m_l} \times R_q^{m_l} \times R_q^{Q_M}$  sampled from  $\chi_e^{2m_l+Q_M}$ , then:

- In SampPRF queries,

$$\mathbf{u}_{x_0}^\top = \mathbf{b}_{x_0}^\top = (\mathbf{s}^\top \mathbf{A}_{x_0} + \mathbf{e}_{x,0}^\top),$$

so

$$\mathcal{Y}[x]^\top = (\mathbf{s}^\top \mathbf{A}_{x_0} + \mathbf{e}_{x,0}^\top) \cdot \mathbf{B}_{x/1} + \mathbf{e}_{x/1}^\top$$

- In the  $i$ 'th SampMLWE query,

$$c := b_{M,i} = \mathbf{s}^\top \mathbf{a}_{M,i} + e_{M,i},$$

where  $\mathbf{a}_{M,i}$  and  $e_{M,i}$  denote the  $i$ th column and coordinate of  $\mathbf{A}_M$  and  $\mathbf{e}_M^\top$ , respectively. Therefore, in this case,  $\mathcal{B}$  simulates to  $\mathcal{A}$  the same view as in **Hybrid<sub>0,1</sub>**, with  $\mathbf{k} := \mathbf{s}$ .

On the other hand, in the case that the input MLWE instance comes from the ‘Random’ MLWE distribution ( $b = 0$ ), then  $\mathbf{u}_{x_0}^\top = \mathbf{b}_{x_0}^\top$  are uniform in  $R_q^{m_l}$  and independent for  $x_0 \in \{0, 1\}$  and  $b_{M,i}$  are also uniform in  $R_q$  and independent. Therefore, in this case,  $\mathcal{B}$  simulates to  $\mathcal{A}$  the same view as in **Hybrid<sub>0,0</sub>**.

The above two observations establish the claim (32).

Starting with  $i = 1$ , we now continue inductively to similarly define  $(L-1)$  pairs of hybrids (**Hybrid<sub>i,1</sub>**, **Hybrid<sub>i,0</sub>**) for  $i = 1, \dots, L-1$  as follows.

**Hybrid<sub>i,1</sub>**: Compared to the previous **Hybrid<sub>i-1,0</sub>** (i.e. **Hybrid<sub>0,0</sub>** in the induction base case  $i = 1$ ), we make the following syntactic change:

- In SampPRF, we rewrite  $\mathbf{u}_{x \setminus i}^\top$  as  $\mathbf{u}_{x \setminus i}^\top := \mathbf{k}_{x \setminus i}^\top \cdot \mathbf{G} + \mathbf{v}_{x \setminus i}^\top$ , where  $\mathbf{k}_{x \setminus i}^\top$  is sampled uniformly at random from  $R_q^m$  (i.e. the distribution  $\chi_k^m$ ) and  $\mathbf{v}_{x \setminus i}^\top$  is sampled uniformly at random from  $R_q^{m_l}/(\mathbf{G}R_q^m)$ . (recall that  $x \setminus i := x_0 x_1 \dots, x_{i-1}$  denotes the  $i$ -bit prefix of  $x$  and is equal to  $x_0$  in the induction base case  $i = 1$ ).

This hybrid is shown in Figure 13.

Assume the induction hypothesis for  $i \geq 1$ , namely that in the previous game **Hybrid<sub>i-1,0</sub>**, the SampPRF output has the form

$$\mathcal{Y}[x]^\top := \mathbf{u}_{x \setminus i}^\top \cdot \mathbf{B}_{x/i} + \mathbf{e}_{x/i}^\top + \mathbf{v}_{x,i-1}^\top \in R_q^{m_l}, \quad (33)$$

with  $\mathbf{u}_{x \setminus i} := \mathcal{U}[x \setminus i]$  uniform in  $R_q^{m_l}$  and independent for distinct  $i$ -bit prefixes  $x \setminus i$  of  $x$  (note that this is satisfied in the induction base case  $i = 1$  by the game **Hybrid<sub>0,0</sub>** with  $\mathbf{v}_{x,0}^\top := \mathbf{0}$ ). Then, the change in the game **Hybrid<sub>i,1</sub>** is syntactic only, since the distribution of  $\mathbf{u}_{x \setminus i} := \mathbf{k}_{x \setminus i}^\top \cdot \mathbf{G} + \mathbf{v}_{x \setminus i}^\top$  in this game is still uniform in  $R_q^{m_l}$  and independent for distinct  $i$ -bit prefixes  $x \setminus i$  of  $x$ , due to the uniformly random choice of  $\mathbf{k}_{x \setminus i}^\top$  in  $R_q^m$  and  $\mathbf{v}_{x \setminus i}^\top$  in  $R_q^{m_l}/(\mathbf{G}R_q^m)$ . Therefore, the view of  $\mathcal{A}$  stays the same as in the previous game, and we have

$$|\Pr[\mathbf{Hybrid}_{i,1} = 1] - \Pr[\mathbf{Hybrid}_{i-1,0} = 1]| = 0. \quad (34)$$

Furthermore, in this game **Hybrid<sub>i,1</sub>**, the SampPRF output has the form

$$\mathcal{Y}[x]^\top := (\mathbf{k}_{x \setminus i}^\top \mathbf{A}_{x_i} + \mathbf{e}_{x,i}^\top) \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top + \mathbf{v}_{x,i}^\top, \quad (35)$$

since  $\mathbf{B}_{x/i} := \mathbf{G}^{-1}(\mathbf{A}_{x_i} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$  so  $\mathbf{k}_{x \setminus i}^\top \mathbf{G} \mathbf{B}_{x/i} = \mathbf{k}_{x \setminus i}^\top \mathbf{A}_{x_i} \mathbf{B}_{x/(i+1)}$ ,  $\mathbf{e}_{x/i}^\top := \mathbf{e}_{x,i}^\top \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top$  and  $\mathbf{v}_{x,i}^\top := \mathbf{v}_{x \setminus i}^\top \mathbf{B}_{x/i} + \mathbf{v}_{x,i-1}^\top$ .

**Hybrid<sub>i,0</sub>**: In this hybrid, we make the following change:

MLWE-PRF $_{\mathbb{H}, Q_P, Q_M, q, m, m_l, L, \vec{\chi}}^b(\lambda)$	SampPRF( $x$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b>
2 : $\mathcal{Y} := \emptyset$ // Initialize empty table	2 : $(\mathbf{B}_x)^\top := (\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x))^\top \in R_q^{m \times m_l}$
3 : $\mathcal{E} := \emptyset$ // Initialize empty table	3 : // $= \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$
4 : $\mathcal{U} := \emptyset$ // Initialize empty table	4 : // $= \mathbf{A}_{x_0} \cdot \mathbf{B}_{x/1}$
5 : $\mathcal{K} := \emptyset$ // Initialize empty table	5 : $\forall j \in \{0, \dots, L-1\}$ :
6 : $\mathcal{V} := \emptyset$ // Initialize empty table	6 : <b>if</b> $\mathcal{E}[x \setminus (j+1)] = \perp$ <b>then</b>
7 : $b' \leftarrow \mathcal{A}^{\text{SampPRF, SampMLWE}}(1^\lambda)$	7 : $\mathcal{E}[x \setminus (j+1)] \xleftarrow{\$} \chi_0^{m_l}$
8 : <b>return</b> $b = b'$	8 : $\mathbf{e}_{x,j} := \mathcal{E}[x \setminus (j+1)]$
<b>SampMLWE()</b>	9 : $\mathbf{e}_{x/i}^\top := \sum_{j=i}^{L-2} \mathbf{e}_{x,j}^\top \cdot \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$
1 : $\mathbf{a} \xleftarrow{\$} R_q^m, e \xleftarrow{\$} \chi_0$	10 : // $= \mathbf{e}_{x,i}^\top \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top$
2 : $c \xleftarrow{\$} R_q$	11 : $\forall j \in \{1, \dots, L-1\}$ :
3 : <b>return</b> $(\mathbf{a}, c)$	12 : <b>if</b> $\mathcal{V}[x \setminus j] = \perp$ <b>then</b>
	13 : $\mathcal{V}[x \setminus j] \xleftarrow{\$} R_q^{m_l} / (\mathbf{G} R_q^m)$
	14 : $\mathbf{v}_{x \setminus j} := \mathcal{V}[x \setminus j]$
	15 : $\mathbf{v}_{x,j}^\top := \mathbf{v}_{x \setminus j}^\top \mathbf{B}_{x/j} + \mathbf{v}_{x,j-1}^\top$
	16 : <b>if</b> $\mathcal{U}[x \setminus i] = \perp$ <b>then</b>
	17 : $\mathbf{k}_{x \setminus i}^\top := \mathcal{K}[x \setminus i]^\top \xleftarrow{\$} R_q^m$
	18 : $\mathcal{U}[x \setminus i]^\top := \mathbf{k}_{x \setminus i}^\top \cdot \mathbf{G} + \mathbf{v}_{x \setminus i}^\top$
	19 : $\mathbf{u}_{x \setminus i}^\top := \mathcal{U}[x \setminus i]^\top$
	20 : $\mathcal{Y}[x]^\top := \mathbf{u}_{x \setminus i}^\top \cdot \mathbf{B}_{x/i} + \mathbf{e}_{x/i}^\top + \mathbf{v}_{x,i-1}^\top \in R_q^{m_l}$
	21 :     // $= (\mathbf{k}_{x \setminus i}^\top \mathbf{A}_{x_i} + \mathbf{e}_{x,i}^\top) \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top + \mathbf{v}_{x,i}^\top$
	22 : <b>return</b> $\mathcal{Y}[x]$

Fig. 13: **Hybrid<sub>1,1</sub>** of the MLWE to MLWE-PRF reduction proof

MLWE-PRF <sub>H,Q_P,Q_M,q,m,m_l,L,\vec{\chi}</sub> <sup>b</sup> ( $\lambda$ )	SampPRF( $x$ )
1 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$	1 : <b>if</b> $\mathcal{Y}[x] = \perp$ <b>then</b>
2 : $\mathcal{Y} := \emptyset$ // Initialize empty table	2 : $(\mathbf{B}_x)^\top := (\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{bp}}(x))^\top \in R_q^{m \times m_l}$
3 : $\mathcal{E} := \emptyset$ // Initialize empty table	3 : $\parallel = \mathbf{A}_{x_0} \cdot \mathbf{G}^{-1}(\mathbf{A}_{x_1} \cdots \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}))$
4 : $\mathcal{U} := \emptyset$ // Initialize empty table	4 : $\parallel = \mathbf{A}_{x_0} \cdot \mathbf{B}_{x/1}$
5 : $\mathcal{K} := \emptyset$ // Initialize empty table	5 : $\forall j \in \{0, \dots, L-1\}$ :
6 : $\mathcal{V} := \emptyset$ // Initialize empty table	6 : <b>if</b> $\mathcal{E}[x \setminus (j+1)] = \perp$ <b>then</b>
7 : $b' \leftarrow \mathcal{A}^{\text{SampPRF, SampMLWE}}(1^\lambda)$	7 : $\mathcal{E}[x \setminus (j+1)] \xleftarrow{\$} \chi_0^{m_l}$
8 : <b>return</b> $b = b'$	8 : $\mathbf{e}_{x,j} := \mathcal{E}[x \setminus (j+1)]$
<b>SampMLWE</b> ( $\lambda$ )	9 : $\mathbf{e}_{x/i}^\top := \sum_{j=i}^{L-2} \mathbf{e}_{x,j}^\top \cdot \mathbf{B}_{x/(j+1)} + \mathbf{e}_{x,L-1}^\top$
1 : $\mathbf{a} \xleftarrow{\$} R_q^m, e \xleftarrow{\$} \chi_0$	10 : $\parallel = \mathbf{e}_{x,i}^\top \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top$
2 : $c \xleftarrow{\$} R_q$	11 : $\forall j \in \{1, \dots, L-1\}$ :
3 : <b>return</b> $(\mathbf{a}, c)$	12 : <b>if</b> $\mathcal{V}[x \setminus j] = \perp$ <b>then</b>
	13 : $\mathcal{V}[x \setminus j] \xleftarrow{\$} R_q^{m_l} / (\mathbf{G} R_q^m)$
	14 : $\mathbf{v}_{x \setminus j} := \mathcal{V}[x \setminus j]$
	15 : $\mathbf{v}_{x,j}^\top := \mathbf{v}_{x \setminus j}^\top \mathbf{B}_{x/j} + \mathbf{v}_{x,j-1}^\top$
	16 : <b>if</b> $\mathcal{U}[x \setminus (i+1)] = \perp$ <b>then</b>
	17 : $\mathcal{U}[x \setminus (i+1)] \xleftarrow{\$} R_q^{m_l}$
	18 : $\mathbf{u}_{x \setminus (i+1)} := \mathcal{U}[x \setminus (i+1)]$
	19 : $\mathcal{Y}[x]^\top := \mathbf{u}_{x \setminus (i+1)}^\top \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top + \mathbf{v}_{x,i}^\top \in R_q^{m_l}$
	20 : <b>return</b> $\mathcal{Y}[x]$

Fig. 14: **Hybrid**<sub>1,0</sub> of the MLWE to MLWE-PRF reduction proof

- In **SampPRF**, we replace  $(\mathbf{k}_{x \setminus i}^\top \mathbf{A}_{x_i} + \mathbf{e}_{x,i}^\top)$  in the computation of  $\mathcal{Y}[x]$  with a uniformly random  $\mathbf{u}_{x \setminus (i+1)}^\top$  in  $R_q^{m_l}$ , where the  $\mathbf{u}_{x \setminus (i+1)}^\top$  is independently sampled for each distinct  $i$ -bit prefix  $x \setminus (i+1) \in \{0,1\}^i$ .

This hybrid is shown in Figure 14.

Similarly to **Hybrid**<sub>0,0</sub>, we now construct an adversary  $\mathcal{B}_i$  against the  $\text{MLWE}_{2m_l, q, m, Q_P, \chi_0, \chi_k}$  problem, such that

$$|\Pr[\mathbf{Hybrid}_{i,0} = 1] - \Pr[\mathbf{Hybrid}_{i,1} = 1]| = \text{Adv}_{\mathcal{B}_i}^{\text{MLWE}}(\lambda). \quad (36)$$

On input MLWE instance  $(\mathbf{A}, \mathbf{B}) \in R_q^{2m_l \times m} \times R_q^{2m_l \times Q_P}$ , adversary  $\mathcal{B}_i$  parses them as  $\mathbf{A}^T := (\mathbf{A}_0, \mathbf{A}_1) \in R_q^{m \times m_l} \times R_q^{m \times m_l}$  and  $\mathbf{B}^T := (\mathbf{B}_0^\top, \mathbf{B}_1^\top) \in R_q^{Q_P \times m_l} \times R_q^{Q_P \times m_l}$ , and runs  $\mathcal{A}$  simulating its oracles as in **Hybrid**<sub>i,0</sub> but with the following modification:

- In **SampPRF**, we replace assignment  $\mathcal{U}[x \setminus (i+1)] \stackrel{\S}{\leftarrow} R_q^{m_l}$  in line 17 with the assignment  $\mathcal{U}[x \setminus (i+1)] := \mathbf{B}_{x_i}^\top[\text{qind}(x \setminus i)] \in R_q^{m_l}$ , where  $\mathbf{B}_{x_i}^\top[j] \in R_q^{m_l}$  denotes the  $j$ th row of  $\mathbf{B}_{x_i}^\top$  for  $j \in [Q_P]$  and  $\text{qind}(x \setminus i)$  returns the index  $j \in [Q_P]$  of the first query  $x'$  to **SampPRF** which has the same  $i$ -bit prefix as  $x$  (i.e.  $x' \setminus i = x \setminus i$ ).

When  $\mathcal{A}$  terminates and outputs  $b'$ ,  $\mathcal{B}$  terminates and outputs  $b'$ .

Observe that if the input MLWE instance comes from the ‘Real’ (multi-secret) MLWE distribution ( $b = 1$ ), then for  $j \in [Q_P]$ , the  $j$ th row of  $\mathbf{B}_{x_i}^\top$  has the form  $\mathbf{B}_{x_i}^\top[j] = \mathbf{S}^\top[j] \mathbf{A}_{x_i} + \mathbf{E}_{x_i}^\top[j]$ , where the  $j$ th MLWE secret vector  $\mathbf{S}^\top[j]$  is independently sampled from  $\chi_k^m$ , and the  $j$ th error vectors  $\mathbf{E}_0^\top[j], \mathbf{E}_1^\top[j]$  are independently sampled from  $\chi_e^{m_l}$ , and therefore for **SampPRF** query  $x$ ,

$$\mathbf{u}_{x \setminus (i+1)}^\top = \mathbf{B}_{x_i}^\top[\text{qind}(x \setminus i)] = \mathbf{S}^\top[\text{qind}(x \setminus i)] \mathbf{A}_{x_i} + \mathbf{E}_{x_i}^\top[\text{qind}(x \setminus i)],$$

so

$$\mathcal{Y}[x]^\top = (\mathbf{S}^\top[\text{qind}(x \setminus i)] \mathbf{A}_{x_i} + \mathbf{E}_{x_i}^\top[\text{qind}(x \setminus i)]) \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top + \mathbf{v}_{x,i}^\top$$

Therefore, in this case, comparing with (35), we see that  $\mathcal{B}$  simulates to  $\mathcal{A}$  the same view as in **Hybrid**<sub>i,1</sub>, with  $\mathbf{k}_{x \setminus i}^\top := \mathbf{S}^\top[\text{qind}(x \setminus i)]$ .

On the other hand, in the case that the input MLWE instance comes from the ‘Random’ MLWE distribution ( $b = 0$ ), then  $\mathbf{u}_{x \setminus (i+1)}^\top = \mathbf{B}_{x_i}^\top[\text{qind}(x \setminus i)]$  are uniform in  $R_q^{m_l}$  and independent for distinct  $(i+1)$ -bit prefixes  $x \setminus (i+1) \in \{0,1\}^{i+1}$ . Therefore, in this case,  $\mathcal{B}$  simulates to  $\mathcal{A}$  the same view as in **Hybrid**<sub>i,0</sub>.

The above two observations establish the claim (36). Furthermore, observe that in this game **Hybrid**<sub>i,1</sub>, the **SampPRF** output has the form

$$\mathcal{Y}[x]^\top := \mathbf{u}_{x \setminus (i+1)}^\top \cdot \mathbf{B}_{x/(i+1)} + \mathbf{e}_{x/(i+1)}^\top + \mathbf{v}_{x,i}^\top \in R_q^{m_l}, \quad (37)$$

which establishes the induction hypothesis (33) for  $i+1$ . Applying the induction  $L-1$  times with  $i = 1, \dots, L-1$ , we arrive at the final hybrid **Hybrid**<sub>L-1,0</sub>, in which the output of **SampPRF** on query  $x$  is

$$\mathcal{Y}[x]^\top := \mathbf{u}_{x \setminus L}^\top \cdot \mathbf{B}_{x/L} + \mathbf{e}_{x/L}^\top + \mathbf{v}_{x,L-1}^\top \quad (38)$$

$$= \mathbf{u}_x^\top + \mathbf{v}_{x,L-1}^\top \in R_q^{m_l}, \quad (39)$$

which is uniformly random in  $R_q^{m_l}$  and independent for distinct  $x \in \{0,1\}^L$ , since  $\mathbf{u}_x^\top$  is uniformly random in  $R_q^{m_l}$  and independent of  $\mathbf{v}_{x,L-1}^\top$  (where, on the second line of (38), we have used the fact that  $x \setminus L := x$ ,  $\mathbf{B}_{x/L} := \mathbf{I}$ , and  $\mathbf{e}_{x/L}^\top = \mathbf{0}$ ). Therefore, the view of  $\mathcal{A}$  in **Hybrid**<sub>L-1,0</sub> is identical to  $\mathcal{A}$ ’s view in the random ( $b = 0$ ) experiment  $\text{MLWE-PRF}^0$ , as required. This completes the description of the  $L$  pairs of hybrids.

Putting together the bounds (32), and (36), (34) for  $i = 1 \dots, L-1$ , we get:

$$\text{Adv}_{\mathcal{A}}^{\text{MLWE-PRF}}(\lambda) = |\Pr[\mathbf{Hybrid}_{L-1,0} = 1] - \Pr[\mathbf{Hybrid}_{0,1} = 1]| \quad (40)$$

$$\leq \sum_{i=0}^{L-1} |\Pr[\mathbf{Hybrid}_{i,1} = 1] - \Pr[\mathbf{Hybrid}_{i,0} = 1]| \quad (41)$$

$$\leq \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda) + \sum_{i=1}^{L-1} \text{Adv}_{\mathcal{B}_i}^{\text{MLWE}}(\lambda) \quad (42)$$

$$\leq \text{Adv}_{\mathcal{B}}^{\text{MLWE}}(\lambda) + (L-1) \cdot \text{Adv}_{\overline{\mathcal{B}}}^{\text{MLWE}}(\lambda), \quad (43)$$

$$(44)$$

where  $\overline{\mathcal{B}}$  denotes the adversary among  $\mathcal{B}_1, \dots, \mathcal{B}_{L-1}$  with the highest advantage. This completes the proof.  $\square$

## D.2 Correctness Analysis

The OPRF correctness analysis remains unaffected whether a part of the PRF input  $x$  in the computation of  $\mathbf{H}(x)$  is public or not. That's why we just write  $\mathbf{B}_x := \mathbf{H}(x)$  in the correctness analysis to keep the notation a bit simpler.

**Definition 19.** We say that function family  $\mathbf{H}$  satisfies  $\epsilon_u$ -uniformity if, for each fixed  $x \in \{0,1\}^L$ , the distribution of any fixed  $\mathbb{Z}_q$  coordinate of  $\mathbf{B}_x \mathbf{k} := \mathbf{H}(x) \mathbf{k} \in R_q^h$  over the choice of  $\mathbf{H}$  (in  $\mathbf{F.Setup}(1^\lambda)$ ) and  $\mathbf{k} \xleftarrow{\$} \chi_k^m$  is within statistical distance

$$\epsilon_u \leq 2^{-(\kappa+2+\log(hd))} \quad (45)$$

from the uniform distribution on  $\mathbb{Z}_q$ .

**Lemma 7.** Fix  $\kappa, h, d$ . Let  $B_f(\kappa, d)$  denote an upper bound on a fixed coordinate of  $\mathbf{e}_f = \mathbf{e}'_s - \mathbf{R}\mathbf{e}_s$  (as in (49)) that holds except with probability  $p_e \leq 2^{-(\kappa+2+\log(hd))}$ . Also, assume that the function family  $\mathbf{H}$  satisfies  $\epsilon_u$ -uniformity (as per Definition 19) with

$$\epsilon_u \leq 2^{-(\kappa+2+\log(hd))}. \quad (23)$$

Then, for any fixed  $x \in \{0,1\}^L$ ,  $2^{-\kappa}$ -correctness holds (as per Definition 2) if

$$q/p \geq 2^{\kappa+2} \cdot hd \cdot (2B_f(\kappa, d) + 1). \quad (24)$$

*Proof.* Fix  $(t, x)$ , and let  $E$  and  $\delta$  denote respectively the event that a correctness error occurs and the correctness error probability. We have

$$\delta := \Pr[E] := \Pr[\lfloor \mathbf{u}_x - \mathbf{R}\mathbf{v}_k \rfloor_p \neq \lfloor \mathbf{B}_x \mathbf{k} \rfloor_p] \quad (46)$$

$$= \Pr[\lfloor \mathbf{B}_x \mathbf{k} + \mathbf{e}_f \rfloor_p \neq \lfloor \mathbf{B}_x \mathbf{k} \rfloor_p] \quad (47)$$

$$(48)$$

where the final error  $\mathbf{e}_f$  satisfies

$$\mathbf{e}_f := \mathbf{u}_x - \mathbf{R}\mathbf{v}_k - \mathbf{B}_x \mathbf{k} = \mathbf{e}'_s - \mathbf{R}\mathbf{e}_s. \quad (49)$$

For  $i \in [h]$  and  $j \in [d]$ , let  $\mathbf{e}_{f,i} \in R$  denote the  $i$ th ring element of  $\mathbf{e}_f \in R^h$  and let  $e_{f,i,j} \in \mathbb{Z}$  denote the  $j$ 'th integer coefficient of  $\mathbf{e}_{f,i}$  in the coefficient embedding of  $R$ . Similarly, viewing  $\mathbf{B}_x$  in its representation over  $\mathbb{Z}$ , we denote by  $\mathbf{b}_{x,i,j}^\top \in \mathbb{Z}_q^{md}$  the  $j$ 'th row of the column rot (negacyclic) matrix corresponding to  $\mathbf{B}_x$ . Let  $B_p \subset [0, q-1]$  denote the set of rounding mod  $p$  interval boundary points, i.e.  $B_p = \{\frac{q}{2p}, \frac{q}{2p} + \frac{q}{p}, \dots, \frac{q}{2p} + (p-1) \cdot \frac{q}{2p}\}$ . If event  $E$  occurs, then there must exist some  $i \in [h]$  and  $j \in [d]$  such that  $\mathbf{b}_{x,i,j}^\top \mathbf{k} \in \mathbb{Z}_q$  lands within distance  $|e_{f,i,j}|$  of a rounding interval boundary point in  $B_p$ . Therefore, we have

$$\begin{aligned} \delta &\leq \Pr[\exists i \in [h], j \in [d] \text{ s.t. } \mathbf{b}_{x,i,j}^\top \mathbf{k} \in B_p \pm |e_{f,i,j}|] \\ &\leq \sum_{i \in [h], j \in [d]} \delta_{i,j}, \text{ where } \delta_{i,j} := \Pr[\mathbf{b}_{x,i,j}^\top \mathbf{k} \in B_p \pm |e_{f,i,j}|]. \end{aligned} \quad (50)$$

Now, for each fixed  $i \in [h]$  and  $j \in [d]$ , we have

$$\begin{aligned}
\delta_{i,j} &\leq \Pr_{u \stackrel{\$}{\leftarrow} \mathbb{Z}_q} [u \in B_p \pm |e_{f,i,j}|] + \epsilon_u \\
&\leq \Pr_{u \stackrel{\$}{\leftarrow} \mathbb{Z}_q} [u \in B_p \pm B_f(\kappa, d)] + p_e + \epsilon_u \\
&= \frac{|(B_p \pm B_f(\kappa, d)) \cap \mathbb{Z}_q|}{|\mathbb{Z}_q|} + p_e + \epsilon_u \\
&\leq \frac{p \cdot (2B_f(\kappa, d) + 1)}{q} + p_e + \epsilon_u \\
&\leq 2^{-\kappa}/(hd),
\end{aligned}$$

where the first inequality above uses the assumed  $\epsilon_u$ -uniformity of  $\mathbf{H}$ , the second inequality uses the bound  $|e_{f,i,j}| \leq B_f(\kappa, d)$  that holds except with probability  $\leq p_e$ , the third inequality uses  $|B_p| = p$  and the last inequality uses the assumed upper bounds  $\leq 2^{-\kappa}/(3hd)$  on the three terms on the right-hand side of the third inequality. It follows that  $\delta \leq \sum_{i \in [h], j \in [d]} \delta_{i,j} \leq 2^{-\kappa}$ , as claimed.  $\square$

**Bounds for BLMR13 PRF Instantiation of  $\mathbf{H}$ .** For BLMR13 (see [SSS23] for the module instantiation), we have  $\chi_k := \mathcal{U}(R_q)$ ,  $\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{blmr}}(\mathbf{x}) := \prod_{i=0}^{L-1} \mathbf{A}_{x_i}$ , where we sample  $\mathbf{A}_0, \mathbf{A}_1 \stackrel{\$}{\leftarrow} R_{\text{bin}}^{m \times m}$  and restart if  $\mathbf{A}_0$  or  $\mathbf{A}_1$  are not invertible over  $R_q$ .

**Lemma 8.** *For BLMR13, the function family  $\mathbf{H}^{\text{blmr}}$  satisfies 0-uniformity. In particular, the distribution of  $\mathbf{B}_x^{\text{blmr}} \cdot \mathbf{k} := \mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{blmr}}(x) \cdot \mathbf{k} \in R_q^{m \times m}$  over the choice of  $\mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{blmr}}$  and  $\mathbf{k} \stackrel{\$}{\leftarrow} \chi_k^m$  is the uniform distribution on  $R_q^{m \times m}$ . The result carries over to the case  $\mathbf{H}(x) := \mathbf{B}_x^{\text{blmr}}[0]$  (as in  $\mathbf{F}.\text{Setup}(1^\lambda)$ ).*

*Proof.* Follows immediately from the uniformly random distribution  $\chi_k$  and the invertibility of the matrix  $\mathbf{B}_x^{\text{blmr}} = \mathbf{H}_{\mathbf{A}_0, \mathbf{A}_1}^{\text{blmr}}(x) = \prod_{i=0}^{L-1} \mathbf{A}_{x_i}$  over  $R_q$ , thanks to the invertibility of  $\mathbf{A}_0$  and  $\mathbf{A}_1$  over  $R_q$ . The result carries over if we simply focus on the first coordinate of  $\mathbf{B}_x^{\text{blmr}} \cdot \mathbf{k}$ , i.e.,  $\mathbf{B}_x^{\text{blmr}}[0] \cdot \mathbf{k}$ .  $\square$

**Bounds for BP14 PRF Instantiation of  $\mathbf{H}$ .** For BP14 with  $\chi_k := \mathcal{U}(R_q)$ ,  $\mathbf{B}_x = \mathbf{H}(\mathbf{x}) := (\mathbf{A}_{x_0} \mathbf{B}'_1)^\top \in R_q^{1 \times m}$ , and  $\mathbf{l} := \lceil \log_\gamma q \rceil$  with  $\gamma \geq 2$ ,  $\mathbf{B}_{L-1} := \mathbf{G}^{-1}(\mathbf{A}_{x_{L-1}}) \in R_q^{m \times m \mathbf{l}}$ ,  $\mathbf{B}_i := \mathbf{G}^{-1}(\mathbf{A}_{x_i} \mathbf{B}_{i+1}) \in R_q^{m \mathbf{l} \times m \mathbf{l}}$  for  $i = L-2, \dots, 1$ .  $\mathbf{B}'_1$  consists of the leftmost  $h$  columns of  $\mathbf{B}_1$ , and  $\mathbf{F}.\text{Setup}(1^\lambda)$  samples  $\mathbf{A}_0, \mathbf{A}_1 \stackrel{\$}{\leftarrow} R_q^{m \times m \mathbf{l}}$ .

**Lemma 9.** *Assume that  $q$  is prime. If the (truncated) BP14 function family  $\mathbf{H}$  does not satisfy  $\epsilon_u$ -uniformity for some non-negligible  $\epsilon_u$ , then there exists a poly-time 1-query adversary against the pseudorandomness of the BP14 PRF with non-negligible advantage  $\geq (1 - 1/p) \cdot \epsilon_u$ .*

*Proof.* By hypothesis, there exists  $x \in \{0, 1\}^L$ ,  $i \in [h]$  and  $j \in [d]$  such that the statistical distance  $\epsilon$  between the distribution of  $\mathbf{b}_{x,i,j}^\top \cdot \mathbf{k} \in \mathbb{Z}_q$  and  $\mathcal{U}(\mathbb{Z}_q)$  is  $\geq \epsilon_u$ , where  $\mathbf{b}_{x,i,j}^\top$  denotes the  $j$ 'th row of the column rot (negacyclic) matrix representation of  $\mathbf{B}_x$  over  $\mathbb{Z}$  of the  $i$ 'th row of  $\mathbf{B}_x \in R_q^{h \times m}$ , and similarly  $\mathbf{k}$  is viewed over  $\mathbb{Z}$  as the concatenation of ring element coefficient vectors. Let  $\mathbf{Z}$  denote the event that  $\mathbf{b}_{x,i,j}^\top = \mathbf{0}$ . Then we have  $\epsilon = \epsilon_{nz} \cdot (1 - \Pr[\mathbf{Z}]) + \epsilon_z \cdot \Pr[\mathbf{Z}]$ , where  $\epsilon_{nz}$  (resp.  $\epsilon_z$ ) denotes the statistical distance between the distribution of  $\mathbf{b}_{x,i,j}^\top \cdot \mathbf{k} \in \mathbb{Z}_q$  and  $\mathcal{U}(\mathbb{Z}_q)$  conditioned on the event that  $\mathbf{Z}$  does not occur (resp.  $\mathbf{Z}$  occurs). Conditioned on  $\mathbf{Z}$  not occurring, we have  $\mathbf{b}_{x,i,j}^\top \neq \mathbf{0}$ , and thus  $\mathbf{b}_{x,i,j}^\top$  has a non-zero and hence (by primality of  $q$ ) invertible component in  $\mathbb{Z}_q$ . Thanks to the independence and uniformity of  $\mathbf{k}$ 's components in  $\mathbb{Z}_q$ , it follows that in this case  $\mathbf{b}_{x,i,j}^\top \cdot \mathbf{k}$  is uniformly distributed in  $\mathbb{Z}_q$ , so  $\epsilon_{nz} = 0$ . Therefore  $\epsilon = \epsilon_z \cdot \Pr[\mathbf{Z}] \leq \Pr[\mathbf{Z}]$ . It follows that  $\Pr[\mathbf{Z}] \geq \epsilon \geq \epsilon_u$  is non-negligible. Then there exists a pseudorandomness adversary  $\mathcal{A}$  against the BP14 PRF that works as follows: it queries  $x$  to its oracle to get  $\mathbf{z} \in R_p^h$ . If  $\mathbf{Z}$  occurs then  $\mathcal{A}$  outputs 1 if  $z_{i,j} = 0$  and output 0 else. If  $\mathbf{Z}$  does not occur,  $\mathcal{A}$  outputs a random coin. In the case that  $\mathcal{A}$ 's oracle is the BP14 PRF, if the event  $\mathbf{Z}$  occurs, the oracle will output 0 with probability 1 and hence  $\mathcal{A}$  will output 1 with probability  $p_P = \Pr[\mathbf{Z}] + 1/2(1 - \Pr[\mathbf{Z}])$ . In the other case that  $\mathcal{A}$ 's oracle is a uniformly random function with range  $R_p^h$ , the event  $z_{i,j} = 0$  will occur with probability  $1/p$  independently of event  $\mathbf{Z}$ , and hence in this case

$\mathcal{A}$  will output 1 with probability  $p_U = \Pr[Z] \cdot 1/p + 1/2(1 - \Pr[Z])$ . It follows that the distinguishing advantage of  $\mathcal{A}$  against the BP14 PRF security is lower bounded as  $|p_P - p_U| \geq (1 - 1/p) \cdot \Pr[Z] > (1 - 1/p) \cdot \epsilon_u$ , as claimed.  $\square$

*Remark 7.* We note that the  $\epsilon_u$ -uniformity Lemma above assumes  $\chi_k$  has uniformly random coordinates in  $\mathbb{Z}_q$ , but the BP14 PRF could also be instantiated with  $\chi_k$  restricted to small coefficients. In the latter case, the  $\epsilon_u$ -uniformity seems more difficult to prove but we heuristically expect it to still be satisfied.

*Remark 8.* We note that our correctness analysis extends to a *malicious* server setting by relying on the fact that the NIZK proof  $\pi_s$  by the server ensures that the error terms  $(\mathbf{e}_s, \mathbf{e}'_s)$  have small coefficients (say bounded by  $\beta$  and  $\beta_1$ , respectively). Particularly, using the worst-case bound in (25) for  $B_f$  gives the correctness analysis requirement against a malicious server.

### D.3 Pseudorandomness Analysis

**Theorem 5.** *The POPRF construction  $\mathbf{F}$  from Figures 8 and 9 satisfies pseudorandomness given in Definition 3, with random oracles  $\text{RO}_r$  and  $\text{RO}_z$ , if:*

- The client argument system  $\text{NIZK}_1$  is computationally sound and the server argument system  $\text{NIZK}_2$  is computationally zero-knowledge (Definition 18),
- The commitment scheme  $\text{COM}$  is computationally hiding and extractable (Definitions 15 and 17), and
- The  $\text{MLWE-PRF-RU}_{\text{param}}$  assumption (Definition 10) holds for  $\text{param} := (\mathbf{H}, Q, Q_{x,t}^\infty, Q_M, q, m, \ell + m, h, L, \beta_r, \vec{\chi})$  and  $\vec{\chi} := (\chi, \chi_1, \chi_k)$ , where  $Q$  corresponds to the number of  $\text{BlindEval}$  queries,  $Q_{x,t}^\infty$  to the maximum number of identical  $(x||t)$  queries to  $\text{BlindEval}$ , and  $Q_M$  to the number of  $\text{RO}_r$  queries, respectively.

More precisely, for any PPT adversary  $\mathcal{A}$ , there exist PPT adversaries  $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$  and  $\mathcal{B}_6$  against the computational zero-knowledge of  $\text{NIZK}_2$ , computational soundness of  $\text{NIZK}_1$ , hiding and extractability of  $\text{COM}$ , and  $\text{MLWE-PRF-RU}_{\text{param}}$  assumption, respectively, such that

$$\begin{aligned} \text{Adv}_{\mathbf{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^{\text{po-prf}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_2, \mathcal{B}_1}^{\text{CZK}}(\lambda) + Q_B \cdot \text{Adv}_{\text{NIZK}_1, \mathcal{B}_2}^{\text{CS}}(\lambda) + \text{Adv}_{\text{COM}, \mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + Q_{\text{RO}_r} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_4}^{\text{Ext}}(\lambda) + Q_B \cdot \text{Adv}_{\text{COM}, \mathcal{B}_5}^{\text{Ext}}(\lambda) + \text{Adv}_{\mathcal{B}_6}^{\text{iMLWE}_{\text{param}}}(\lambda) \\ &\quad + Q_{\text{RO}_z} \left( \frac{1}{p} + \frac{1}{q} \right)^{dh}, \end{aligned}$$

where  $Q_B, Q_{\text{RO}_r}$  and  $Q_{\text{RO}_z}$  denote the number of  $\text{BlindEval}, \text{RO}_r$  and  $\text{RO}_z$  queries, respectively, that the adversary  $\mathcal{A}$  makes.

*Proof.* We consider the POPRF pseudorandomness game given in Definition 3.

We note that the adversary  $\mathcal{A}$  has access to oracles  $\text{Eval}, \text{PreProcServer}, \text{BlindEval}$  and  $\text{Prim}$ , where  $\text{Prim}$  denotes the random oracles  $\text{RO}_r$  and  $\text{RO}_z$ . The random oracles  $\text{RO}_r$  and  $\text{RO}_z$  internally keep lists  $\mathcal{H}$  and  $\mathcal{F}$ , which are initially set to  $\perp$ . On input  $x$  from the domain of  $\text{RO}_r$  (resp.  $\text{RO}_z$ ), the oracle first checks if  $\mathcal{H}[x] \neq \perp$  (resp.  $\mathcal{F}[x] \neq \perp$ ), i.e., checks if it has already been defined via a prior query. If it is, then it returns  $\mathcal{H}[x]$  (resp.  $\mathcal{F}[x]$ ), and otherwise, it sets  $\mathcal{H}[x]$  (resp.  $\mathcal{F}[x]$ ) to a uniformly random value in the range of  $\text{RO}_r$  (resp.  $\text{RO}_z$ ), and returns it. The same logic applies to the other internal lists,  $\mathcal{Y}$  and  $\mathcal{Z}$ , used during the security proof.

The security proof follows in a series of hybrids, where we start with the pseudorandomness game  $\text{POPRF}_{\mathbf{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^1$ , i.e.,  $b = 1$ , and we gradually apply changes until we end up with the game  $\text{POPRF}_{\mathbf{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^0$ , i.e.,  $b = 0$  and the oracle calls are answered by the simulator  $\mathcal{S}$ . Let  $\text{Adv}_{\mathcal{A}, i}(\lambda)$  denote the advantage of  $\mathcal{A}$  in **Hybrid** $_i$ . We consider the hybrids listed below, and for each hybrid we depict the changes using figures, and since the adversarially controlled algorithms are not changed, we do not include them in the figures.

<p><b>F.Setup(<math>1^\lambda</math>)</b></p> <hr/> <p>1 : <math>\forall i \in [3]: \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)</math>  2 : <math>\text{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)</math>  3 : <math>(\text{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)</math>  4 : <b>return</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></p>	<p><b>F.KeyGen(pp)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <math>\mathbf{k} \xleftarrow{\\$} \chi_{\mathbf{k}}^m</math>  3 : <math>\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)</math>  4 : <b>return</b> <math>(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))</math></p>
<p><b>F.PreProcServer(pp, sk, preq)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{sk} := (\mathbf{k}, \rho_k)</math>  3 : <b>parse</b> <math>\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}</math>  4 : <math>\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}</math>  5 : <math>\forall i \in [T]: \mathbf{e}_{s,i} \xleftarrow{\\$} \chi^{\ell+m}</math>  6 : <math>\forall i \in [T]: \mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i} \in R_q^{\ell+m}</math>  7 : <math>\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}</math>  8 : <b>return</b> <math>(\text{sts}, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})</math></p>	<p><b>F.BlindEval(pp, pk, sk, t, req, sts)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k)</math>  3 : <b>parse</b> <math>\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}</math>  4 : <b>parse</b> <math>\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)</math>  5 : <math>\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)</math>  6 : <b>if</b> <math>\text{NIZK}_1.V(\text{crs}_1, \pi_c, \text{stmt}_1) \neq 1</math> <b>then</b>  7 :     <b>return</b> <math>\perp</math>  8 : <math>\mathbf{e}'_s \xleftarrow{\\$} \chi_1^h</math>  9 : <math>\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s \in R_q^h</math>  10 : <math>\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})</math>  11 : <math>\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\text{crs}_2, \tau_2, \text{stmt}_2)</math>  12 : <b>return</b> <math>\text{rep} := (\mathbf{u}_x, \pi_s)</math></p>
<p><b>RO<sub>r</sub>(<math>\mathbf{c}_{r,i}</math>)</b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{H}[\mathbf{c}_{r,i}] = \perp</math> <b>then</b>  2 :     <math>\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\\$} R_q^{(\ell+m) \times m}</math>  3 : <b>return</b> <math>\mathcal{H}[\mathbf{c}_{r,i}]</math></p>	<p><b>F.Eval(sk, t, x)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{sk} := (\mathbf{k}, \rho_k)</math>  2 : <math>\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}</math>  3 : <math>\mathbf{z} := [\mathbf{B}_{x,t} \mathbf{k}]_p \in R_p^h</math>  4 : <math>y := \text{RO}_z(t, x, \mathbf{z})</math>  5 : <b>return</b> <math>y</math></p>
<p><b>RO<sub>z</sub>(<math>t, z, \mathbf{z}</math>)</b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{F}[t, x, \mathbf{z}] = \perp</math> <b>then</b>  2 :     <math>\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\\$} \{0, 1\}^\lambda</math>  3 : <b>return</b> <math>\mathcal{F}[t, x, \mathbf{z}]</math></p>	

Fig. 15: **Hybrid<sub>1</sub>** of pseudorandomness proof.

**Hybrid<sub>0</sub>**: This corresponds to the pseudorandomness game  $\text{POPRF}_{\text{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^1$  (Definition 3), and the interaction follows as in Figures 8 and 9. Hence, it follows that

$$\text{Adv}_{\text{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^{\text{po-prf}}(\lambda) = \text{Adv}_{\mathcal{A}, 0}(\lambda).$$

**Hybrid<sub>1</sub>**: In this hybrid, we replace  $\text{crs}_2$  and proof  $\pi_s$  with a simulated CRS and proof  $(\text{crs}_2^*, \pi_s^*)$ , as shown in Figure 15.

In order to argue indistinguishability between **Hybrid<sub>0</sub>** and **Hybrid<sub>1</sub>**, we construct a reduction  $\mathcal{B}_1$  to the computational zero-knowledge property of  $\text{NIZK}_2$ . We note that all queries are answered as in **Hybrid<sub>0</sub>** with the exception of **BlindEval** queries. The only difference here is that  $\mathcal{B}_1$  receives the CRS  $\text{crs}_2$  from its zero-knowledge challenger instead of generating it via the  $\text{NIZK}_2.\text{Setup}$  algorithm. When  $\mathcal{A}$  makes a **BlindEval** query,  $\mathcal{B}_1$  proceeds as in **Hybrid<sub>0</sub>** to compute  $\mathbf{v}_{k,j}$  and  $\mathbf{u}_x$ , but it makes an oracle call to its zero-knowledge challenger with the input  $((\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j}), (\mathbf{k}, \mathbf{e}_{s,j}, \mathbf{e}'_s, \rho_k))$  to obtain the proof  $\pi_s$ .

If the zero-knowledge challenger of  $\text{NIZK}_2$  used the honest setup and prover algorithms to generate  $\text{crs}_2$  and  $\pi_s$ , then we are exactly in **Hybrid<sub>0</sub>**, and if it used the simulator, then we are in **Hybrid<sub>1</sub>**. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_1$  can break the computational

zero-knowledge property of  $\text{NIZK}_2$ . Hence, it follows that

$$|\text{Adv}_{\mathcal{A},0}(\lambda) - \text{Adv}_{\mathcal{A},1}(\lambda)| \leq \text{Adv}_{\text{NIZK}_2, \mathcal{B}_1}^{\text{CZK}}(\lambda).$$

**Hybrid<sub>2</sub>**: Let  $\text{BAD}_1$  be the event that for a blind evaluation query  $(t, \text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j))$ , it holds that  $\text{NIZK}_1.V(\text{crs}_1, \pi_c, (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \mathbf{H}, t, j)) = 1$  (where  $\mathbf{c}_{r,j}$  is in  $\text{sts}$ ), but  $\nexists (\mathbf{R}_j, x, \rho_{r,j}, \rho_x)$  such that

$$\begin{aligned} \mathbf{C}_x &= \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{H}(x, t) \bmod q \\ \wedge \mathbf{c}_{r,j} &= \text{COM.Commit}(\text{ck}_3, \mathbf{R}_j; \rho_{r,j}) \\ \wedge \mathbf{d}_x &= \text{COM.Commit}(\text{ck}_1, x; \rho_x). \end{aligned}$$

Next, we prove that the  $\text{BAD}_1$  event can only happen with negligible probability due to the soundness of  $\text{NIZK}_1$  argument system. Looking ahead, we need here such a reduction to the soundness of  $\text{NIZK}_1$  because later in **Hybrid<sub>4</sub>** we do a reduction to the extractability of the commitment scheme, and there we need to ensure that the adversarial commitment  $\mathbf{c}_{r,j}$  is well-formed.

Clearly, we have that **Hybrid<sub>1</sub>** and **Hybrid<sub>2</sub>** are identical until the event  $\text{BAD}_1$  happens. Hence, we show that we can bound the probability  $\Pr[\text{BAD}_1]$  by constructing a reduction  $\mathcal{B}_2$  to the computational soundness of  $\text{NIZK}_1$  argument system. We note that all queries are answered as in **Hybrid<sub>1</sub>**. Let  $Q_B$  denote the maximum number of  $\text{BlindEval}$  queries that  $\mathcal{A}$  performs.  $\mathcal{B}_2$  randomly chooses an index  $i^* \in [Q_B]$  to serve as its guess for the query where the proof  $\pi_c$  will be forged by the adversary  $\mathcal{A}$ . For  $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$  that corresponds to the  $i^*$ -th blind evaluation query,  $\mathcal{B}_2$  simply outputs  $\pi_c$ . We emphasize that  $\mathcal{B}_2$  avoids recovering any parts of the witness  $(\mathbf{R}_j, x, \rho_{r,j}, \rho_x)$  (which are computationally infeasible to compute), and instead just guesses which of the proofs would have caused the event  $\text{BAD}_1$  to happen. Hence, we have that

$$|\text{Adv}_{\mathcal{A},1}(\lambda) - \text{Adv}_{\mathcal{A},2}(\lambda)| \leq Q_B \cdot \text{Adv}_{\text{NIZK}_1, \mathcal{B}_2}^{\text{CS}}(\lambda).$$

**Hybrid<sub>3</sub>**: In this hybrid, during key generation we compute the commitment  $\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{0}_m)$ , where  $\mathbf{0}_m$  is an all zero vector of length  $m$ , instead of computing a commitment of the secret key  $\mathbf{k}$  as in the previous hybrid. This change is shown in Figure 16.

In order to show indistinguishability between **Hybrid<sub>2</sub>** and **Hybrid<sub>3</sub>**, we construct a reduction  $\mathcal{B}_3$  to the computational hiding property of  $\text{COM}$ . For computing the public key  $\text{pk}$ ,  $\mathcal{B}_3$  samples a random  $\mathbf{k} \xleftarrow{\$} \chi_k^m$ , sets  $(m_0 := \mathbf{k}, m_1 := \mathbf{0}_m)$  and submits the pair  $(m_0, m_1)$  to the computational hiding challenger of  $\text{COM}$ , which responds with  $\mathbf{c}_k$ . At this point  $\mathcal{B}_3$  sets  $\text{pk} := \mathbf{c}_k$ , and proceeds to answer the oracle queries as in the previous hybrid.

If the computational hiding challenger of  $\text{COM}$  provided a commitment to  $m_0$ , then we are exactly in **Hybrid<sub>2</sub>**, and if it provided a commitment to  $m_1$ , then we are in **Hybrid<sub>3</sub>**. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_3$  can break the computational hiding property of  $\text{COM}$ . Hence, it follows that

$$|\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3}(\lambda)| \leq \text{Adv}_{\text{COM}, \mathcal{B}_3}^{\text{CH}}(\lambda).$$

**Hybrid<sub>4</sub>**: In this hybrid, we replace the commitment key  $\text{ck}_3$  with a trapdoor variant  $(\text{ck}_3, \text{td}_3) \leftarrow \text{COM.E}_1(1^\lambda)$ , and use the trapdoor  $\text{td}_3$  to extract the committed values queried to  $\text{RO}_r$  oracle, i.e., run  $\mathbf{R}_i \leftarrow \text{COM.E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})$ , and store the  $\mathbf{R}_i$  in the list  $\mathcal{R}$  for later use, as depicted in Figure 17.

In order to prove indistinguishability between **Hybrid<sub>3</sub>** and **Hybrid<sub>4</sub>** we construct an adversary  $\mathcal{B}_4$  against the extractability property of  $\text{COM}$ . We note that all queries are answered as in **Hybrid<sub>3</sub>** with the exception of  $\text{RO}_r$  queries.  $\mathcal{B}_4$  samples a trapdoor commitment key  $(\text{ck}_3, \text{td}_3) \leftarrow \text{COM.E}_1(1^\lambda)$ , instead of generating  $\text{ck}_3$  via the  $\text{COM.Setup}$  algorithm. Let  $Q_{\text{RO}_r}$  denote the number of  $\text{RO}_r$  queries that the adversary  $\mathcal{A}$  makes. For each such query  $\mathbf{c}_{r,i}$ ,  $\mathcal{B}_4$  uses the trapdoor  $\text{td}$  to extract  $\mathbf{R}_i \leftarrow \text{COM.E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})$  and stores it in the list  $\mathcal{R}[\mathbf{c}_{r,i}] := \mathbf{R}_i$ . We remark that these  $\mathbf{R}_i$  values will be used at a later hybrid.

Due to the soundness of the client proof  $\pi_c$  (as proven in **Hybrid<sub>2</sub>**) we know that the adversarial commitment  $\mathbf{c}_{r,i}$  is well-formed, and hence, the extractor extracts the correct  $\mathbf{R}_i$ . Hence, if  $\mathcal{A}$  can distinguish

<b>F.Setup(<math>1^\lambda</math>)</b> <hr/> 1 : $\forall i \in [3] : \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)$ 2 : $\text{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)$ 3 : $(\text{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)$ 4 : <b>return</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$	<b>F.KeyGen(pp)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$ 3 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{0}_m; \rho_k)$ 4 : <b>return</b> $(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))$
<b>F.PreProcServer(pp, sk, preq)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}$ 4 : $\forall i \in [T] : \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\forall i \in [T] : \mathbf{e}_{s,i} \xleftarrow{\$} \chi^{\ell+m}$ 6 : $\forall i \in [T] : \mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i} \in R_q^{\ell+m}$ 7 : $\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 8 : <b>return</b> $(\text{sts}, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})$	<b>F.BlindEval(pp, pk, sk, t, req, sts)</b> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 4 : <b>parse</b> $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 5 : $\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)$ 6 : <b>if</b> $\text{NIZK}_1.V(\text{crs}_1, \pi_c, \text{stmt}_1) \neq 1$ <b>then</b> 7 : <b>return</b> $\perp$ 8 : $\mathbf{e}'_s \xleftarrow{\$} \chi_1^h$ 9 : $\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s \in R_q^h$ 10 : $\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 11 : $\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\text{crs}_2, \tau_2, \text{stmt}_2)$ 12 : <b>return</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$
<b>RO<sub>r</sub>(<math>\mathbf{c}_{r,i}</math>)</b> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$	<b>F.Eval(sk, t, x)</b> <hr/> 1 : <b>parse</b> $\text{sk} := (\mathbf{k}, \rho_k)$ 2 : $\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}$ 3 : $\mathbf{z} := [\mathbf{B}_{x,t} \mathbf{k}]_p \in R_p^h$ 4 : $y := \text{RO}_z(t, x, \mathbf{z})$ 5 : <b>return</b> $y$
<b>RO<sub>z</sub>(<math>t, z, \mathbf{z}</math>)</b> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$	

Fig. 16: **Hybrid<sub>3</sub>** of pseudorandomness proof.

between the two hybrids with non-negligible advantage, then  $\mathcal{B}_4$  can break the extractability property of COM. Moreover, since the adversary  $\mathcal{A}$  makes at most  $Q_{\text{RO}_r}$  queries to  $\text{RO}_r$  oracle, it follows that

$$|\text{Adv}_{\mathcal{A},3}(\lambda) - \text{Adv}_{\mathcal{A},4}(\lambda)| \leq Q_{\text{RO}_r} \cdot \text{Adv}_{\text{COM},\mathcal{B}_4}^{\text{Ext}}(\lambda).$$

**Hybrid<sub>5</sub>**: In this hybrid, we replace the commitment key  $\text{ck}_1$  with a trapdoor variant  $(\text{ck}_1, \text{td}_1) \leftarrow \text{COM}.\mathcal{E}_1(1^\lambda)$ , and use the trapdoor  $\text{td}_1$  to extract  $x \leftarrow \text{COM}.\mathcal{E}_2(\text{ck}_1, \text{td}_1, \mathbf{d}_x)$  when receiving **BlindEval** queries, as depicted in Figure 18. We note that the extracted input  $x$  will be used in the next hybrid.

The indistinguishability between **Hybrid<sub>4</sub>** and **Hybrid<sub>5</sub>** follows from the extractability property of COM as in the previous hybrid. Concretely, due to the soundness of the client proof  $\pi_c$  (as proven in **Hybrid<sub>2</sub>**) we know that the adversarial commitment  $\mathbf{d}_x$  is well-formed, and hence, the extractor extracts the correct input  $x$ . Moreover, since the adversary  $\mathcal{A}$  makes at most  $Q_{\mathcal{B}}$  queries to the **BlindEval** oracle, it follows that

$$|\text{Adv}_{\mathcal{A},4}(\lambda) - \text{Adv}_{\mathcal{A},5}(\lambda)| \leq Q_{\mathcal{B}} \cdot \text{Adv}_{\text{COM},\mathcal{B}_5}^{\text{Ext}}(\lambda),$$

where  $\mathcal{B}_5$  is the reduction to the computational hiding of COM.

<p><b>F.Setup(<math>1^\lambda</math>)</b></p> <hr/> <p>1 : <math>\forall i \in [2] : \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)</math>  2 : <math>(\text{ck}_3, \text{td}_3) \leftarrow \text{COM.E}_1(1^\lambda)</math>  3 : <math>\text{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)</math>  4 : <math>(\text{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)</math>  5 : <b>return</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></p> <p><b>F.PreProcServer(pp, sk, preq)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{sk} := (\mathbf{k}, \rho_k)</math>  3 : <b>parse</b> <math>\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}</math>  4 : <math>\forall i \in [T] : \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}</math>  5 : <math>\forall i \in [T] : \mathbf{e}_{s,i} \xleftarrow{\\$} \chi^{\ell+m}</math>  6 : <math>\forall i \in [T] : \mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i} \in R_q^{\ell+m}</math>  7 : <math>\text{st}_5 := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}</math>  8 : <b>return</b> <math>(\text{st}_5, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})</math></p> <p><b><math>\text{RO}_r(\mathbf{c}_{r,i})</math></b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{H}[\mathbf{c}_{r,i}] = \perp</math> <b>then</b>  2 : <math>\mathbf{A}_{r,i} := \mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\\$} R_q^{(\ell+m) \times m}</math>  3 : <math>\mathbf{R}_i \leftarrow \text{COM.E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})</math>  4 : <math>\mathcal{R}[\mathbf{c}_{r,i}] := \mathbf{R}_i</math>  5 : <b>return</b> <math>\mathcal{H}[\mathbf{c}_{r,i}]</math></p> <p><b><math>\text{RO}_z(t, z, \mathbf{z})</math></b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{F}[t, x, \mathbf{z}] = \perp</math> <b>then</b>  2 : <math>\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\\$} \{0, 1\}^\lambda</math>  3 : <b>return</b> <math>\mathcal{F}[t, x, \mathbf{z}]</math></p>	<p><b>F.KeyGen(pp)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <math>\mathbf{k} \xleftarrow{\\$} \chi_k^m</math>  3 : <math>\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{0}_m; \rho_k)</math>  4 : <b>return</b> <math>(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))</math></p> <p><b>F.BlindEval(pp, pk, sk, t, req, st<sub>5</sub>)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k)</math>  3 : <b>parse</b> <math>\text{st}_5 := (\mathbf{c}_{r,j}, \{\mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]})</math>  4 : <b>parse</b> <math>\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)</math>  5 : <math>\text{stmt}_1 := (\mathbf{c}_r, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)</math>  6 : <b>if</b> <math>\text{NIZK}_1.\text{V}(\text{crs}_1, \pi_c, \text{stmt}_1) \neq 1</math> <b>then</b>  7 : <b>return</b> <math>\perp</math>  8 : <math>\mathbf{e}'_s \xleftarrow{\\$} \chi_1^h</math>  9 : <math>\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s \in R_q^h</math>  10 : <math>\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})</math>  11 : <math>\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\text{crs}_2, \tau_2, \text{stmt}_2)</math>  12 : <b>return</b> <math>\text{rep} := (\mathbf{u}_x, \pi_s)</math></p> <p><b>F.Eval(sk, t, x)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{sk} := (\mathbf{k}, \rho_k)</math>  2 : <math>\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}</math>  3 : <math>\mathbf{z} := \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p \in R_p^h</math>  4 : <math>y := \text{RO}_z(t, x, \mathbf{z})</math>  5 : <b>return</b> <math>y</math></p>
---	---

Fig. 17: **Hybrid<sub>4</sub>** of pseudorandomness proof.

**Hybrid<sub>6</sub>**: In this hybrid, we remove the secret key  $\mathbf{k}$  and change the way we respond to the PreProcServer, BlindEval and Eval oracle queries. Concretely, for answering Eval queries, instead of computing  $\mathbf{z} := \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p$ , we sample a uniform  $\mathcal{Z}[\mathbf{B}_{x,t}] \xleftarrow{\$} R_q^h$  and return  $\mathbf{z} := \lfloor \mathcal{Z}[\mathbf{B}_{x,t}] \rfloor_p$ . For answering PreProcServer queries, instead of computing  $\mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i}$  we sample a uniformly random  $\mathbf{v}_{k,i} \xleftarrow{\$} R_q^{\ell+m}$ . Additionally, when answering the BlindEval queries, instead of computing  $\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s$  we compute  $\mathbf{u}_x := \mathbf{R}_j \mathbf{v}_{k,j} + \mathcal{Z}[\mathbf{B}_{x,t}] + \tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}}$ , for some  $\tilde{\mathbf{e}}' \xleftarrow{\$} \chi_1^h$  and  $\tilde{\mathbf{e}} \xleftarrow{\$} \chi^{\ell+m}$ . These changes are depicted in Figure 19.

In order to show indistinguishability between **Hybrid<sub>5</sub>** and **Hybrid<sub>6</sub>** we construct an adversary  $\mathcal{B}_6$  against the MLWE-PRF-RU<sub>param</sub> assumption, for  $\text{param} := (\text{H}, Q, Q_{x,t}^\infty, Q_M, q, m, \ell + m, h, L, \beta_r, \vec{\chi})$ , for  $\vec{\chi} := (\chi, \chi_1, \chi_k)$ . For answering  $\text{RO}_r$  query  $(\mathbf{c}_{r,i})$ ,  $\mathcal{B}_6$  extracts  $\mathbf{R}_i \leftarrow \text{COM.E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})$  (as in the previous hybrid), then queries the MLWE-PRF-RU oracle  $(\mathbf{A}, \mathbf{c}) \leftarrow \text{SampMLWecom}(\mathbf{R}_i)$ , and sets  $\mathcal{H}[\mathbf{c}_{r,i}] := \mathbf{A}$  and  $\mathbf{v}_{k,i} := \mathbf{c}$  (this is without loss of generality since when answering the PreProcServer oracle queries  $\mathcal{B}_6$  will make a call to the  $\text{RO}_r$  oracle). Upon receiving a BlindEval query  $(i, t, (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j))$ ,  $\mathcal{B}_6$  extracts  $x \leftarrow$

<p><b>F.Setup</b>(<math>1^\lambda</math>)</p> <hr/> 1 : $(\mathbf{ck}_1, \mathbf{td}_1) \leftarrow \text{COM.E}_1(1^\lambda)$ 2 : $\mathbf{ck}_2 \leftarrow \text{COM.Setup}(1^\lambda)$ 3 : $(\mathbf{ck}_3, \mathbf{td}_3) \leftarrow \text{COM.E}_1(1^\lambda)$ 4 : $\mathbf{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)$ 5 : $(\mathbf{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)$ 6 : <b>return</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$	<p><b>F.KeyGen</b>(pp)</p> <hr/> 1 : <b>parse</b> pp := $(\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{k} \xleftarrow{\$} \chi_k^m$ 3 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\mathbf{ck}_2, \mathbf{0}_m; \rho_k)$ 4 : <b>return</b> $(\mathbf{pk} := \mathbf{c}_k, \mathbf{sk} := (\mathbf{k}, \rho_k))$
<p><b>F.PreProcServer</b>(pp, sk, preq)</p> <hr/> 1 : <b>parse</b> pp := $(\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> sk := $(\mathbf{k}, \rho_k)$ 3 : <b>parse</b> preq := $\{\mathbf{c}_{r,i}\}_{i \in [T]}$ 4 : $\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\forall i \in [T]: \mathbf{e}_{s,i} \xleftarrow{\$} \chi^{\ell+m}$ 6 : $\forall i \in [T]: \mathbf{v}_{k,i} := \mathbf{A}_{r,i} \mathbf{k} + \mathbf{e}_{s,i} \in R_q^{\ell+m}$ 7 : $\mathbf{st}_5 := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 8 : <b>return</b> $(\mathbf{st}_5, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})$	<p><b>F.BlindEval</b>(pp, pk, sk, t, req, st<sub>5</sub>)</p> <hr/> 1 : <b>parse</b> pp := $(\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> pk := $\mathbf{c}_k, \mathbf{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> st <sub>5</sub> := $(\mathbf{c}_{r,j}, \{\mathbf{A}_{r,i}, \mathbf{e}_{s,i}, \mathbf{v}_{k,i}\}_{i \in [T]})$ 4 : <b>parse</b> req := $(\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 5 : $\text{stmt}_1 := (\mathbf{c}_r, \mathbf{C}_x, \mathbf{d}_x, \mathbf{ck}_1, \mathbf{ck}_3, \mathbf{A}_{r,j}, \mathbf{H}, t, j)$ 6 : <b>if</b> $\text{NIZK}_1.\mathcal{V}(\mathbf{crs}_1, \pi_c, \text{stmt}_1) \neq 1$ <b>then</b> 7 : <b>return</b> $\perp$ 8 : $\mathbf{x} \leftarrow \text{COM.E}_2(\mathbf{ck}_1, \mathbf{td}_1, \mathbf{d}_x)$ 9 : $\mathbf{e}'_s \xleftarrow{\$} \chi_1^h$ 10 : $\mathbf{u}_x := \mathbf{C}_x \mathbf{k} + \mathbf{e}'_s \in R_q^h$ 11 : $\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \mathbf{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 12 : $\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\mathbf{crs}_2, \tau_2, \text{stmt}_2)$ 13 : <b>return</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$
<p><math>\text{RO}_r(\mathbf{c}_{r,i})</math></p> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathbf{A}_{r,i} := \mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : $\mathbf{R}_i \leftarrow \text{COM.E}_2(\mathbf{ck}_3, \mathbf{td}_3, \mathbf{c}_{r,i})$ 4 : $\mathcal{R}[\mathbf{c}_{r,i}] := \mathbf{R}_i$ 5 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$	<p><b>F.Eval</b>(sk, t, x)</p> <hr/> 1 : <b>parse</b> sk := $(\mathbf{k}, \rho_k)$ 2 : $\mathbf{B}_{x,t} := \mathbf{H}(x, t) \in R_q^{h \times m}$ 3 : $\mathbf{z} := [\mathbf{B}_{x,t} \mathbf{k}]_p \in R_p^h$ 4 : $y := \text{RO}_z(t, x, \mathbf{z})$ 5 : <b>return</b> y
<p><math>\text{RO}_z(t, x, \mathbf{z})</math></p> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$	

Fig. 18: **Hybrid<sub>5</sub>** of pseudorandomness proof.

$\text{COM.E}_2(\mathbf{ck}_1, \mathbf{td}_1, \mathbf{d}_x)$  (as in the previous hybrid) and makes an oracle call  $\mathbf{z} \leftarrow \text{SampPRF-RU}(\mathcal{H}[\mathbf{c}_r, j], x || t)$  to the MLWE-PRF-RU challenger. Then,  $\mathcal{B}_6$  sets  $\mathbf{u}_x := \mathbf{R}_j \mathbf{v}_{k,j} + \mathbf{z}$ . The rest of the queries are answered as in the previous hybrid.

If the MLWE-PRF-RU challenger used  $b = 1$ , then the view of  $\mathcal{A}$  simulated by  $\mathcal{B}_6$  is exactly as in **Hybrid<sub>5</sub>**, and if it used  $b = 0$ , then the view of  $\mathcal{A}$  is as in **Hybrid<sub>6</sub>**. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_6$  can break the  $\text{MLWE-PRF-RU}_{\text{param}}$  assumption, where  $Q$  in the MLWE-PRF-RU assumption denotes the total number of BlindEval queries,  $Q_{x,t}^\infty$  denotes the maximum number of identical  $(x || t)$  queries and  $Q_M$  denotes the total number of  $\text{RO}_r$  queries, respectively, that the adversary  $\mathcal{A}$  makes. Combined with the above claim we get that

$$|\text{Adv}_{\mathcal{A},5}(\lambda) - \text{Adv}_{\mathcal{A},6}(\lambda)| \leq \text{Adv}_{\mathcal{B}_6}^{\text{iMLWE}_{\text{param}}}(\lambda).$$

**Hybrid<sub>7</sub>**: In this hybrid we replace the Eval algorithm with a random oracle, which is depicted in Figure 20. This corresponds to the pseudorandomness game  $\text{POPRF}_{\mathcal{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^0$  (Definition 3).

Let  $\text{BAD}_2$  be the event that the adversary  $\mathcal{A}$  queries  $(t, x, \mathbf{z})$  to  $\text{RO}_z$  oracle, for some  $\mathbf{z} := \lfloor \mathcal{Z}[t, x] \rfloor_p$ , before ever making a query to the  $\text{RO}_r$  oracle. Clearly, **Hybrid<sub>6</sub>** and **Hybrid<sub>7</sub>** are identical until the event  $\text{BAD}_2$  happens. Hence, we observe that the only distinguishing advantage the adversary  $\mathcal{A}$  has between **Hybrid<sub>6</sub>** and **Hybrid<sub>7</sub>** is if  $\mathcal{A}$  queries  $\text{RO}_z$  oracle with some tuple  $(t, x, \mathbf{z})$ , where initially it holds that  $\mathbf{z} \neq \lfloor \mathcal{Z}[\mathbf{B}_{x,t}] \rfloor_p$  (this covers the case when  $\mathcal{Z}[\mathbf{B}_{x,t}] = \perp$ ), but after making a call to the **BlindEval** oracle, we have that it is set to correct  $\mathcal{Z}[\mathbf{B}_{x,t}]$  such that  $\mathbf{z} := \lfloor \mathcal{Z}[\mathbf{B}_{x,t}] \rfloor_p$ , then  $\mathcal{A}$  will obtain two different outputs for the same  $(t, x, \mathbf{z})$  input to  $\text{RO}_z$  oracle (i.e., before and after making a call to the **BlindEval** oracle), and therefore, it can distinguish the hybrids. The probability of this happening corresponds to the aforementioned  $\text{BAD}_2$  event happening, which we now argue that it can only happen with a negligible probability.

Note that inside the **PreProcServer** oracle, we sample  $\mathbf{v}_{k,j}$  uniformly from  $R_q^{\ell+m}$ , and inside the **BlindEval** oracle we sample  $\mathcal{Z}[\mathbf{B}_{x,t}]$  uniformly from  $R_q^h$ . Then, we compute  $\mathbf{u}_x := \mathbf{R}_j \mathbf{v}_{k,j} + \mathcal{Z}[\mathbf{B}_{x,t}] + \tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}} \in R_q^h$ . Therefore, the only value that is under the control of the adversary is  $\mathbf{R}_j$ , and since the rest of the values are chosen uniformly randomly, we have that  $\mathbf{u}_x$  is distributed uniformly over  $R_q^h$ .

Concretely, we have that the adversarially controlled term of  $\mathbf{R}_j \mathbf{v}_{k,j}$  will actually cancel out during the computation of  $\lfloor \mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j} \rfloor_p$  (in the **Finalize** algorithm). Since  $\mathcal{Z}[\mathbf{B}_{x,t}]$  is sampled independently of  $\tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}}$  and after  $\mathbf{R}_j$  is fixed by the adversary, the distribution of each coefficient of  $\mathcal{Z}[\mathbf{B}_{x,t}] + \tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}}$  will be uniform mod  $q$ . Then, the chance that the adversary can hit the correct value for a coefficient after rounding is at most  $1/p + 1/q$  (the  $1/q$  terms appears when  $q$  is not a multiple of  $p$ ). Since  $\mathcal{A}$  is a PPT algorithm, it can only make polynomially many queries to  $\text{RO}_z$  oracle. Let  $Q_{\text{RO}_z}$  denote the total number of queries that  $\mathcal{A}$  makes to  $\text{RO}_z$  oracle, then

$$\Pr[\text{BAD}_2] \leq Q_{\text{RO}_z} \cdot \left( \frac{1}{p} + \frac{1}{q} \right)^{dh} =: \epsilon_{\text{BAD}}.$$

Finally, we note that the simulator  $\mathcal{S}$  answers the oracle calls as in **Hybrid<sub>7</sub>**, and putting everything together, we obtain

$$\begin{aligned} \text{Adv}_{\mathcal{F}, \mathcal{A}, \mathcal{S}, \text{RO}_r, \text{RO}_z}^{\text{po-prf}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_2, \mathcal{B}_1}^{\text{CZK}}(\lambda) + Q_{\mathcal{B}} \cdot \text{Adv}_{\text{NIZK}_1, \mathcal{B}_2}^{\text{CS}}(\lambda) + \text{Adv}_{\text{COM}, \mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + Q_{\text{RO}_r} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_4}^{\text{Ext}}(\lambda) + Q_{\mathcal{B}} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_5}^{\text{Ext}}(\lambda) + \text{Adv}_{\mathcal{B}_6}^{\text{iMLWE}_{\text{param}}}(\lambda) \\ &\quad + \epsilon_{\text{BAD}}, \end{aligned}$$

as claimed. This completes the proof of Theorem 5. □

<p><b>F.Setup</b>(<math>1^\lambda</math>)</p> <hr/> 1 : $(\text{ck}_1, \text{td}_1) \leftarrow \text{COM}.\mathcal{E}_1(1^\lambda)$ 2 : $\text{ck}_2 \leftarrow \text{COM}.\text{Setup}(1^\lambda)$ 3 : $(\text{ck}_3, \text{td}_3) \leftarrow \text{COM}.\mathcal{E}_1(1^\lambda)$ 4 : $\text{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)$ 5 : $(\text{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)$ 6 : <b>return</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$	<p><b>F.KeyGen</b>(<math>\text{pp}</math>)</p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{c}_k \leftarrow \text{COM}.\text{Commit}(\text{ck}_2, \mathbf{0}_m; \rho_k)$ 3 : <b>return</b> $(\text{pk} := \mathbf{c}_k, \text{sk} := \perp)$
<p><b>F.PreProcServer</b>(<math>\text{pp}, \text{sk}, \text{preq}</math>)</p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}$ 4 : $\forall i \in [T] : \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\forall i \in [T] : \mathbf{v}_{k,i} \xleftarrow{\$} R_q^{\ell+m}$ 6 : $\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 7 : <b>return</b> $(\text{sts}, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})$	<p><b>F.BlindEval</b>(<math>\text{pp}, \text{pk}, \text{sk}, t, \text{req}, \text{sts}</math>)</p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{sts} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 4 : <b>parse</b> $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 5 : $\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)$ 6 : <b>if</b> $\text{NIZK}_1.\text{V}(\text{crs}_1, \pi_c, \text{stmt}_1) \neq 1$ <b>then</b> 7 : <b>return</b> $\perp$ 8 : $x \leftarrow \text{COM}.\mathcal{E}_2(\text{ck}_1, \text{td}_1, \mathbf{d}_x)$ 9 : $\mathbf{R}_j := \mathcal{R}[\mathbf{c}_{r,j}]$ 10 : $\mathbf{B}_{x,t} := \mathbf{C}_x - \mathbf{R}_j \mathbf{A}_{r,j} \in R_q^{h \times m}$ 11 : <b>if</b> $\mathcal{Z}[\mathbf{B}_{x,t}] = \perp$ <b>then</b> 12 : $\mathbf{z}' \xleftarrow{\$} R_q^h$ 13 : $\mathcal{Z}[\mathbf{B}_{x,t}] := \mathbf{z}'$ 14 : $\tilde{\mathbf{e}}' \xleftarrow{\$} \chi_1^h, \tilde{\mathbf{e}} \xleftarrow{\$} \chi^{\ell+m}$ 15 : $\mathbf{e}^* := \tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}} \in R^h$ 16 : $\mathbf{u}_x := \mathbf{R}_j \mathbf{v}_{k,j} + \mathcal{Z}[\mathbf{B}_{x,t}] + \mathbf{e}^* \in R_q^h$ 17 : $\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 18 : $\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\text{crs}_2, \tau_2, \text{stmt}_2)$ 19 : <b>return</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$
<p><b>RO<sub>r</sub></b>(<math>\mathbf{c}_{r,i}</math>)</p> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathbf{A}_{r,i} := \mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : $\mathbf{R}_i \leftarrow \text{COM}.\mathcal{E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})$ 4 : $\mathcal{R}[\mathbf{c}_{r,i}] := \mathbf{R}_i$ 5 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$	<p><b>F.Eval</b>(<math>\text{sk}, t, x</math>)</p> <hr/> 1 : $\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}$ 2 : <b>if</b> $\mathcal{Z}[\mathbf{B}_{x,t}] = \perp$ <b>then</b> 3 : $\mathbf{z}' \xleftarrow{\$} R_q^h$ 4 : $\mathcal{Z}[\mathbf{B}_{x,t}] := \mathbf{z}'$ 5 : $\mathbf{z} := \lfloor \mathcal{Z}[\mathbf{B}_{x,t}] \rfloor_p$ 6 : $y := \text{RO}_z(t, x, \mathbf{z})$ 7 : <b>return</b> $y$
<p><b>RO<sub>z</sub></b>(<math>t, z, \mathbf{z}</math>)</p> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$	

Fig. 19: **Hybrid<sub>6</sub>** of pseudorandomness proof.

F.Setup( $1^\lambda$ )	F.KeyGen(pp)
1 : $(\text{ck}_1, \text{td}_1) \leftarrow \text{COM.E}_1(1^\lambda)$ 2 : $\text{ck}_2 \leftarrow \text{COM.Setup}(1^\lambda)$ 3 : $(\text{ck}_3, \text{td}_3) \leftarrow \text{COM.E}_1(1^\lambda)$ 4 : $\text{crs}_1 \leftarrow \text{NIZK}_1.\text{Setup}(1^\lambda)$ 5 : $(\text{crs}_2, \tau_2) \leftarrow \text{NIZK}_2.\mathcal{S}_1(1^\lambda)$ 6 : <b>return</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$	1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{0}_m; \rho_k)$ 3 : <b>return</b> $(\text{pk} := \mathbf{c}_k, \text{sk} := \perp)$
F.PreProcServer(pp, sk, preq)	F.BlindEval(pp, pk, sk, t, req, st <sub>S</sub> )
1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]}$ 4 : $\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\forall i \in [T]: \mathbf{v}_{k,i} \xleftarrow{\$} R_q^{\ell+m}$ 6 : $\text{st}_S := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 7 : <b>return</b> $(\text{st}_S, \text{prep} := \{\mathbf{v}_{k,i}\}_{i \in [T]})$	1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k)$ 3 : <b>parse</b> $\text{st}_S := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{v}_{k,i}\}_{i \in [T]}$ 4 : <b>parse</b> $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 5 : $\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)$ 6 : <b>if</b> $\text{NIZK}_1.V(\text{crs}_1, \pi_c, \text{stmt}_1) \neq 1$ <b>then</b> 7 : <b>return</b> $\perp$ 8 : $x \leftarrow \text{COM.E}_2(\text{ck}_1, \text{td}_1, \mathbf{d}_x)$ 9 : $\mathbf{R}_j := \mathcal{R}[\mathbf{c}_{r,j}]$ 10 : $\mathbf{B}_{x,t} := \mathbf{C}_x - \mathbf{R}_j \mathbf{A}_{r,j} \in R_q^{h \times m}$ 11 : <b>if</b> $\mathcal{Z}[\mathbf{B}_{x,t}] = \perp$ <b>then</b> 12 : $\mathbf{z}' \xleftarrow{\$} R_q^h$ 13 : $\mathcal{Z}[\mathbf{B}_{x,t}] := \mathbf{z}'$ 14 : $\tilde{\mathbf{e}}' \xleftarrow{\$} \chi_1^h, \tilde{\mathbf{e}} \xleftarrow{\$} \chi^{\ell+m}$ 15 : $\mathbf{e}^* := \tilde{\mathbf{e}}' - \mathbf{R}_j \tilde{\mathbf{e}} \in R^h$ 16 : $\mathbf{u}_x := \mathbf{R}_j \mathbf{v}_{k,j} + \mathcal{Z}[\mathbf{B}_{x,t}] + \mathbf{e}^* \in R_q^h$ 17 : $\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 18 : $\pi_s \leftarrow \text{NIZK}_2.\mathcal{S}_2(\text{crs}_2, \tau_2, \text{stmt}_2)$ 19 : <b>return</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$
$\text{RO}_r(\mathbf{c}_{r,i})$	
1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathbf{A}_{r,i} := \mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : $\mathbf{R}_i \leftarrow \text{COM.E}_2(\text{ck}_3, \text{td}_3, \mathbf{c}_{r,i})$ 4 : $\mathcal{R}[\mathbf{c}_{r,i}] := \mathbf{R}_i$ 5 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$	
$\text{RO}_z(t, x, \mathbf{z})$	
1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}$ 3 : <b>if</b> $\mathcal{Z}[\mathbf{B}_{x,t}] \neq \perp$ <b>and</b> $\mathbf{z} = \lfloor \mathcal{Z}[\mathbf{B}_{x,t}] \rfloor_p$ <b>then</b> 4 : <b>if</b> $\mathcal{Y}[t, x] = \perp$ <b>then</b> 5 : $\mathcal{Y}[t, x] \xleftarrow{\$} \{0, 1\}^\lambda$ 6 : $\mathcal{F}[t, x, \mathbf{z}] := \mathcal{Y}[t, x]$ 7 : <b>else</b> 8 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 9 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$	
	F.Eval(sk, t, x)
	1 : <b>if</b> $\mathcal{Y}[t, x] = \perp$ <b>then</b> 2 : $\mathcal{Y}[t, x] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{Y}[t, x]$

Fig. 20: **Hybrid<sub>7</sub>** of pseudorandomness proof.

## D.4 Request Privacy Analysis

**Theorem 6.** *The POPRF construction F from Figures 8 and 9 satisfies the request privacy against malicious servers (POPRIV2), given in Definition 4, with random oracles  $\text{RO}_r$  and  $\text{RO}_z$ , if:*

- The client argument system  $\text{NIZK}_1$  is computationally zero-knowledge and the server argument system  $\text{NIZK}_2$  is computationally sound (Definition 18),
- The commitment scheme COM is computationally hiding and extractable (Definitions 15 and 17), and
- The  $\text{knMLWE}_{\ell+m,m,h,\chi}$  assumption holds (Definition 7).

More precisely, for any PPT adversary  $\mathcal{A}$ , there exist PPT adversaries  $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4, \mathcal{B}_5$  and  $\mathcal{B}_6$  against the computational zero-knowledge of  $\text{NIZK}_1$ , computational hiding of COM, computational soundness of  $\text{NIZK}_2$ , extractability of COM and  $\text{knMLWE}_{m,\ell+m,h,\chi}$  assumption, respectively, such that

$$\begin{aligned} \text{Adv}_{\text{F},\mathcal{A},\text{RO}_r,\text{RO}_z}^{\text{po-priv-2}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_1,\mathcal{B}_1}^{\text{CZK}}(\lambda) + T \cdot Q_P \cdot \text{Adv}_{\text{COM},\mathcal{B}_2}^{\text{CH}}(\lambda) + 2Q_R \cdot \text{Adv}_{\text{COM},\mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + 2Q_F \cdot \text{Adv}_{\text{NIZK}_2,\mathcal{B}_4}^{\text{CS}}(\lambda) + Q_R^{\text{pk}} \cdot \text{Adv}_{\text{COM},\mathcal{B}_5}^{\text{Ext}}(\lambda) + 2Q_R \cdot \text{Adv}_{\mathcal{B}_6}^{\text{knMLWE}}(\lambda), \end{aligned}$$

where  $Q_P, Q_R$  and  $Q_F$  denote the number of PreProcClient, Request and Finalize oracle queries, respectively, and  $Q_R^{\text{pk}}$  denotes the number of queries to the Request oracle with distinct pk inputs that the adversary  $\mathcal{A}$  makes.

*Proof.* We consider the POPRF request privacy against malicious servers game given in Definition 4. We note that the adversary  $\mathcal{A}$  has access to the oracles PreProcClient, Request, Finalize and the random oracles  $\text{RO}_r$  and  $\text{RO}_z$ . The security proof follows in a series of hybrids, where we start with the request privacy against malicious servers game  $\text{POPRIV2}_{\text{F},\mathcal{A},\text{RO}_r,\text{RO}_z}^b$ , for a random challenge bit  $b \in \{0, 1\}$ , and we gradually apply changes until we end up with a game where the transcript observed by  $\mathcal{A}$  is independent of the challenge bit  $b$ . Let  $\text{Adv}_{\mathcal{A},i}(\lambda)$  denote the advantage of  $\mathcal{A}$  in **Hybrid** <sub>$i$</sub> . We consider the hybrids listed below, and for each hybrid we depict the changes using figures, and since the adversary adversarially controlled algorithms are not changed, we do not include them in the figures.

**Hybrid**<sub>0</sub>: This corresponds to the request privacy against malicious servers game  $\text{POPRIV2}_{\text{F},\mathcal{A},\text{RO}_r,\text{RO}_z}^b$  (Definition 4), for a random challenge bit  $b \in \{0, 1\}$ , and the interaction follows as in Figures 8 and 9. Hence, it follows that

$$\text{Adv}_{\text{F},\mathcal{A},\text{RO}_r,\text{RO}_z}^{\text{po-priv-2}}(\lambda) = \text{Adv}_{\mathcal{A},0}(\lambda).$$

**Hybrid**<sub>1</sub>: In this hybrid, we replace  $\text{crs}_1$  and proof  $\pi_c$  with a simulated CRS and proof as shown in Figure 21.

In order to indistinguishability between **Hybrid**<sub>0</sub> and **Hybrid**<sub>1</sub>, we construct a reduction  $\mathcal{B}_1$  to the computational zero-knowledge property of  $\text{NIZK}_1$ . We note that all queries are answered as in **Hybrid**<sub>0</sub> with the exception of the Request queries. The only difference here is that  $\mathcal{B}_1$  receives the CRS  $\text{crs}_1$  from its zero-knowledge challenger instead of computing it via the  $\text{NIZK}_1.\text{Setup}$  algorithm. Moreover, when  $\mathcal{A}$  makes a Request query  $(i, \text{pk}, \text{prep}, t, x_0, x_1)$ ,  $\mathcal{B}_1$  uses  $\text{st}_{c,i}$  (that is the state produced from the  $i$ -th PreProcClient oracle call), and proceeds as in **Hybrid**<sub>0</sub> to compute  $(\mathbf{d}_{x_0}, \mathbf{C}_{x_0})$  and  $(\mathbf{d}_{x_1}, \mathbf{C}_{x_1})$ , but it makes two oracle calls to its zero-knowledge challenger with the inputs  $((\mathbf{c}_{r,j}, \mathbf{C}_{x_0}, \mathbf{d}_{x_0}, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \mathbf{H}, t, j), (\mathbf{R}_j, x_0, \rho_{r,j}, \rho_{x_0}))$  and  $((\mathbf{c}_{r,j+1}, \mathbf{C}_{x_1}, \mathbf{d}_{x_1}, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j+1}, \mathbf{H}, t, j+1), (\mathbf{R}_{j+1}, x_1, \rho_{r,j+1}, \rho_{x_1}))$ , to obtain the proofs  $\pi_{c,0}$  and  $\pi_{c,1}$  (without loss of generality we assume  $j < T$  here, otherwise  $\mathcal{B}_1$  returns  $\perp$  and aborts). Finally,  $\mathcal{B}_1$  responds to the Request query with  $(\text{req}_b := (\mathbf{d}_{x,b}, \mathbf{C}_{x_b}, \pi_{c,b}, j_b), \text{req}_{1-b} := (\mathbf{d}_{x,1-b}, \mathbf{C}_{x_{1-b}}, \pi_{c,1-b}, j_{1-b}))$ .

If the zero-knowledge challenger of  $\text{NIZK}_1$  used the honest setup and prover algorithms to generate  $\text{crs}_1$  and  $(\pi_{c,0}, \pi_{c,1})$ , then we are exactly in **Hybrid**<sub>0</sub>, and if it used the simulator, then we are in **Hybrid**<sub>1</sub>. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_1$  can break the zero-knowledge property of  $\text{NIZK}_1$ . Hence, it follows that

$$|\text{Adv}_{\mathcal{A},0}(\lambda) - \text{Adv}_{\mathcal{A},1}(\lambda)| \leq \text{Adv}_{\text{NIZK}_1,\mathcal{B}_1}^{\text{CZK}}(\lambda).$$

<p><b>F.Setup(<math>1^\lambda</math>)</b></p> <hr/> <p>1 : <math>\forall i \in [3]: \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)</math>  2 : <math>(\text{crs}_1, \tau) \leftarrow \text{NIZK}_1.\mathcal{S}_1(1^\lambda)</math>  3 : <math>\text{crs}_2 \leftarrow \text{NIZK}_2.\text{Setup}(1^\lambda)</math>  4 : <b>return</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></p>	<p><b>F.KeyGen(pp)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <math>\mathbf{k} \xleftarrow{\\$} \chi_k^m</math>  3 : <math>\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)</math>  4 : <b>return</b> <math>(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))</math></p>
<p><b>F.PreProcClient(pp, pk, T)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <math>\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\\$} \chi_r^{h \times (\ell+m)}</math>  3 : <math>\forall i \in [T]: \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{R}_i; \rho_{r,i})</math>  4 : <math>\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}</math>  5 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}</math>  6 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j := 1, \perp)</math>  7 : <b>return</b> <math>(\text{st}_C, \text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]})</math></p>	<p><b>F.Request(pp, pk, t, x, st<sub>C</sub>)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j, \cdot)</math>  3 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math>  4 : <b>if</b> <math>j &gt; T</math> <b>then return</b> <math>\perp</math>  5 : <math>\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}</math>  6 : <math>\mathbf{C}_x := \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{B}_{x,t} \in R_q^{h \times m}</math>  7 : <math>\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, \mathbf{C}_x; \rho_x)</math>  8 : <math>\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)</math>  9 : <math>\pi_c \leftarrow \text{NIZK}_1.\mathcal{S}_2(\text{crs}_1, \tau, \text{stmt}_1)</math>  10 : <math>\text{st}_j := (t, x, \mathbf{C}_x, \text{pk})</math>  11 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j+1, \text{st}_j)</math>  12 : <math>\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)</math>  13 : <b>return</b> <math>(\text{st}_C, \text{req})</math></p>
<p><b>F.Finalize(pp, rep, st<sub>C</sub>)</b></p> <hr/> <p>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math>  2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j+1, \text{st}_j)</math>  3 : <b>parse</b> <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math>  4 : <b>parse</b> <math>\text{st}_j := (t, x, \mathbf{C}_x, \mathbf{c}_k)</math>  5 : <b>parse</b> <math>\text{rep} := (\mathbf{u}_x, \pi_s)</math>  6 : <math>\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})</math>  7 : <b>if</b> <math>\text{NIZK}_2.\mathcal{V}(\text{crs}_2, \pi_s, \text{stmt}_2) \neq 1</math> <b>then</b>  8 :     <b>return</b> <math>\perp</math>  9 : <math>\mathbf{z} := \lfloor \mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j} \rfloor_p \in R_p^h \quad // = \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p</math>  10 : <math>y := \text{RO}_z(t, x, \mathbf{z})</math>  11 : <b>return</b> <math>y</math></p>	<p><b>RO<sub>r</sub>(c<sub>r,i</sub>)</b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{H}[\mathbf{c}_{r,i}] = \perp</math> <b>then</b>  2 :     <math>\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\\$} R_q^{(\ell+m) \times m}</math>  3 : <b>return</b> <math>\mathcal{H}[\mathbf{c}_{r,i}]</math></p> <p><b>RO<sub>z</sub>(t, z, z)</b></p> <hr/> <p>1 : <b>if</b> <math>\mathcal{F}[t, x, \mathbf{z}] = \perp</math> <b>then</b>  2 :     <math>\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\\$} \{0, 1\}^\lambda</math>  3 : <b>return</b> <math>\mathcal{F}[t, x, \mathbf{z}]</math></p>

Fig. 21: **Hybrid<sub>1</sub>** of request privacy proof.

**Hybrid<sub>2</sub>**: In this hybrid, we commit to  $\mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{0}_{h \times (\ell+m)})$ , where  $\mathbf{0}_{h \times (\ell+m)}$  is an all zero matrix of dimension  $h \times (\ell + m)$ , instead of computing the commitment  $\mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{R}_i)$  as in the previous hybrid. This change is shown in Figure 22.

In order to show indistinguishability between **Hybrid<sub>1</sub>** and **Hybrid<sub>2</sub>**, we construct a reduction  $\mathcal{B}_2$  to the computational hiding property of COM. We note that the public parameters  $\text{pp}$  and the responses to all queries are computed as in **Hybrid<sub>1</sub>**, with the exception of the Request queries. Let  $Q_R$  denote the number of PreProcClient queries that the adversary  $\mathcal{A}$  makes. Upon receiving a PreProcClient query  $(\text{pk}, T)$ ,  $\mathcal{B}_2$  samples  $R_i \xleftarrow{\$} \chi_r^{h \times (\ell+m)}$  and submits the pair  $(m_{i,0} := \mathbf{R}_i, m_{i,1} := \text{vec} \mathbf{0}_{h \times (\ell+m)})$  to the computational hiding challenger of COM, which responds with  $\mathbf{c}_i$ . At this point  $\mathcal{B}_2$  sets  $\mathbf{c}_{r,i} := \mathbf{c}_i$ , and proceeds as in the previous hybrid to simulate the rest of the oracles.

F.Setup( $1^\lambda$ )	F.KeyGen(pp)
<ol style="list-style-type: none"> <li>1 : <math>\forall i \in [3]: \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)</math></li> <li>2 : <math>(\text{crs}_1, \tau) \leftarrow \text{NIZK}_1.\mathcal{S}_1(1^\lambda)</math></li> <li>3 : <math>\text{crs}_2 \leftarrow \text{NIZK}_2.\text{Setup}(1^\lambda)</math></li> <li>4 : <b>return</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <math>\mathbf{k} \xleftarrow{\\$} \chi_{\mathbf{k}}^m</math></li> <li>3 : <math>\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)</math></li> <li>4 : <b>return</b> <math>(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))</math></li> </ol>
F.PreProcClient(pp, pk, T)	F.Request(pp, pk, t, x, st <sub>C</sub> )
<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <math>\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\\$} \chi_r^{h \times (\ell+m)}</math></li> <li>3 : <math>\forall i \in [T]: \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{0}_{h \times (\ell+m)}; \rho_{r,i})</math></li> <li>4 : <math>\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}</math></li> <li>5 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}</math></li> <li>6 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j := 1, \perp)</math></li> <li>7 : <b>return</b> <math>(\text{st}_C, \text{req} := \{\mathbf{c}_{r,i}\}_{i \in [T]})</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j, \cdot)</math></li> <li>3 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math></li> <li>4 : <b>if</b> <math>j &gt; T</math> <b>then return</b> <math>\perp</math></li> <li>5 : <math>\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}</math></li> <li>6 : <math>\mathbf{C}_x := \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{B}_{x,t} \in R_q^{h \times m}</math></li> <li>7 : <math>\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, x; \rho_x)</math></li> <li>8 : <math>\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)</math></li> <li>9 : <math>\pi_c \leftarrow \text{NIZK}_1.\mathcal{S}_2(\text{crs}_1, \tau, \text{stmt}_1)</math></li> <li>10 : <math>\text{st}_j := (t, x, \mathbf{C}_x, \text{pk})</math></li> <li>11 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)</math></li> <li>12 : <math>\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)</math></li> <li>13 : <b>return</b> <math>(\text{st}_C, \text{req})</math></li> </ol>
F.Finalize(pp, rep, st <sub>C</sub> )	RO <sub>r</sub> ( $\mathbf{c}_{r,i}$ )
<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)</math></li> <li>3 : <b>parse</b> <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math></li> <li>4 : <b>parse</b> <math>\text{st}_j := (t, x, \mathbf{C}_x, \mathbf{c}_k)</math></li> <li>5 : <b>parse</b> <math>\text{rep} := (\mathbf{u}_x, \pi_s)</math></li> <li>6 : <math>\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})</math></li> <li>7 : <b>if</b> <math>\text{NIZK}_2.\mathcal{V}(\text{crs}_2, \pi_s, \text{stmt}_2) \neq 1</math> <b>then</b></li> <li>8 :     <b>return</b> <math>\perp</math></li> <li>9 : <math>\mathbf{z} := \lfloor \mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j} \rfloor_p \in R_p^h \quad // = \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p</math></li> <li>10 : <math>y := \text{RO}_z(t, x, \mathbf{z})</math></li> <li>11 : <b>return</b> <math>y</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>if</b> <math>\mathcal{H}[\mathbf{c}_{r,i}] = \perp</math> <b>then</b></li> <li>2 :     <math>\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\\$} R_q^{(\ell+m) \times m}</math></li> <li>3 : <b>return</b> <math>\mathcal{H}[\mathbf{c}_{r,i}]</math></li> </ol>
	RO <sub>z</sub> ( $t, z, \mathbf{z}$ )
	<ol style="list-style-type: none"> <li>1 : <b>if</b> <math>\mathcal{F}[t, x, \mathbf{z}] = \perp</math> <b>then</b></li> <li>2 :     <math>\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>3 : <b>return</b> <math>\mathcal{F}[t, x, \mathbf{z}]</math></li> </ol>

Fig. 22: **Hybrid<sub>2</sub>** of request privacy proof.

If the computational hiding challenger of COM responds with a commitment to  $m_{i,0}$ , then we are exactly in **Hybrid<sub>1</sub>**, and if it responds with a commitment to  $m_{i,1}$ , then we are in **Hybrid<sub>2</sub>**. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_2$  can break the computational hiding of COM. Since we consider here a multi-challenge variant of the computational hiding property (concretely, the  $T$ -challenge variant, where without of loss generality we assume the adversary invokes the PreProcClient oracle with the same value  $T$ ) and the adversary  $\mathcal{A}$  makes at most  $Q_{\mathcal{P}}$  queries to the PreProcClient oracle, it follows by a standard argument that

$$|\text{Adv}_{\mathcal{A},1}(\lambda) - \text{Adv}_{\mathcal{A},2}(\lambda)| \leq T \cdot Q_{\mathcal{P}} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_2}^{\text{CH}}(\lambda).$$

F.Setup( $1^\lambda$ )	F.KeyGen(pp)
<ol style="list-style-type: none"> <li>1 : <math>\forall i \in [3]: \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)</math></li> <li>2 : <math>(\text{crs}_1, \tau) \leftarrow \text{NIZK}_1.\mathcal{S}_1(1^\lambda)</math></li> <li>3 : <math>\text{crs}_2 \leftarrow \text{NIZK}_2.\text{Setup}(1^\lambda)</math></li> <li>4 : <b>return</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <math>\mathbf{k} \xleftarrow{\\$} \chi_{\mathbf{k}}^m</math></li> <li>3 : <math>\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)</math></li> <li>4 : <b>return</b> <math>(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))</math></li> </ol>
F.PreProcClient(pp, pk, T)	F.Request(pp, pk, t, x, st <sub>C</sub> )
<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <math>\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\\$} \chi_r^{h \times (\ell+m)}</math></li> <li>3 : <math>\forall i \in [T]: \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{0}_{h \times (\ell+m)}; \rho_{r,i})</math></li> <li>4 : <math>\forall i \in [T]: \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}</math></li> <li>5 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}</math></li> <li>6 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j := 1, \perp)</math></li> <li>7 : <b>return</b> <math>(\text{st}_C, \text{req} := \{\mathbf{c}_{r,i}\}_{i \in [T]})</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j, \cdot)</math></li> <li>3 : <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math></li> <li>4 : <b>if</b> <math>j &gt; T</math> <b>then return</b> <math>\perp</math></li> <li>5 : <math>\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}</math></li> <li>6 : <math>\mathbf{C}_x := \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{B}_{x,t} \in R_q^{h \times m}</math></li> <li>7 : <math>\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, 0^{ x }; \rho_x)</math></li> <li>8 : <math>\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)</math></li> <li>9 : <math>\pi_c \leftarrow \text{NIZK}_1.\mathcal{S}_2(\text{crs}_1, \tau, \text{stmt}_1)</math></li> <li>10 : <math>\text{st}_j := (t, x, \mathbf{C}_x, \text{pk})</math></li> <li>11 : <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)</math></li> <li>12 : <math>\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)</math></li> <li>13 : <b>return</b> <math>(\text{st}_C, \text{req})</math></li> </ol>
F.Finalize(pp, rep, st <sub>C</sub> )	RO <sub>r</sub> ( $\mathbf{c}_{r,i}$ )
<ol style="list-style-type: none"> <li>1 : <b>parse</b> <math>\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})</math></li> <li>2 : <b>parse</b> <math>\text{st}_C := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)</math></li> <li>3 : <b>parse</b> <math>\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}</math></li> <li>4 : <b>parse</b> <math>\text{st}_j := (t, x, \mathbf{C}_x, \mathbf{c}_k)</math></li> <li>5 : <b>parse</b> <math>\text{rep} := (\mathbf{u}_x, \pi_s)</math></li> <li>6 : <math>\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})</math></li> <li>7 : <b>if</b> <math>\text{NIZK}_2.\mathcal{V}(\text{crs}_2, \pi_s, \text{stmt}_2) \neq 1</math> <b>then</b></li> <li>8 :     <b>return</b> <math>\perp</math></li> <li>9 : <math>\mathbf{z} := \lfloor \mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j} \rfloor_p \in R_p^h \quad // = \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p</math></li> <li>10 : <math>y := \text{RO}_z(t, x, \mathbf{z})</math></li> <li>11 : <b>return</b> <math>y</math></li> </ol>	<ol style="list-style-type: none"> <li>1 : <b>if</b> <math>\mathcal{H}[\mathbf{c}_{r,i}] = \perp</math> <b>then</b></li> <li>2 :     <math>\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\\$} R_q^{(\ell+m) \times m}</math></li> <li>3 : <b>return</b> <math>\mathcal{H}[\mathbf{c}_{r,i}]</math></li> </ol>
	RO <sub>z</sub> ( $t, z, \mathbf{z}$ )
	<ol style="list-style-type: none"> <li>1 : <b>if</b> <math>\mathcal{F}[t, x, \mathbf{z}] = \perp</math> <b>then</b></li> <li>2 :     <math>\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\\$} \{0, 1\}^\lambda</math></li> <li>3 : <b>return</b> <math>\mathcal{F}[t, x, \mathbf{z}]</math></li> </ol>

Fig. 23: **Hybrid<sub>3</sub>** of request privacy proof.

**Hybrid<sub>3</sub>**: In this hybrid, we commit to  $\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, 0^{|x|})$ , where  $0^{|x|}$  is an all zero bitstring of length  $|x|$ , instead of computing the commitment  $\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, x)$  as in the previous hybrid. This change is shown in Figure 23.

The indistinguishability between **Hybrid<sub>2</sub>** and **Hybrid<sub>3</sub>** follow by the computational hiding of COM. Using a similar argumentation as in the previous hybrid, we get that

$$|\text{Adv}_{\mathcal{A},2}(\lambda) - \text{Adv}_{\mathcal{A},3}(\lambda)| \leq 2Q_R \cdot \text{Adv}_{\text{COM},\mathcal{B}_3}^{\text{CH}}(\lambda),$$

where  $\mathcal{B}_3$  is the reduction to the computational hiding of COM,  $Q_R$  denotes the number of Request oracle queries that  $\mathcal{A}$  makes, and 2 comes from the fact that here we consider a 2-challenge variant of the computational hiding property (since inside the Request oracle we compute two separate requests).

**Hybrid<sub>4</sub>**: Let BAD be the event that for a Finalize query  $(i, j, \text{rep}_0 := (\mathbf{u}_{x,0}, \pi_{s,0}), \text{rep}_1 := (\mathbf{u}_{x,1}, \pi_{s,1}))$ , it holds that either

- $\text{NIZK}_2.V(\text{crs}_2, \pi_{s,0}, (\mathbf{u}_{x,0}, \mathbf{C}_{x,i,j,b}, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,i,j,b}, \mathbf{A}_{r,i,j,b})) = 1$ , but  $\nexists (\mathbf{k}, \mathbf{e}_s, \mathbf{e}'_s, \rho_k)$  such that  $(\mathbf{c}_k = \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k) \wedge \mathbf{v}_{k,i,j,b} = \mathbf{A}_{r,i,j,b} \mathbf{k} + \mathbf{e}_s \wedge \mathbf{u}_{x,0} = \mathbf{C}_{x,i,j,b} \mathbf{k} + \mathbf{e}'_s)$ ; or
- $\text{NIZK}_2.V(\text{crs}_2, \pi_{s,1}, (\mathbf{u}_{x,1}, \mathbf{C}_{x,i,j,1-b}, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,i,j,1-b}, \mathbf{A}_{r,i,j,1-b})) = 1$ , but  $\nexists (\mathbf{k}, \mathbf{e}_s, \mathbf{e}'_s, \rho_k)$  such that  $(\mathbf{c}_k = \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k) \wedge \mathbf{v}_{k,i,j,1-b} = \mathbf{A}_{r,i,j,1-b} \mathbf{k} + \mathbf{e}_s \wedge \mathbf{u}_{x,1} = \mathbf{C}_{x,i,j,1-b} \mathbf{k} + \mathbf{e}'_s)$ ,

where  $\mathbf{C}_{x,i,j}$ ,  $\mathbf{A}_{r,i,j}$  and  $\mathbf{v}_{k,i,j}$  are values from the  $j$ -th query to the Request oracle with input  $i$  (i.e., for the  $i$ -th PreProcClient session).

We will prove that the BAD event can only happen with negligible probability due to the soundness of  $\text{NIZK}_2$  argument system. Looking ahead, we need here such a reduction to the soundness of  $\text{NIZK}_2$  because later in **Hybrid<sub>5</sub>** we do a reduction to the extractability of the commitment scheme, and there we need to ensure that the adversarial inputs, specifically the public key  $\text{pk} := \mathbf{c}_k$  is well-formed.

Clearly, we have that **Hybrid<sub>3</sub>** and **Hybrid<sub>4</sub>** are identical until the event BAD happens. Hence, we show that we can bound the probability  $\Pr[\text{BAD}]$  by constructing a reduction  $\mathcal{B}_4$  to the computational soundness of  $\text{NIZK}_2$  argument system. We note that all queries are answered as in **Hybrid<sub>3</sub>**. Let  $Q_F$  denote the number of Finalize queries that  $\mathcal{A}$  performs.  $\mathcal{B}_4$  randomly chooses an index  $q^* \in [2Q_F]$  to serve as its guess for the the proof  $\pi_s$  that will be forged by the adversary  $\mathcal{A}$  (note that each Finalize query includes two proofs, hence the guess from  $[2Q_F]$ ). Let  $(i, j, \text{rep}_0 := (\mathbf{u}_{x,0}, \pi_{s,0}), \text{rep}_1 := (\mathbf{u}_{x,1}, \pi_{s,1}))$  be the  $k$ -th query to Finalize, where  $k = \lceil q^*/2 \rceil$ , then  $\mathcal{B}_3$  outputs  $\pi_{s,0}$  if  $q^*$  is odd and outputs  $\pi_{s,1}$  otherwise. We emphasize that  $\mathcal{B}_3$  avoids recovering any part of the witness  $(\mathbf{k}, \mathbf{e}_s, \mathbf{e}'_s, \rho_k)$  (which are computationally infeasible to compute), and instead just guesses which proof would have caused the event BAD to happen. Hence, we have that

$$|\text{Adv}_{\mathcal{A},3}(\lambda) - \text{Adv}_{\mathcal{A},4}(\lambda)| \leq 2Q_F \cdot \text{Adv}_{\text{NIZK}_2, \mathcal{B}_4}^{\text{CS}}(\lambda).$$

**Hybrid<sub>5</sub>**: In this hybrid, we replace the commitment key  $\text{ck}_2$  with a trapdoor variant  $(\text{ck}_2, \text{td}) \leftarrow \text{COM.E}_1(1^\lambda)$ , and use the trapdoor  $\text{td}$  to extract the secret key  $\mathbf{k}$  from the public key  $\text{pk} := \mathbf{c}_k$ , i.e., run  $\mathbf{k} \leftarrow \text{COM.E}_2(\text{ck}_2, \text{td}, \mathbf{c}_k)$ , and save  $\mathbf{k}$  for later use. These changes are depicted in Figure 24.

In order to prove indistinguishability between **Hybrid<sub>4</sub>** and **Hybrid<sub>5</sub>** we construct an adversary  $\mathcal{B}_5$  against the extractability property of COM. We note that all queries are answered as in **Hybrid<sub>4</sub>**, and the only change here is that  $\mathcal{B}_4$  samples a trapdoor commitment key  $(\text{ck}_2, \text{td}) \leftarrow \text{COM.E}_1(1^\lambda)$ , instead of generating  $\text{ck}_2$  via the COM.Setup algorithm. Let  $Q_R^{\text{pk}}$  denote the number of Request queries that the adversary  $\mathcal{A}$  makes with distinct  $\text{pk}$  inputs. For each such query  $(i, \text{pk}, t, x_0, x_1)$ , upon  $\mathcal{B}_5$  receiving a Finalize query  $(i, j, \text{rep}_0, \text{rep}_1)$ , where it corresponds to one of the  $Q_R^{\text{pk}}$  queries made to the Request oracle,  $\mathcal{B}_4$  computes  $\mathbf{k} \leftarrow \text{COM.E}_2(\text{ck}_2, \text{td}, \text{pk} := \mathbf{c}_k)$ , and stores it in the table  $\mathcal{K}[\text{pk}] := \mathbf{k}$  (for later use). The rest of the interaction is identical to **Hybrid<sub>4</sub>**.

Due to the soundness of the server proof  $\pi_s$  we know that the adversarial inputs, and especially the public keys  $\text{pk} := \mathbf{c}_k$ , are well-formed. Hence, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_5$  can break the extractability property of COM. Moreover, since the adversary  $\mathcal{A}$  makes at most  $Q_R^{\text{pk}}$  queries to the Request oracle with distinct  $\text{pk}$ , it follows that

$$|\text{Adv}_{\mathcal{A},4}(\lambda) - \text{Adv}_{\mathcal{A},5}(\lambda)| \leq Q_R^{\text{pk}} \cdot \text{Adv}_{\text{COM}, \mathcal{B}_5}^{\text{Ext}}(\lambda).$$

**Hybrid<sub>6</sub>**: In this hybrid, we replace  $\mathbf{C}_x := \mathbf{R}\mathbf{A}_r + \mathbf{B}_{x,t} \in R_q^{h \times m}$  with  $\mathbf{C}_x^* := \mathbf{U} + \mathbf{B}_{x,t} \in R_q^{h \times m}$  for a uniformly random matrix  $\mathbf{U} \xleftarrow{\$} R_q^{h \times m}$ , as shown in Figure 25.

In order to prove indistinguishability between **Hybrid<sub>5</sub>**  $\approx$  **Hybrid<sub>6</sub>** we construct an adversary  $\mathcal{B}_6$  against the  $\text{knMLWE}_{\ell+m, m, h, \chi}$  assumption (Definition 7). We note that the public parameters  $\text{pp}$  and responses to all queries are computed as in **Hybrid<sub>5</sub>** with the exception of the Request queries. For each Request query  $(i, \text{pk}, t, x_0, x_1)$ ,  $\mathcal{B}_6$  receives from the  $\text{knMLWE}_{\ell+m, m, h, \chi}$  challenger the pairs  $(\mathbf{A}_0, \mathbf{U}_0)$  and  $(\mathbf{A}_1, \mathbf{U}_1)$ , which it uses to compute  $\mathbf{C}_{x,0} := \mathbf{U}_0 + \text{H}(x_0, t)$  and  $\mathbf{C}_{x,1} := \mathbf{U}_1 + \text{H}(x_1, t)$ . Note that  $\mathbf{C}_{x,0}$  and  $\mathbf{C}_{x,1}$  form part of the request that  $\mathcal{B}_6$  returns to  $\mathcal{A}$  as output of the Request query. Moreover,  $\mathcal{B}_6$  programs

<p><b>F.Setup(<math>1^\lambda</math>)</b></p> <hr/> 1 : $\forall i \in \{1, 3\} : \mathbf{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)$ 2 : $(\mathbf{ck}_2, \mathbf{td}) \leftarrow \text{COM.E}_1(1^\lambda)$ 3 : $(\mathbf{crs}_1, \tau) \leftarrow \text{NIZK}_1.\mathcal{S}_1(1^\lambda)$ 4 : $\mathbf{crs}_2 \leftarrow \text{NIZK}_2.\text{Setup}(1^\lambda)$ 5 : <b>return</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$	<p><b>F.KeyGen(pp)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{k} \xleftarrow{\$} \chi_{\mathbf{k}}^m$ 3 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\mathbf{ck}_2, \mathbf{k}; \rho_k)$ 4 : <b>return</b> $(\mathbf{pk} := \mathbf{c}_k, \mathbf{sk} := (\mathbf{k}, \rho_k))$
<p><b>F.PreProcClient(pp, pk, T)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\$} \chi_r^{h \times (\ell+m)}$ 3 : $\forall i \in [T] : \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\mathbf{ck}_3, \mathbf{0}_{h \times (\ell+m)}; \rho_{r,i})$ 4 : $\forall i \in [T] : \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\mathbf{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}$ 6 : $\mathbf{st}_C := (\mathbf{st}_{\text{pre}}, T, j := 1, \perp)$ 7 : <b>return</b> $(\mathbf{st}_C, \text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]})$	<p><b>F.Request(pp, pk, t, x, st<sub>C</sub>)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\mathbf{st}_C := (\mathbf{st}_{\text{pre}}, T, j, \cdot)$ 3 : $\mathbf{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}$ 4 : <b>if</b> $j > T$ <b>then return</b> $\perp$ 5 : $\mathbf{B}_{x,t} := \mathbf{H}(x, t) \in R_q^{h \times m}$ 6 : $\mathbf{C}_x := \mathbf{R}_j \mathbf{A}_{r,j} + \mathbf{B}_{x,t} \in R_q^{h \times m}$ 7 : $\mathbf{d}_x \leftarrow \text{COM.Commit}(\mathbf{ck}_1, 0^{ x }; \rho_x)$ 8 : $\mathbf{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \mathbf{ck}_1, \mathbf{ck}_3, \mathbf{A}_{r,j}, \mathbf{H}, t, j)$ 9 : $\pi_c \leftarrow \text{NIZK}_1.\mathcal{S}_2(\mathbf{crs}_1, \tau, \mathbf{stmt}_1)$ 10 : $\mathbf{st}_j := (t, x, \mathbf{C}_x, \mathbf{pk})$ 11 : $\mathbf{st}_C := (\mathbf{st}_{\text{pre}}, T, j + 1, \mathbf{st}_j)$ 12 : $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 13 : <b>return</b> $(\mathbf{st}_C, \text{req})$
<p><b>F.Finalize(pp, rep, st<sub>C</sub>)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\mathbf{H}, \{\mathbf{ck}_i\}_{i \in [3]}, \{\mathbf{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\mathbf{st}_C := (\mathbf{st}_{\text{pre}}, T, j + 1, \mathbf{st}_j)$ 3 : <b>parse</b> $\mathbf{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}$ 4 : <b>parse</b> $\mathbf{st}_j := (t, x, \mathbf{C}_x, \mathbf{c}_k)$ 5 : <b>parse</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$ 6 : $\mathbf{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \mathbf{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 7 : <b>if</b> $\text{NIZK}_2.\mathbf{V}(\mathbf{crs}_2, \pi_s, \mathbf{stmt}_2) \neq 1$ <b>then</b> 8 : <b>return</b> $\perp$ 9 : <b>if</b> $\mathcal{K}[\mathbf{pk}] = \perp$ <b>then</b> 10 : $\mathbf{k} \leftarrow \text{COM.E}_2(\mathbf{ck}_2, \mathbf{td}, \mathbf{c}_k)$ 11 : $\mathcal{K}[\mathbf{pk}] := \mathbf{k}$ 12 : $\mathbf{z} := \lfloor \mathbf{u}_x - \mathbf{R}_j \mathbf{v}_{k,j} \rfloor_p \in R_p^h \quad // = \lfloor \mathbf{B}_{x,t} \mathbf{k} \rfloor_p$ 13 : $y := \text{RO}_z(t, x, \mathbf{z})$ 14 : <b>return</b> $y$	<p><b>RO<sub>r</sub>(c<sub>r,i</sub>)</b></p> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$
	<p><b>RO<sub>z</sub>(t, z, z)</b></p> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$

Fig. 24: **Hybrid<sub>5</sub>** of request privacy proof.

the random oracle  $\text{RO}_r$ , such that on input the pair  $\mathbf{c}_{r,i}$  it returns  $\mathbf{A}_i$  provided by the  $\text{knMLWE}_{\ell+m, m, h, \chi}$  challenger (we note that this change is not highlighted in Figure 25, because the  $\mathbf{A}_i$  values returned by the  $\text{knMLWE}_{\ell+m, m, h, \chi}$  challenger are also uniformly random values from  $R_q^{(\ell+m) \times m}$ ). Upon receiving a **Finalize** query  $(i, j, \text{rep}_0 := (\mathbf{u}_{x,0}, \pi_{s,0}), \text{rep}_1 := (\mathbf{u}_{x,0}, \pi_{s,0}))$  from  $\mathcal{A}$ ,  $\mathcal{B}_6$  proceeds as in the previous hybrid, to obtain  $\mathbf{pk} := \mathbf{c}_k$  from the state  $\mathbf{st}_j$ , and then extracts  $\mathbf{k}$  from  $\mathbf{pk}$  using the commitment trapdoor. Additionally,  $\mathcal{B}_5$  uses the secret key  $\mathbf{k}$  to compute the values  $\mathbf{z}_0 := \lfloor \mathbf{u}_{x,0} - \mathbf{U}_0 \mathbf{k} - \mathbf{R}_{i,j,0} (\mathbf{v}_{k,i,j,0} - \mathbf{A}_{r,i,j,0} \mathbf{k}) \rfloor_p$  and  $\mathbf{z}_1 := \lfloor \mathbf{u}_{x,1} - \mathbf{U}_1 \mathbf{k} - \mathbf{R}_{i,j,1} (\mathbf{v}_{k,i,j,1} - \mathbf{A}_{r,i,j,1} \mathbf{k}) \rfloor_p$ , which constitute an alternative way of computing the correct  $\mathbf{z}_0$  and  $\mathbf{z}_1$  values with the help of the secret key  $\mathbf{k}$  and the randomnesses  $\mathbf{R}_{i,j,0}$  and  $\mathbf{R}_{i,j,1}$  that were

<p><b>F.Setup(<math>1^\lambda</math>)</b></p> <hr/> 1 : $\forall i \in \{1, 3\} : \text{ck}_i \leftarrow \text{COM.Setup}(1^\lambda)$ 2 : $(\text{ck}_2, \text{td}) \leftarrow \text{COM.E}_1(1^\lambda)$ 3 : $(\text{crs}_1, \tau) \leftarrow \text{NIZK}_1.\mathcal{S}_1(1^\lambda)$ 4 : $\text{crs}_2 \leftarrow \text{NIZK}_2.\text{Setup}(1^\lambda)$ 5 : <b>return</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$	<p><b>F.KeyGen(pp)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{k} \xleftarrow{\$} \chi_{\mathbf{k}}^m$ 3 : $\mathbf{c}_k \leftarrow \text{COM.Commit}(\text{ck}_2, \mathbf{k}; \rho_k)$ 4 : <b>return</b> $(\text{pk} := \mathbf{c}_k, \text{sk} := (\mathbf{k}, \rho_k))$
<p><b>F.PreProcClient(pp, pk, T)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : $\mathbf{R}_1, \dots, \mathbf{R}_T \xleftarrow{\$} \chi_r^{h \times (\ell+m)}$ 3 : $\forall i \in [T] : \mathbf{c}_{r,i} \leftarrow \text{COM.Commit}(\text{ck}_3, \mathbf{0}_{h \times (\ell+m)}; \rho_{r,i})$ 4 : $\forall i \in [T] : \mathbf{A}_{r,i} := \text{RO}_r(\mathbf{c}_{r,i}) \in R_q^{(\ell+m) \times m}$ 5 : $\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \rho_{r,i}\}_{i \in [T]}$ 6 : $\text{st}_{\text{C}} := (\text{st}_{\text{pre}}, T, j := 1, \perp)$ 7 : <b>return</b> $(\text{st}_{\text{C}}, \text{preq} := \{\mathbf{c}_{r,i}\}_{i \in [T]})$	<p><b>F.Request(pp, pk, t, x, st<sub>C</sub>)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{st}_{\text{C}} := (\text{st}_{\text{pre}}, T, j, \cdot)$ 3 : $\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}$ 4 : <b>if</b> $j > T$ <b>then return</b> $\perp$ 5 : $\mathbf{B}_{x,t} := \text{H}(x, t) \in R_q^{h \times m}$ 6 : $\mathbf{U} \xleftarrow{\$} R_q^{h \times m}$ 7 : $\mathbf{C}_x := \mathbf{U} + \mathbf{B}_{x,t} \in R_q^{h \times m}$ 8 : $\mathbf{d}_x \leftarrow \text{COM.Commit}(\text{ck}_1, 0^{ \mathbf{x} }; \rho_x)$ 9 : $\text{stmt}_1 := (\mathbf{c}_{r,j}, \mathbf{C}_x, \mathbf{d}_x, \text{ck}_1, \text{ck}_3, \mathbf{A}_{r,j}, \text{H}, t, j)$ 10 : $\pi_c \leftarrow \text{NIZK}_1.\mathcal{S}_2(\text{crs}_1, \tau, \text{stmt}_1)$ 11 : $\text{st}_j := (t, x, \mathbf{C}_x, \text{pk})$ 12 : $\text{st}_{\text{C}} := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)$ 13 : $\text{req} := (\mathbf{d}_x, \mathbf{C}_x, \pi_c, j)$ 14 : <b>return</b> $(\text{st}_{\text{C}}, \text{req})$
<p><b>F.Finalize(pp, rep, st<sub>C</sub>)</b></p> <hr/> 1 : <b>parse</b> $\text{pp} := (\text{H}, \{\text{ck}_i\}_{i \in [3]}, \{\text{crs}_i\}_{i \in [2]})$ 2 : <b>parse</b> $\text{st}_{\text{C}} := (\text{st}_{\text{pre}}, T, j + 1, \text{st}_j)$ 3 : <b>parse</b> $\text{st}_{\text{pre}} := \{\mathbf{c}_{r,i}, \mathbf{A}_{r,i}, \mathbf{R}_i, \mathbf{v}_{k,i}, \rho_{r,i}\}_{i \in [T]}$ 4 : <b>parse</b> $\text{st}_j := (t, x, \mathbf{C}_x, \mathbf{c}_k)$ 5 : <b>parse</b> $\text{rep} := (\mathbf{u}_x, \pi_s)$ 6 : $\text{stmt}_2 := (\mathbf{u}_x, \mathbf{C}_x, \mathbf{c}_k, \text{ck}_2, \mathbf{v}_{k,j}, \mathbf{A}_{r,j})$ 7 : <b>if</b> $\text{NIZK}_2.\text{V}(\text{crs}_2, \pi_s, \text{stmt}_2) \neq 1$ <b>then</b> 8 : <b>return</b> $\perp$ 9 : <b>if</b> $\mathcal{K}[\text{pk}] = \perp$ <b>then</b> 10 : $\mathbf{k} \leftarrow \text{COM.E}_2(\text{ck}_2, \text{td}, \mathbf{c}_k)$ 11 : $\mathcal{K}[\text{pk}] := \mathbf{k}$ 12 : $\mathbf{z} := \lfloor \mathbf{u}_x - \text{UK}[\text{pk}] - \mathbf{R}_j(\mathbf{v}_{k,j} - \mathbf{A}_{r,j}\mathbf{k}) \rfloor_p \in R_p^h$ 13 : $y := \text{RO}_z(t, x, \mathbf{z})$ 14 : <b>return</b> $y$	<p><b>RO<sub>r</sub>(c<sub>r,i</sub>)</b></p> <hr/> 1 : <b>if</b> $\mathcal{H}[\mathbf{c}_{r,i}] = \perp$ <b>then</b> 2 : $\mathcal{H}[\mathbf{c}_{r,i}] \xleftarrow{\$} R_q^{(\ell+m) \times m}$ 3 : <b>return</b> $\mathcal{H}[\mathbf{c}_{r,i}]$
	<p><b>RO<sub>z</sub>(t, z, z)</b></p> <hr/> 1 : <b>if</b> $\mathcal{F}[t, x, \mathbf{z}] = \perp$ <b>then</b> 2 : $\mathcal{F}[t, x, \mathbf{z}] \xleftarrow{\$} \{0, 1\}^\lambda$ 3 : <b>return</b> $\mathcal{F}[t, x, \mathbf{z}]$

Fig. 25: **Hybrid<sub>6</sub>** of request privacy proof.

sampled during the Request query. Finally,  $\mathcal{B}$  computes the output  $y_0$  and  $y_1$  as before and returns it to the adversary  $\mathcal{A}$ .

If the  $\text{knMLWE}_{\ell+m,m,h,\chi}$  challenger provided a knapsack MLWE instance, i.e.,  $\mathbf{U} := \mathbf{R}\mathbf{A}$ , for some  $\mathbf{R} \in R_q^{h \times (\ell+m)}$ , then the view is as in **Hybrid<sub>5</sub>**, and if it provided a uniformly random matrix  $\mathbf{U} \in R_q^{h \times m}$ , then the view is as in **Hybrid<sub>6</sub>**. Therefore, if  $\mathcal{A}$  can distinguish between the two hybrids with non-negligible advantage, then  $\mathcal{B}_6$  can break the  $\text{knMLWE}_{\ell+m,m,h,\chi}$  assumption. Since we consider here a multi-challenge variant of the  $\text{knMLWE}_{\ell+m,m,h,\chi}$  problem (concretely, the 2-challenge variant) and the adversary  $\mathcal{A}$  makes

at most  $Q_R$  queries to the Request oracle, it follows by a standard argument that

$$|\text{Adv}_{\mathcal{A},5}(\lambda) - \text{Adv}_{\mathcal{A},6}(\lambda)| \leq 2Q_R \cdot \text{Adv}_{\mathcal{B}_6}^{\text{knMLWE}_{\ell+m,m,h,x}}(\lambda).$$

We note that at this point the transcript observed by the adversary  $\mathcal{A}$  is independent of the challenge bit  $b$ , and hence  $\text{Adv}_{\mathcal{A},6}(\lambda) = 1/2$ . Putting everything together, we obtain

$$\begin{aligned} \text{Adv}_{\mathcal{F},\mathcal{A},\text{RO}_r,\text{RO}_z}^{\text{po-priv-2}}(\lambda) &\leq \text{Adv}_{\text{NIZK}_1,\mathcal{B}_1}^{\text{CZK}}(\lambda) + T \cdot Q_P \cdot \text{Adv}_{\text{COM},\mathcal{B}_2}^{\text{CH}}(\lambda) + 2Q_R \cdot \text{Adv}_{\text{COM},\mathcal{B}_3}^{\text{CH}}(\lambda) \\ &\quad + 2Q_F \cdot \text{Adv}_{\text{NIZK}_2,\mathcal{B}_4}^{\text{CS}}(\lambda) + Q_R^{\text{pk}} \cdot \text{Adv}_{\text{COM},\mathcal{B}_5}^{\text{Ext}}(\lambda) + 2Q_R \cdot \text{Adv}_{\mathcal{B}_6}^{\text{knMLWE}}(\lambda), \end{aligned}$$

as claimed. This completes the proof of Theorem 6.  $\square$

## E Issues in Prior Constructions and Models

### E.1 Issues in [ADDS21] and [ADDG24]

We describe here two flaws present in the security proofs of the existing lattice-based OPRF constructions. The first flaw is from the [ADDS21] paper, which aims to prove their construction secure in the simulation based setting (we refer the reader to [ADDS21] for details regarding the security model). The flaw appears during the Query phase of the malicious client security proof (given in [ADDS21, Lemma 8]), where for each query  $(\mathbf{c}_x, \pi)$  from the malicious client, the simulator needs to extract the client's secret input  $x$  from the NIZK argument  $\pi$  in order to complete the reduction. However, the authors make use of a NIZK argument system that is based on Fiat-Shamir transform [YAZ<sup>+</sup>19], and hence, is not straight-line extractable, i.e., cannot extract the witness without rewinding. This means that each query of the malicious client incurs a loss in the soundness of the NIZK argument system, and since a malicious client can make  $Q = \text{poly}(\lambda)$  distinct queries, it means we will end up with an exponential decay in soundness. We note that this issue can be fixed by using the Katsumata transform [Kat21] to obtain a lattice-based straight-line extractable NIZK argument system, albeit by incurring a performance penalty.

The second flaw appears in the recent [ADDG24] paper. Specifically, during the security proof of request privacy against malicious servers (POPRIV2) (given in [ADDG24, Theorem 4]), one of the reductions requires access to the server's secret key  $\text{sk}$ . The authors attempt to solve this dilemma by requiring the server to give its  $\text{sk}$  to the reduction algorithm, however, this is in stark contrast to the request privacy game given in Definition 4 (also used in [ADDG24]). Since their construction already includes NIZK proofs, one potential way to patch this issue is to instead extract the  $\text{sk}$  from a NIZK proof. However, similar to the previously explained issue, one needs to consider here a straight-line extractable proof system, which degrades the performance of the scheme. An alternative solution is to use an extractable commitment scheme and include a commitment to  $\text{sk}$  as part of the public key  $\text{pk}$ , which is exactly what we do in our construction given in Section 2.1, and during our request privacy security proof given in Appendix D.4.

### E.2 Uniqueness and Key Binding

In [TCR<sup>+</sup>22], the authors defined a property called *uniqueness*, which ensures that POPRF outputs are unique even in the case of a malicious server.

**Definition 20 (Uniqueness [TCR<sup>+</sup>22]).** *We say that a partial oblivious PRF (POPRF)  $F$  is unique against malicious servers, if for all PPT adversaries  $\mathcal{A}$  the following advantage is  $\text{negl}(\lambda)$ ,*

$$\text{Adv}_{\mathcal{F},\mathcal{A},\text{RO}}^{\text{po-uniq}}(\lambda) = \Pr [\text{POUNIQ}_{\mathcal{F},\mathcal{A},\text{RO}}(\lambda) = 1],$$

where the experiment POUNIQ is defined as follows:

POUNIQ <sub>F,A,RO</sub> ( $\lambda$ )	Request( $\text{sk}, \text{aux}, \text{pk}, t, x$ )
1 : $R := [], q := 0$	1 : <b>assert</b> F.Wellformed( $\text{sk}, \text{pk}, \text{aux}$ )
2 : $\text{pp} \leftarrow \text{F.Setup}(1^\lambda)$	2 : $(\text{st}, \text{req}) \leftarrow \text{F.Request}_{\text{pp}}^{\text{RO}}(\text{pk}, t, x)$
3 : $\text{st}_{\text{RO}} \leftarrow \text{RO.Init}(1^\lambda)$	3 : $q := q + 1$
4 : $(i, j, \text{rep}_i, \text{rep}_j) \leftarrow \mathcal{A}^{\text{Request,RO}}(\text{pp})$	4 : $R[q] := (\text{st}, \text{pk}, t, x)$
5 : <b>assert</b> $1 \leq i \leq j \leq q$	5 : <b>return req</b>
6 : $(\text{st}_i, \text{pk}_i, t_i, x_i) := R[i]$	
7 : $(\text{st}_j, \text{pk}_j, t_j, x_j) := R[j]$	
8 : $y_i \leftarrow \text{F.Finalize}_{\text{pp}}^{\text{RO}}(\text{rep}_i, \text{st}_i)$	
9 : $y_j \leftarrow \text{F.Finalize}_{\text{pp}}^{\text{RO}}(\text{rep}_j, \text{st}_j)$	
10 : <b>if</b> $y_i = \perp \vee y_j = \perp$ <b>then</b>	
11 : <b>return</b> 0	
12 : <b>return</b> $((\text{pk}_i, t_i, x_i) = (\text{pk}_j, t_j, x_j)) \wedge (y_i \neq y_j)$	

The highlighted part is included to address rogue key attacks using the knowledge of secret key model [Bol03]. If the Wellformed predicate does not require any auxiliary information to verify the secret/public key pair  $(\text{sk}, \text{pk})$ , then we consider that  $\text{aux} := \perp$ .

Tyagi et al. [TCR+22] showed that uniqueness follows from the correctness and request privacy against malicious servers (POPRIV2) of the POPRF. However, since correctness only holds against well-formed public keys and uniqueness definition considers a malicious adversary, they proved this in the knowledge of secret key (KOSK) model [Bol03]. In KOSK, the adversary reveals its secret key and this is used to check the well-formedness of the secret/public key pair (e.g., by using the Wellformed predicate in Definition 20). We note that such a model can be instantiated by including an extractable proof of knowledge or commitment, from which the secret key can be extracted and the proof may proceed as in the KOSK model.

Nevertheless, the reduction from uniqueness to request privacy that was given by Tyagi et al. [TCR+22, Appendix H], only works when there is a bijection between the secret and public keys, i.e., for each public key  $\text{pk}$  there is a unique  $\text{sk}$ . This comes from the fact that during the reduction they consider that the output of a valid uniqueness adversary would result in two tuples  $(\text{pk}_i, \text{sk}_i, t_i, x_i) = (\text{pk}_j, \text{sk}_j, t_j, x_j)$ . However, the uniqueness definition (see Definition 20) only guarantees that  $(\text{pk}_i, t_i, x_i) = (\text{pk}_j, t_j, x_j)$ , i.e., we do not have the guarantee that  $\text{sk}_i = \text{sk}_j$ . Although, this holds in group-based OPRF constructions, such as in [TCR+22], this does not necessarily hold in lattice-based constructions. For example, in [ADDS21] the public key is of the form  $\mathbf{c} := \mathbf{a} \cdot k + \mathbf{e} \bmod q$ , such that for a fixed (non-zero)  $\mathbf{a}$ , there may exist many pairs of secret keys  $(k, \mathbf{e})$  that satisfy the public key equation.

In order to provide a more robust uniqueness guarantees, we define here an additional property for a POPRF called *key binding*, which ensures that a public key is bound to a particular secret key.

**Definition 21 (Key Binding).** We say that a partial oblivious PRF (POPRF)  $F$  is key binding with respect to a Wellformed predicate, if for all PPT adversaries  $\mathcal{A}$  the following advantage is  $\text{negl}(\lambda)$ ,

$$\text{Adv}_{F,\mathcal{A},\text{RO}}^{\text{po-key-bind}}(\lambda) = \Pr[\text{POKBIND}_{F,\mathcal{A},\text{RO}}(\lambda) = 1],$$

where the experiment POKBIND is defined as follows:

$\text{POKBIND}_{F,\mathcal{A},\text{RO}}(\lambda)$ <hr style="border: 0.5px solid black;"/> $1 : \text{pp} \leftarrow F.\text{Setup}(1^\lambda)$ $2 : \text{st}_{\text{RO}} \leftarrow \text{RO}.\text{Init}(1^\lambda)$ $3 : (\text{pk}, \text{sk}, \text{sk}', \text{aux}, \text{aux}') \leftarrow \mathcal{A}^{\text{RO}}(\text{pp})$ $4 : \text{assert } \text{sk} \neq \text{sk}'$ $5 : \text{return } F.\text{Wellformed}(\text{sk}, \text{pk}, \text{aux}) \wedge F.\text{Wellformed}(\text{sk}', \text{pk}, \text{aux}')$
--

Next, we show that our construction from Figure 9 satisfies key binding.

**Theorem 7.** *If COM is computationally binding, then the OPRF construction F from Figure 9 is key binding.*

*Proof.* First we define a Wellformed predicate for our key binding definition. In our construction the secret key is of the form  $\text{sk} := \mathbf{k}$ , for  $\mathbf{k} \xleftarrow{\$} \chi_k^m$ , and the public key is computed as  $\text{pk} := \mathbf{c}_k \leftarrow \text{COM}.\text{Commit}(\text{ck}_2, \mathbf{k})$ , for some commitment key  $\text{ck}_2 \leftarrow \text{COM}.\text{Setup}(1^\lambda)$ . Hence, our Wellformed predicate can be defined as the verification algorithm of the commitment scheme COM, i.e.,  $F.\text{Wellformed}(\text{sk}, \text{pk}, \text{aux}) = \text{COM}.\text{Verify}(\text{ck}_2, \mathbf{c}_k, \rho_k, \mathbf{k})$ , where  $\text{aux} := \rho_k$  is the decommitment information.

In order to show that our construction satisfies key binding we construct a reduction  $\mathcal{B}$  to the computational binding property of COM.  $\mathcal{B}$  receives the commitment key  $\text{ck}^*$  from the computational binding challenger of COM, sets  $\text{ck}_2 := \text{ck}^*$ , generates the rest of the public parameters  $\text{pp}$  as in Figure 9, and sends  $\text{pp}$  to  $\mathcal{A}$ . Upon receiving the tuple  $(\text{pk} := \mathbf{c}_k, \text{sk} := \mathbf{k}, \text{sk}' := \mathbf{k}', \text{aux} := \rho_k, \text{aux}' := \rho'_k)$  from  $\mathcal{A}$ ,  $\mathcal{B}$  forwards the tuple  $(\mathbf{c}_k, \rho_k, \rho'_k, \mathbf{k}, \mathbf{k}')$  to the computational binding challenger of COM.

We note that the winning conditions for the key binding of F (Definition 21) and computational binding of COM (Definition 16) are identical in this case. Hence, if  $\mathcal{A}$  can win the key binding experiment with non-negligible advantage, then  $\mathcal{B}$  can win the computational binding experiment with non-negligible advantage.  $\square$

**Theorem 8.** *For every PPT adversary  $\mathcal{A}$  against the uniqueness of F, there exist PPT adversaries  $\mathcal{B}_1$  and  $\mathcal{B}_2$  against the key binding and request privacy of F, respectively, such that*

$$\text{Adv}_{F,\mathcal{A},\text{RO}}^{\text{po-uniq}}(\lambda) \leq \text{Adv}_{F,\mathcal{B}_1,\text{RO}}^{\text{po-key-bind}}(\lambda) + \text{Adv}_{F,\mathcal{B}_2,\text{RO}}^{\text{po-priv-2}}(\lambda).$$

*Proof.* We note that the adversary  $\mathcal{A}$  has access to the oracles Request and RO. We consider a single intermediate hybrid before reducing the uniqueness to the request privacy against malicious servers (POPRIV2). Let  $\text{Adv}_{\mathcal{A},i}(\lambda)$  denote the advantage of  $\mathcal{A}$  in **Hybrid** <sub>$i$</sub>  and let **Hybrid** <sub>$i$</sub>   $\approx$  **Hybrid** <sub>$i+1$</sub>  denote  $|\Pr[\text{Hybrid}_i = 1] - \Pr[\text{Hybrid}_{i+1} = 1]| \leq \text{negl}(\lambda)$ .

**Hybrid**<sub>0</sub>: This corresponds to the uniqueness experiment  $\text{POUNIQ}_{F,\mathcal{A},\text{RO}}$  (Definition 20). Hence, we have that

$$\text{Adv}_{F,\mathcal{A},\text{RO}}^{\text{po-uniq}}(\lambda) = \text{Adv}_{\mathcal{A},0}(\lambda).$$

**Hybrid**<sub>1</sub>: Let BAD be the event that the adversary queries the Request oracle with two queries  $(\text{sk}, \text{aux}, \text{pk}, t, x)$  and  $(\text{sk}', \text{aux}', \text{pk}', t', x')$ , such that  $\text{pk} = \text{pk}'$  and  $\text{sk} \neq \text{sk}'$ . If BAD happens, then the challenger aborts.

Clearly, we have that **Hybrid**<sub>0</sub> and **Hybrid**<sub>1</sub> are identical until the event BAD happens. Hence, we show that we can bound the probability  $\Pr[\text{BAD}]$  by constructing a reduction  $\mathcal{B}_1$  to the key binding of F.  $\mathcal{B}_1$  receives the public parameters  $\text{pp}$  from the key binding challenger, and forwards  $\text{pp}$  to  $\mathcal{A}$ . When  $\mathcal{A}$  makes a RO query, then  $\mathcal{B}_1$  forwards the same query to its own RO oracle and relays the answer back to  $\mathcal{A}$ . For answering the Request oracle queries,  $\mathcal{B}_1$  proceeds exactly as in **Hybrid**<sub>0</sub>, but when  $\mathcal{A}$  submits two queries  $(\text{sk}, \text{aux}, \text{pk}, t, x)$  and  $(\text{sk}', \text{aux}', \text{pk}', t', x')$  that satisfy  $\text{pk} = \text{pk}'$  and  $\text{sk} \neq \text{sk}'$ , then  $\mathcal{B}$  submits the tuple  $(\text{pk}, \text{sk}, \text{sk}', \text{aux}, \text{aux}')$  to the key binding challenger. Hence, we have that  $\mathcal{B}_1$  wins the key binding experiment with the same probability that the event BAD happens, and therefore, we have that

$$|\text{Adv}_{\mathcal{A},0}(\lambda) - \text{Adv}_{\mathcal{A},1}(\lambda)| \leq \text{Adv}_{F,\mathcal{B}_1,\text{RO}}^{\text{po-key-bind}}(\lambda),$$

and in particular  $\mathbf{Hybrid}_0 \approx \mathbf{Hybrid}_1$ .

All that remains is to bound the advantage of the adversary  $\mathcal{A}$  against the uniqueness experiment in  $\mathbf{Hybrid}_1$ , i.e., bound  $\text{Adv}_{\mathcal{A},1}(\lambda)$ . We can do so by constructing a reduction  $\mathcal{B}_2$  to the request privacy of  $F$ . Analogous to the previous reduction,  $\mathcal{B}_2$  receives the public parameters  $\mathbf{pp}$  from the request privacy challenger, and forwards  $\mathbf{pp}$  to  $\mathcal{A}$ . Moreover, RO queries of  $\mathcal{A}$  are forwarded to the RO oracle of  $\mathcal{B}_2$  and the response relayed back to  $\mathcal{A}$ . Upon receiving a Request query  $(\text{sk}, \text{aux}, \text{pk}, t, x)$  from  $\mathcal{A}$ ,  $\mathcal{B}_2$  makes a Request query to its oracle (provided by the request privacy challenger) with the input  $(\text{pk}, t, x, x)$ , received back a pair of request messages  $(\text{req}_0, \text{req}_1)$  and sends  $\text{req}_0$  to  $\mathcal{A}$ . Moreover,  $\mathcal{B}_2$  stores the tuple  $(\text{pk}, \text{sk}, t, x, \text{req}_1)$  associated with the Request query.

When  $\mathcal{A}$  outputs the responses  $(i, j, \text{rep}_{i,0}, \text{rep}_{j,0})$ , then  $\mathcal{B}_2$  retrieves the tuples  $(\text{pk}_i, \text{sk}_i, t_i, x_i, \text{req}_{i,1})$  and  $(\text{pk}_j, \text{sk}_j, t_j, x_j, \text{req}_{j,1})$  that it previously stored. We consider the case where  $\mathcal{A}$  wins the uniqueness experiment, and hence, in this case we can drop the subscripts as  $(\text{pk}_i, \text{sk}_i, t_i, x_i) = (\text{pk}_j, \text{sk}_j, t_j, x_j)$  (note that  $\text{sk}_i = \text{sk}_j$  is guaranteed by the key binding property of  $F$  that we previously proved).  $\mathcal{B}_2$  calls its Finalize oracle with  $(i, \text{rep}_{i,0}, \text{rep}_{i,1})$ , where  $\text{rep}_{i,0}$  is provided by  $\mathcal{A}$  and  $\text{rep}_{i,1}$  is generated honestly by  $\mathcal{B}_2$ . By the correctness of  $F$ , we have that the value returned from the honest flow with  $\text{rep}_{i,1}$  is equal to  $y := F.\text{Eval}(\text{sk}, t, x)$ . If Finalize oracle returns two different values, then by matching which position  $y$  is returned reveals the challenge bit and allows  $\mathcal{B}_2$  to win the request privacy experiment. On the other hand, if Finalize returns two identical values, then  $\mathcal{B}_2$  repeats the above for the  $j$ -th query. If  $\mathcal{A}$  wins the uniqueness experiment, then we know that Finalize will not return identical values for the  $j$ -th query, and hence, the challenge bit will be revealed, allowing  $\mathcal{B}_2$  to win the request privacy experiment. Therefore,  $\mathcal{B}_2$  outputs the correct challenge bit in the request privacy experiment, whenever  $\mathcal{A}$  wins uniqueness. Hence, we have that

$$\text{Adv}_{\mathcal{A},1}(\lambda) \leq \text{Adv}_{F, \mathcal{B}_2, \text{RO}}^{\text{po-priv-2}}(\lambda).$$

Putting everything together, we obtain

$$\text{Adv}_{F, \mathcal{A}, \text{RO}}^{\text{po-uniq}}(\lambda) \leq \text{Adv}_{F, \mathcal{B}_1, \text{RO}}^{\text{po-key-bind}}(\lambda) + \text{Adv}_{F, \mathcal{B}_2, \text{RO}}^{\text{po-priv-2}}(\lambda),$$

as claimed. This completes the proof of Theorem 8.  $\square$

## F Instantiation of the Commitment Scheme

From our analyses in Section 4.1, we can observe that the main properties we need from the commitment scheme are hiding the message (without trapdoor), binding, and enabling extraction of the message using the trapdoor. This is effectively an encryption scheme, and we will employ a variant of Regev's encryption [Reg05] over module lattices for our purposes. Given this Regev-style instantiation is quite well-known by now, we do not delve into too many technical details here and refer the reader to the references provided below for more details. We also note again that for the  $\mathbf{d}_x$  commitment to the binary input  $x \in \{0, 1\}^L$ , we can simply employ a standard lattice-based encryption scheme. Therefore, the rest of the section focuses on the other commitments.

For certain instantiations of the message mapping  $H$  (such as BLMR13 PRF [BLMR13]), we need the server's secret key  $\mathbf{k}$  to have large coefficients (i.e., unbounded mod  $q$ ). As a result, we need to be able to allow the decryption of large message inputs in the encryption scheme. For this, we can employ the technique (implicit in [GSW13]) described more explicitly in [LNPS21]. Suppose we want to commit to an  $N$ -dimensional vector over  $R_q$ , and define  $M := N$  if the committed message has small coefficients or  $M := 2N$  if the committed message has large coefficients. For readability, we write  $\sqrt{q}$  to mean  $\lceil \sqrt{q} \rceil$  below.

- **Setup** samples  $\mathbf{A} \xleftarrow{\$} R_q^{n \times (n+M+w)}$ . In the case we want to generate the commitment key with a trapdoor, then we set  $\mathbf{b}_{i,j}^\top = \mathbf{s}_{i,j}^\top \mathbf{A} + \mathbf{e}_{i,j}^\top$  for  $j = 0, 1$  and  $i = 1, \dots, N$ , where  $\mathbf{s}_{i,j}$ 's are MLWE secret keys where each entry is sampled from a secret distribution  $\chi_s$  and  $\mathbf{e}_{i,j}$ 's are short error vectors where each entry is sampled from an error distribution  $\chi_e$ . If we don't want the trapdoor, then we simply sample  $\mathbf{b}_{i,j}$ 's uniformly at random. The commitment key is then set as  $\text{ck} = (\mathbf{A}, \mathbf{b}_{1,0}, \mathbf{b}_{1,1}, \dots, \mathbf{b}_{N,0}, \mathbf{b}_{N,1})$ .

- To **Commit** to a message  $\mathbf{m} = (m_1, \dots, m_N) \in R^N$ , we proceed as follows:
  1. If  $\mathbf{m}$  has *small* coefficients compared to  $q$  with  $\|m_i\|_\infty \leq q_m$ , then we first sample a short randomness vector  $\mathbf{r} \leftarrow \chi_e^{n+N+w}$  and set  $\mathbf{c} = \mathbf{A}\mathbf{r}$ . Then, we compute  $c_i = q_m \mathbf{b}_{i,0}^\top \mathbf{r} + m_i$  for  $i = 1, \dots, N$ . The function outputs  $(\mathbf{c}, c_1, \dots, c_N) \in R_q^{n+N}$ .
  2. If  $\mathbf{m}$  has *large* coefficients compared to  $q$  (i.e., unbounded mod  $q$ ), then we first sample a short randomness vector  $\mathbf{r} \in \chi_e^{n+2N+w}$  and set  $\mathbf{c} = \mathbf{A}\mathbf{r}$ . Then, we compute  $c_{i,0} = \mathbf{b}_{i,0}^\top \mathbf{r} + m_i$  and  $c_{i,1} = \mathbf{b}_{i,1}^\top \mathbf{r} + \sqrt{q}m_i$ . The function outputs  $(\mathbf{c}, c_{1,0}, c_{1,1}, \dots, c_{N,0}, c_{N,1}) \in R_q^{n+2N}$ .
- To **Extract** a message  $m_i$  from  $c_i$  for  $i \in [N]$  using the trapdoor  $\mathbf{s}_{i,j}$ 's,
  1. If  $\mathbf{m}$  has *small* coefficients<sup>14</sup> compared to  $q$ , we will be given  $(\mathbf{c}, c_1, \dots, c_N) \in R_q^{n+N}$ . Then, we compute  $u_i := c_i - q_m \mathbf{s}_{i,0}^\top \mathbf{c} = q_m \mathbf{e}_{i,0}^\top \mathbf{r} + m_i \bmod q$ . Then, we round off the error term  $\mathbf{e}_{i,0}^\top \mathbf{r}$  by computing  $u_i \bmod q_m$  to recover each  $m_i$ .
  2. If  $\mathbf{m}$  has *large* coefficients<sup>15</sup> compared to  $q$  (i.e., unbounded mod  $q$ ), we will be given  $(\mathbf{c}, c_{1,0}, c_{1,1}, \dots, c_{N,0}, c_{N,1}) \in R_q^{n+2N}$ . Then, we compute  $u_{i,0} = c_{i,0} - \mathbf{s}_{i,0}^\top \mathbf{c} = \mathbf{e}_{i,0}^\top \mathbf{r} + m_i$ , and  $u_{i,1} = c_{i,1} - \mathbf{s}_{i,1}^\top \mathbf{c} = \mathbf{e}_{i,1}^\top \mathbf{r} + \sqrt{q}m_i$ . Now compute  $e' := u_1 - \sqrt{q}u_0 \bmod \sqrt{q}$  and  $m_i := (u_1 + e')/\sqrt{q}$ .

We note that the server's commitment to the key  $\mathbf{k}$  can have small-coefficient message input when BP14 PRF is used; while for BLMR13 PRF, large-coefficient message support is needed. On the other hand, the client's commitment to  $(\mathbf{R}, x)$  has always small coefficients in the message (regardless of the PRF) when e.g. the bits of the input message are represented as the coefficients of a polynomial(s).

As already observed in earlier work such as [LNPS21, ESZ22a], we note that the above encryption scheme is effectively an extractable version of the BDLOP commitment scheme [BDL<sup>+</sup>18]. The (plain) BDLOP commitment is employed for instance in LNP22 proof [LNP22]. Therefore, we can optimize the performance by calculating a single BDLOP commitment (instead of multiple times) when both the underlying NIZK and our OPRF protocol require such a commitment.

As shown in [BDL<sup>+</sup>18], the commitment scheme is

1. (computationally) hiding<sup>16</sup> if  $\text{MLWE}_{n+M, q, w, \chi_e, \chi_s}$  is hard, and
2. (computationally) binding if  $\text{MSIS}_{n, n+M+w, 2\beta_r}$  is hard where  $\beta_r$  denotes the bound on  $\|\mathbf{r}\|$  proven by the NIZK proof.

We note that in our case, we prove the commitment opening relations by the client ( $\mathbf{c}_r = \text{COM.Commit}(\dots)$ ) and the server ( $\mathbf{c}_k = \text{COM.Commit}(\dots)$ ) *without* a relaxation/approximation factor. Hence, the MSIS bound  $\beta_{\text{SIS}}$  is tighter (compared to what is provided in [BDL<sup>+</sup>18]) and does not involve terms depending on the size of the relaxation factor. We also note that commitment to  $\mathbf{R}$  is done by committing to the rows of  $\mathbf{R}$  (under one commitment).

**Extraction/decryption correctness.** In our NIZK proofs employed in LeOPaRd, we prove that all error, randomness, and/or message (if small) coefficients are much smaller than the system modulus  $q$ . Therefore, the error terms described above to be removed will have a quite small infinity norm compared to the modulus sizes we need for our OPRF. Therefore, extraction/decryption correctness requirements are easily met for well-formed commitments (with probability 1) under typical parameters. Nevertheless, we explicitly state the required bounds below.

1. For *small* messages with  $\|m_i\|_\infty \leq q_m$ , we need

$$\|q_m \mathbf{e}_{i,0}^\top \mathbf{r} + m_i\|_\infty < q/2.$$

<sup>14</sup> Note that well-formedness of the commitment along with having small coefficients in error, randomness, and message will be proven via a NIZK proof.

<sup>15</sup> Note that well-formedness of the commitment along with “double-encryption” of  $(m, \sqrt{q}m)$  will be proven via a NIZK proof. The double-encryption proof often comes for free in communication since it is a simple linear proof over  $R_q$  (see, e.g., [LNP22, Section 3]).

<sup>16</sup> In fact, the commitment outputs are (computationally) indistinguishable from uniformly random vectors over  $R_q$ .

2. For *large* messages, we need

$$\|\mathbf{e}_{i,j}^\top \mathbf{r}\|_\infty < \sqrt{q}/4.$$

To argue that a commitment key with a trapdoor is indistinguishable from a regular commitment key, we require the hardness of  $\text{MLWE}_{q,n,\chi_e,\chi_s}$  so that  $\mathbf{b}_{i,j}^\top = \mathbf{s}_{i,j}^\top \mathbf{A} + \mathbf{e}_{i,j}^\top$  is indistinguishable from random. For the client's commitment  $\mathbf{c}_r$ , we assume the error entries are sampled from the same distribution  $\chi_r$  as the entries of  $\mathbf{R}$  in **LeOPaRd** as the client security already relies on **MLWE** with its secret dimension  $\ell$  and error distribution  $\chi_r$ . Therefore, we get the dimension parameter  $n_c$  for the client's commitment as  $n_c = \ell$  (provided the relevant  $\text{MSIS}_{n_c,2\beta_r}$  problem is hard, which is easily satisfied for our parameter settings). For the server's commitment  $\mathbf{c}_k$  as its public key, we assume the error coefficients are sampled from a uniform distribution with standard deviation  $\sigma_0$  as the server security already relies on **MLWE** with such errors and secret key dimension  $m$ . Therefore, we get the dimension parameter  $n_s$  for the server's commitment as  $n_s = m$  (provided the relevant  $\text{MSIS}_{n_s,2\beta_r}$  problem is hard, which is easily satisfied for our parameter settings).