

Homomorphic Evaluation of LWR-based PRFs and Application to Transciphering

Amit Deo
Zama, France

Marc Joye
Zama, France

Benoît Libert
Zama, France

Benjamin R. Curtis
Zama, France

Mayeul de Bellabre
Zama, France

Abstract. Certain applications such as FHE transciphering require randomness while operating over encrypted data. This randomness has to be obliviously generated in the encrypted domain and remain encrypted throughout the computation. Moreover, it should be guaranteed that independent-looking random coins can be obliviously generated for different computations.

In this work, we consider the homomorphic evaluation of pseudorandom functions (PRFs) with a focus on practical lattice-based candidates. In the homomorphic PRF evaluation setting, given a fully homomorphic encryption of the PRF secret key s , it should be possible to homomorphically compute encryptions of PRF evaluations $\{\text{PRF}_s(x_i)\}_{i=1}^M$ for public inputs $\{x_i\}_{i=1}^M$. We consider this problem for PRF families based on the hardness of the Learning-With-Rounding (LWR) problem introduced by Banerjee, Peikert and Rosen (Eurocrypt '12). We build on the random-oracle variant of a PRF construction suggested by Banerjee *et al.* and demonstrate that it can be evaluated using only two sequential programmable bootstraps in the TFHE homomorphic encryption scheme. We also describe several modifications of this PRF—which we prove as secure as the original function—that support homomorphic evaluations using only one programmable bootstrap per slot.

Numerical experiments were conducted using practically relevant FHE parameter sets from the TFHE-rs library. Our benchmarks show that a throughput of about 1000 encrypted pseudorandom bits per second (resp. 900 encrypted pseudorandom bits per second) can be achieved on an AWS hpc7a.96xlarge machine (resp. on a standard laptop with an Apple M2 chip), on a single thread. The PRF evaluation keys in our experiments have sizes roughly 40% and 60% of a bootstrapping key. Applying our solution to transciphering enables important bandwidth savings, typically trading 64-bit values for 4-bit values per transmitted ciphertext.

Keywords: Fully homomorphic encryption (FHE), Pseudorandom functions (PRFs), Learning-with-rounding (LWR), Oblivious randomness generation

1 Introduction

Fully homomorphic encryption (FHE) provides a method of outsourcing computation on sensitive data [30]. In particular, a client may encrypt data (e.g., patient medical records) using an FHE scheme and outsource some computation/evaluation (e.g., some diagnosis) on that data. Importantly, the server never learns the initial data, or the result of its computation.

A recurring problem with existing FHE schemes [15, 17, 19, 26, 32] is their ciphertext expansion: a ciphertext is at least an order of magnitude larger than its corresponding plaintext. This poses a significant problem when storing a large number of FHE ciphertexts in a remote database. A solution inspired by the performance of symmetric key ciphers is called *transciphering*. The idea is to first store symmetric key ciphertexts in the remote database long term. Then,

whenever computation on some database element(s) is requested, the server homomorphically evaluates the secret-key deciphering algorithm to obtain FHE ciphertext(s) encrypting the same database plaintext(s). To preserve privacy, the server is given an encrypted version of the client’s symmetric key. In summary, transciphering essentially allows a server to use an encrypted version of the client’s symmetric key to transform a symmetric key ciphertext into an FHE one. As pseudorandom functions (PRFs) are a key building block of secret-key cryptography, the problem of transciphering boils down to homomorphically evaluating a PRF as efficiently as possible.

Another application of homomorphically evaluating PRFs is producing encrypted randomness for blockchain smart contracts. An example of where one might require this is for simulating dice rolls or more generally playing games with randomness on the blockchain. In slightly more detail, a privacy-preserving blockchain may consist of a series of FHE ciphertexts encrypted under a key shared amongst an assigned group of validators. The task of these validators is to decrypt *particular* ciphertexts in a distributed manner. An example of an implementation of such a blockchain is fhEVM [23]. During a game/smart contract execution, clients can produce encryptions of random values by homomorphically evaluating a PRF. In doing so, the plaintexts remain hidden from the client but can still be considered random by relying on the security of the underlying PRF. These ciphertexts can then be used as dice rolls or even fed into an encrypted shuffling algorithm in a card game. Shuffling based on FHE-encrypted PRF outputs is also useful [27] in the context of private information retrieval [20].

When it comes to homomorphically computing secret-key pseudorandom objects, one may end up evaluating a complex circuit, which may be time-consuming and lead to impractical parameters without resorting to bootstrapping. Indeed, all existing FHE schemes involve ciphertexts containing a noise that grows during homomorphic evaluations. At some point, the noise grows too large to enable correct decryption, so that FHE schemes specify a bootstrapping algorithm which resets the noise to some predefined size. FHE schemes such as FHEW/TFHE [19, 25] take bootstrapping one step further. In particular, in their basic version, they enable the application of a *negacyclic* univariate function to the plaintext during the bootstrapping operation. This is referred to as *programmable* bootstrapping (PBS). Essentially, for any *negacyclic* univariate function f , applying the programmable bootstrapping algorithm takes an encryption of m and outputs an encryption of $f(m)$ with a predefined noise level. Another important point is that bootstrapping in FHEW/TFHE is very efficient in terms of latency (i.e., takes milliseconds) compared to BFV/BGV [15, 17, 26] where bootstrapping takes multiple seconds.

1.1 Our Contributions and Techniques

We consider the question of how efficiently we can homomorphically evaluate pseudorandom functions based on lattice assumptions. More precisely, we address the problem of evaluating PRFs based on the difficulty of the Learning-With-Rounding (LWR) problem [9], which can be seen as a variant of the Learning-With-Errors (LWE) problem where the noise is deterministically generated. For a public matrix $\mathbf{A} \in \mathbb{Z}_Q^{n \times m}$ with moduli p and Q such that $p < Q$ and a secret vector $\mathbf{s} \in \mathbb{Z}^n$, the LWR problem is to distinguish $\lceil (p/Q) \cdot (\mathbf{A}^\top \cdot \mathbf{s} \bmod Q) \rceil$ from a uniformly random vector in \mathbb{Z}_p^m . Here, the notation $\lceil \cdot \rceil$ denotes rounding to the nearest integer (rounding upwards in the case of a tie). The conjectured hardness of LWR naturally leads to a pseudorandom generator [9, Section 1.1] which expands a seed \mathbf{s} into a

longer pseudorandom string $\lceil (p/Q) \cdot (\mathbf{A}^\top \cdot \mathbf{s} \bmod Q) \rceil$ using a public matrix \mathbf{A} . By the GGM construction [34], it also implies a pseudorandom function family. Furthermore, it yields a more direct PRF construction (explicitly described [14] but already implicit in [9]) in the random oracle model. Given input x and a secret key $\mathbf{s} \in \mathbb{Z}^n$, the PRF evaluation is defined to be $\lceil (p/Q) \cdot (\mathbf{A}(x)^\top \cdot \mathbf{s} \bmod Q) \rceil$, where the matrix $\mathbf{A}(x) = H(x) \in \mathbb{Z}_Q^{n \times m}$ is derived from a random oracle H .

In this paper, we show that the latter construction can be evaluated efficiently using a small number of sequential PBSes if we restrict the ratio Q/p to be small. We note that the known reductions from LWE to LWR either assume that Q/p is super-polynomial [9] or that the number m of samples given to the distinguisher is a priori bounded [5, 11]. In fact, a certain class of reductions for $Q/p = \text{poly}(\lambda)$ and an unbounded number of samples was shown to be impossible [50]. However, even for an unbounded number of samples, LWR still appears to be exponentially hard (as commented in [9, 12]) in the parameter regime $Q/p = \Omega(\sqrt{n})$ when $p \mid Q$.

In the following, we show that, for a polynomial ratio Q/p , we can evaluate the above LWR-based PRF using a small number of sequential bootstraps. The idea is to compute an input-dependent vector $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ and view $\mathbf{c} = (-\mathbf{a}, 0) \in \mathbb{Z}_Q^{n+1}$ as an LWE ciphertext (with plaintext modulus p) that has a very large noise, but still decrypts to $\lceil \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \rceil$ under the LWE secret key $\mathbf{s} \in \mathbb{Z}^n$. So, if we have an FHE bootstrapping key encrypting the PRF secret key \mathbf{s} , we can obtain a low-noise encryption of the same value $\lceil \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \rceil$ by bootstrapping $\mathbf{c} = (-\mathbf{a}, 0) \in \mathbb{Z}_Q^{n+1}$.

With the bootstrapping techniques of FHEW [25] and TFHE [19], one difficulty is that we must find the appropriate *negacyclic* functions to evaluate in order to perform this operation. As a warm-up, we use techniques from Liu *et al.* [45] to evaluate the original PRF using three sequential PBSes (Lemma C.1) for each slot of $\log p$ pseudorandom bits. This evaluation procedure is detailed in Appendix C since it is not our most efficient solution. To improve on this first attempt, we use a recent “Full-domain functional bootstrapping” method from Ma *et al.* [47] so as to reduce the PBS depth (i.e., the number of sequential bootstraps per pseudorandom plaintext slot) to 2.

Finally, we suggest a modified version of the random-oracle-based PRF of [9, 14] which supports homomorphic evaluations in depth one and dispenses with the need to sequentially evaluate different negacyclic functions. By “depth-one”, we mean that, if the output space of the PRF is \mathbb{Z}_p^ℓ , each slot of $\log p$ output bits only costs one PBS to evaluate and ℓ slots can be processed in parallel in order to obtain a long output in \mathbb{Z}_p^ℓ . Our construction essentially applies a PBS using a single negacyclic version of the usual rounding function. As a result, the “ciphertext” $(-\mathbf{a}, 0)$ from above gets mapped to an encryption of

$$(-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2Q)} \cdot \left\lceil \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \right\rceil. \quad (1)$$

When p and Q are powers-of-2, we can show that the above function is a PRF based on the pseudorandomness of LWR. The reduction works by deriving the base-2 digits and the sign of the modified PRF from LWR samples whose moduli are scaled up by appropriate powers-of-2. An insignificant technicality of the reduction is that it is convenient to assume LWR holds with respect to the floor function, rather than the rounding function. As mentioned by Banerjee *et al.* [9, Sect. 2], there is no reason to suspect that the use of $\lceil \cdot \rceil$ as a rounding function impacts the hardness of LWR, even for $Q/p = \text{poly}(\lambda)$. We prove the pseudorandomness of the modified function (1) in Lemma 5.5 and its floor function analogue in Lemma 5.4. Note

that the latter achieves slightly better parameters. Although the base 2 is appropriate for practical implementations of TFHE, one could replace it by other bases if required.

A point worth mentioning is that due to the structure of TFHE, we typically end up using LWR with power-of-two moduli $Q \in \{2^9, 2^{10}, 2^{11}\}$ and $p \in \{2, \dots, 2^5\}$ in our depth-1 construction. In particular, p is often very small which allows LWR to remain hard for relatively small LWR dimension n_{LWR} compared to the TFHE LWE dimension n_{LWE} . This can be leveraged by truncating PBS operation; i.e., by performing n_{LWR} blind rotation steps during the PBS rather than the usual n_{LWE} steps. Since blind rotation dominates PBS latency, we may achieve a factor $n_{\text{LWR}}/n_{\text{LWE}}$ improvement in terms of latency.

1.2 Application to Transciphering

As an application of our homomorphic PRF evaluation scheme, we present a new efficient transciphering method whose security relies solely on a well-studied lattice problem. In particular, although we are using a tweaked LWR-based PRF for our most efficient instantiation, there is a tight mathematical reduction to a standard and well-known LWR assumption. This is in contrast to most previous works that tweak cipher/PRF design to suit FHE evaluation without provable reductions from well-studied problems. Furthermore, our most efficient scheme is very simple to implement using existing TFHE libraries as it simply takes advantage of the native PBS operation (more specifically, the blind rotation). The PRF evaluation key takes the form of blind rotation key material and essentially consists of ring GSW encryptions [32] of the PRF key bits. A positive characteristic of our construction is that the noise of output ciphertexts is as small as in bootstrapped ciphertexts. This is not true for schemes that rely on e.g. leveled BFV/BGV operations, where extra time-consuming bootstrapping may be necessary to use the outputs in certain applications. We finally mention the flexibility of our transciphering scheme. Should we be required to use an alternative plaintext modulus for some application, we may simply adjust the dimension of the LWR PRF secret to retain security. In particular, we simply need to find the smallest dimension such that LWR with moduli $Q = 2N$ and $2p$ remains sufficiently hard (for our depth-1 construction). Here, N is a parameter dictating the degree of a cyclotomic polynomial. For correctness of the PBS, $Q/2p$ will always be sufficiently large and we may find this LWR dimension. For example, the TFHE-rs [53] parameter sets typically satisfy $Q/2p \in \{32, 64\}$. Then, it is trivial to apply our transciphering scheme with a range of plaintext spaces by making small adjustments to the number of PRF key bits (and thus, adjusting the number of blind rotation steps). This means that the plaintext space can be dictated by the wider application, rather than the transciphering scheme. Further, our scheme has good granularity in the sense that each individual plaintext may be processed independently, rather than in blocks of a fixed size.

Implementation and Comparisons. To highlight the practicality of our scheme, we use the TFHE-rs [53] library. In particular, we run experiments on an AWS hpc7a.96xlarge instance with AVX512 using a single thread. Using the TFHE-rs parameter set PARAM_MESSAGE_2_CARRY_2, we obtain 5 bits with 5.2 ms latency leading to a throughput of 970 encrypted pseudorandom bits per second (bits/s). For the PARAM_MESSAGE_1_CARRY_1 parameters, we obtain 3 bits in 2.8 ms leading to a throughput of 1070 bits/s. The corresponding latencies on a Macbook Pro with an Apple M2 processor and 8 GB RAM are just 10% and 30% higher for the two parameter sets respectively. Our experiments show that the PRF evaluation key size is roughly 60% of a bootstrapping key for both of these parameter sets. Note that these results are single thread and would benefit from parallelization in the obvious way.

It is interesting to compare our results with state-of-the-art TFHE alternatives. Among these, the Trivium (which was only designed to provide 80 bits of security) and Kreyvium ciphers achieve 64 bits at a latency of over 100 ms (or around 500 bits/s of throughput) using more than 64 parallel threads [8, Table 2] on an AWS m6i.metal machine. This is ignoring the expensive warm-up phase for these ciphers. Our results show that our construction is faster, even when using a single thread on a standard laptop.

In an even more recent single-thread homomorphic evaluation of Trivium on a laptop with an Intel i5-1245U CPU and 16GB RAM, Bon *et al.* [13] obtained a throughput of 35 bits per second, which is ≈ 25 times lower than our throughput of 881 bits/s (also on a laptop using a single thread). A caveat of their approach is that the output of the evaluation algorithm is a ciphertext with noise larger than a bootstrapped ciphertext. This is not the case in our construction. The authors mention that this inconvenience is handled by performing a PBS in the transciphering or application phase. Shifting this PBS back into the evaluation phase for a fairer comparison suggests a throughput of 26 bits per second which is approximately 34 times lower than our reported throughput. Our increased throughput is nevertheless obtained at the expense of encrypting the PRF secret key in a larger ciphertext (with size $\approx 37\%$ and 58% of a regular TFHE bootstrapping key for the parameter sets used in our experiments). Fortunately, this is acceptable if the encrypted PRF secret key is only sent once and amortized over many transciphering operations.

1.3 Related Work

The idea of using bootstrapping to obviously generate FHE encryptions of random bits was previously used in the past (see, e.g., [1]). In this paper, we consider a derandomized version of the process where we prove that obviously generated ciphertexts indeed encrypt pseudorandom messages that are uniquely determined by an encrypted seed and a public input. For this purpose, we also aim at relying on the pseudorandomness of a well-studied PRF family.

Homomorphic evaluation of PRFs/ciphers using FHE has received a lot of attention in the research literature. With transciphering in mind, a natural task was to optimize the evaluation of AES [31]. Unfortunately, the efficiency of this approach is questionable and the current record stands at around 30 seconds for the evaluation of a single AES block [52] on a powerful server. This has led to the development of FHE-friendly cipher/PRF constructions such as LowMC [4], PASTA [24], FASTA [21], FLIP [49], FiLIP [48], Elisabeth [22], Rubato [36] and Chaghri [6]. Unfortunately, the security level of such schemes is not well understood and attacks are still being discovered [33, 35, 44]. In an attempt to avoid this problem, the standardized cipher Trivium [38] and the subsequent cipher Kreyvium [18] have been investigated as good options for efficient transciphering [8]. The homomorphic evaluation of Trivium and other symmetric primitives (including SIMON, AES and Keccak) was also considered via a framework [13] allowing to evaluate more complex Boolean functions. However, even Trivium and Kreyvium have recently been subjected to improved attacks [37]. For the sake of not putting all one's eggs in the same basket, it is desirable to have alternative solutions based on more stable and number theoretic assumptions, in particular if they enable higher throughputs than the homomorphic evaluation of stream ciphers. This motivates us to consider LWR-based PRFs and exploit the fact that their structure blends quite well with the bootstrapping paradigm of FHEW/TFHE.

In a direction somewhat analogous to ours, a recent work shows how to homomorphically

evaluate an adaptation of the LWR PRF using the BGV/BFV scheme [27]. In particular, this work tweaks the LWR-based PRF in a way that replaces the exact rounding function with an alternative based on the Legendre symbol. The advantage of this is that the resulting PRF is realizable using a reasonably small number of leveled homomorphic multiplications (concretely ≈ 20). It should be noted that these leveled multiplications lead to an output ciphertext with a noise larger than that of a freshly bootstrapped ciphertext which may need to be considered in certain applications. On the downside, their modified rounding function makes the resulting PRF significantly deviate from the well-known construction and introduces a novel variant of the LWR assumption. As of today, this assumption does not appear to be implied by the original assumption and has not undergone much cryptanalytic effort. In contrast, we rely on an LWR assumption that has been standing for over a decade.

Our approach has a common feature with the HERMES transciphering scheme [7] in that the latter also allows switching from a lattice-based symmetric cipher (with the difference that they use an LWE-based one while we rely on LWR) to CKKS/BGV/BFV (instead of TFHE in our case) without relying on any ad-hoc assumption. Yet, their 1.58 expansion factor is larger than ours. Moreover, they achieve a latency of around 25 seconds before any output is computed with around 60 bits per second of amortized throughput.

Brakerski *et al.* [16] took a different (non-transciphering-based) approach allowing to reduce the expansion rate of FHE schemes by constructing a ciphertext compression mechanism leading to rate- $(1 - o(1))$ FHE. Their technique applies to packed LWE ciphertexts sharing a common header (typically, a vector over \mathbb{Z}_q^n) where ℓ message-carrying slots encrypt ℓ distinct binary plaintexts under *distinct* LWE secret keys. It shrinks each of these ℓ slots down to a single bit, thus replacing a vector in \mathbb{Z}_q^ℓ by a binary string $\{0, 1\}^\ell$. For $\ell = \tilde{\Omega}(\lambda^2) = \text{poly}(\lambda)$ plaintexts, the achieved rate (i.e. plaintext size divided by ciphertext size) is $1 - O(1/\lambda)$ when expanding the ciphertext header from a seed. In other words, the size of a compressed output ciphertext on ℓ plaintext inputs approaches the size of around ℓ plaintexts as ℓ grows to infinity. However, subsequent FHE evaluations cannot be done on compressed ciphertexts and require bootstrapping to “undo” the compression. Applying this method to plain TFHE thus requires $\ell = \text{poly}(\lambda)$ bootstrapping keys (one for each of the original ℓ distinct keys). This overhead may be avoided by introducing a ring structure to TFHE as in [40] with a larger than usual ring dimension ℓ to achieve reasonable expansion factor. However, the sub-optimal granularity (i.e., the fact that a block of $\ell = \text{poly}(\lambda)$ ciphertexts is required before a compression operation can begin) remains. Note that one may also apply the compression to BGV/BFV.

Outline of the paper. The rest of this paper is organized as follows. We begin with background knowledge and definitions in Section 2. Next, in Section 3, we discuss homomorphic evaluation of the standard LWR function in depth 2. Then, in Section 4 we discuss ring-LWR-based PRFs in the context of BFV. In Section 5, we present the depth-1 evaluation of our modified LWR PRF along with the associated security reductions. This is followed by an implementation of our depth-1 construction using the TFHE-rs library in Section 6. Finally, for completeness, we concretize the transciphering application in Section 7. In appendix, we formally define pseudorandom functions and describe GGSW ciphertexts. We also include a less efficient depth-3 evaluation of the standard LWR PRF.

2 Background and Definitions

Notation. In the following, when D is a distribution, $x \sim D$ means that x is a random variable distributed according to D . The notation $x \leftarrow D$ denotes the explicit action of sampling an element x according to the distribution D . For a finite set \mathcal{S} , $U(\mathcal{S})$ stands for the uniform distribution over \mathcal{S} . For any integer $q \geq 2$, \mathbb{Z}_q denotes the ring of integers with addition and multiplication modulo q . For any real number y , we use $\lceil y \rceil$ to denote rounding y to the nearest integer (rounding upwards in the case of a tie) and $\lfloor y \rfloor$ to denote the floor function. If y is replaced by a vector \mathbf{y} , we apply the rounding and floor functions entry-wise.

2.1 Cryptographic Assumptions

LWE/LWR Assumptions. We first recall the *Learning-With-Errors* (LWE) assumption defined by Regev [51].

Definition 2.1 (LWE assumption). Let integers $m \geq n \geq 1$, $q \geq 2$ and let χ_s, χ_e be distributions over \mathbb{Z} . The $\text{LWE}_{n,m,q,\chi_s,\chi_e}$ problem consists in distinguishing between the distributions

$$\{(\mathbf{A}^\top, \mathbf{A}^\top \mathbf{s} + \mathbf{e}) \mid \mathbf{A} \sim U(\mathbb{Z}_q^{n \times m}), \mathbf{s} \leftarrow \chi_s^n, \mathbf{e} \sim \chi_e^m\}$$

and $U(\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m)$.

When the distribution of \mathbf{s} is the uniform distribution $U(\mathbb{Z}_q^n)$, the assumption is sometimes denoted by $\text{LWE}_{n,m,q,\chi_e}$.

We now recall the definition of the *Learning-With-Rounding* (LWR) problem [9].

Definition 2.2 (LWR assumption). Let integers $m \geq n \geq 1$, $q > p \geq 2$ and let χ_s be a distribution over \mathbb{Z} . The *Learning-With-Rounding* ($\text{LWR}_{n,m,q,p,\chi_s}$) problem consists in distinguishing between the distributions

$$\{(\mathbf{A}^\top, \lceil (p/q) \cdot (\mathbf{A}^\top \mathbf{s} \bmod q) \rceil \bmod p) \mid \mathbf{A} \sim U(\mathbb{Z}_q^{n \times m}), \mathbf{s} \leftarrow \chi_s^n\}$$

and $U(\mathbb{Z}_q^{m \times n} \times \mathbb{Z}_p^m)$.

When the number m of samples is a priori bounded, the LWE assumption is known [5, 11] to imply the hardness of LWR for a polynomial ratio $q/p = \text{poly}(\lambda)$. When there is no predetermined upper bound on the number of samples, the only known reduction [9, Theorem 3.2] from LWE to LWR requires a super-polynomial ratio $q/p = \Omega(\lambda^{\omega(1)})$. However, it is quite plausible that LWR remains hard for $q/p = \text{poly}(\lambda)$ even for an a priori unbounded number of samples. As discussed in [9, 12], as long as $q/p = \Omega(\sqrt{n})$ and q/p is an integer (so that $\lceil (p/q) \cdot U(\mathbb{Z}_q) \rceil = U(\mathbb{Z}_p)$), LWR may be exponentially hard even for quantum algorithms. We also note that replacing the rounding function $\lceil \cdot \rceil$ by the floor function is not believed to affect the hardness of the LWR problem [9, Sect. 2].

Ring-LWE/LWR Assumptions. We now recall the definition of the ring Learning-With-Errors problem [46].

Definition 2.3 (RLWE assumption). Take an integer $q \geq 2$. Let $\Phi(X)$ be a cyclotomic polynomial of degree N and let the rings $\mathcal{R} = \mathbb{Z}[X]/(\Phi(X))$ and $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$. Let χ_s, χ_e be distributions over \mathcal{R} . The *Ring LWE* ($\text{RLWE}_{N,m,q,\chi_s,\chi_e}$) problem consists in distinguishing between the distributions

$$\{(\mathbf{a}, \mathbf{a} \cdot \mathbf{s} + \mathbf{e}) \mid \mathbf{a} \sim U(\mathcal{R}_q^m), \mathbf{s} \leftarrow \chi_s, \mathbf{e} \sim \chi_e^m\} \text{ and } U(\mathcal{R}_q^m \times \mathcal{R}_q^m) .$$

The LWR problem has a natural analogue in the ring setting. The *Ring Learning-With-Rounding* (RLWR) problem [9] is defined as follows.

Definition 2.4 (RLWR assumption). Let integers $q > p \geq 2$. Let $\Phi(X)$ be a cyclotomic polynomial of degree N and let the rings $\mathcal{R} = \mathbb{Z}[X]/(\Phi(X))$ and $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$. Let χ_s be a distribution over \mathcal{R} . The *Ring Learning-With-Rounding* (RLWR $_{N,m,q,p,\chi_s}$) problem consists in distinguishing between the distributions

$$\{(\mathbf{a}, \lceil (p/q) \cdot (\mathbf{a} \cdot s \bmod q) \rceil \bmod p) \mid \mathbf{a} \sim U(\mathcal{R}_q^m), s \leftarrow \chi_s\}$$

and $U(\mathcal{R}_q^m \times \mathcal{R}_p^m)$.

Finally, we recall the definition of the *Generalized Learning-With-Errors* (GLWE) problem (also known as the *Module-Learning-With-Errors* problem) studied in [42] that is useful when discussing the TFHE/FHEW [19, 25] FHE schemes.

Definition 2.5 (GLWE assumption). Take an integer $q \geq 2$ and a rank $k \geq 1$. Let $\Phi(X)$ be a cyclotomic polynomial of degree N and let the rings $\mathcal{R} = \mathbb{Z}[X]/(\Phi(X))$ and $\mathcal{R}_q = \mathcal{R}/q\mathcal{R}$. Let χ_s, χ_e be distributions over \mathcal{R} . The *Generalized LWE* (GLWE $_{k,N,m,q,\chi_s,\chi_e}$) problem consists in distinguishing between the distributions

$$\{(\mathbf{A}^\top, \mathbf{A}^\top \cdot \mathbf{s} + \mathbf{e}) \mid \mathbf{A} \sim U(\mathcal{R}_q^{k \times m}), \mathbf{s} \leftarrow \chi_s^k, \mathbf{e} \sim \chi_e^m\}$$

and $U(\mathcal{R}_q^{m \times k} \times \mathcal{R}_q^m)$.

2.2 (Key-Homomorphic) Pseudorandom Functions Based on LWR

In [9] (see also [12]), Banerjee, Peikert and Rosen implicitly describe a weak pseudorandom function based on the hardness of the LWR problem. This weak PRF maps a uniformly random input $\mathbf{a} \in \mathbb{Z}_Q^n$ to the output $\text{wPRF}_s(\mathbf{a}) = \lceil (p/Q) \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \rceil \bmod p$, where $\mathbf{s} \in \mathbb{Z}^n$ is the secret key.

By introducing a random oracle $H: \{0, 1\}^\ell \rightarrow \mathbb{Z}_Q^n$, this weak PRF can be turned into a full PRF by computing $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ when the PRF has to be evaluated on an arbitrary input $x \in \{0, 1\}^\ell$. As pointed out in [14], this PRF turns out to be almost key-homomorphic PRF and was recently used in [29]. It is precisely defined as follows.

The secret key is a vector $\mathbf{s} = (s_1, \dots, s_n) \sim \chi_s^n$ where each s_i is sampled from a distribution χ_s specified by public parameters. These public parameters also contain the description of a hash function $H: \{0, 1\}^\ell \rightarrow \mathbb{Z}_Q^n$ (modeled as a random oracle) and two moduli p and Q where p divides Q . Typically, χ_s is the uniform distribution over \mathbb{Z}_Q . Alternatively, χ_s can be a discrete Gaussian distribution with a suitable standard deviation σ or even the uniform binary distribution. We note that in any of these cases, the parameters (n, q, p) can be chosen carefully to protect against all known attacks.

A function evaluation is then defined as

$$x \mapsto \text{PRF}_s(x) \triangleq \left\lceil \frac{p}{Q} \cdot (\langle H(x), \mathbf{s} \rangle \bmod Q) \right\rceil \bmod p \quad (2)$$

and outputs a scalar in \mathbb{Z}_p . If we need to output t pseudorandom elements in \mathbb{Z}_p , we can extend it as

$$x \mapsto \text{PRF}_s(x) \triangleq (y_1, \dots, y_t)$$

where

$$y_i = \left\lceil \frac{p}{Q} \cdot (\langle H(x, i), \mathbf{s} \rangle \bmod Q) \right\rceil \bmod p \quad \forall i \in \{1, \dots, t\} .$$

When proving the pseudorandomness of the function (2) (in the random oracle model), the reduction (see [9] or [29, Theorem 3.1]) is given from an instance of LWR where the number of samples is *not* a priori bounded since each queried input x is mapped to a different sample (i.e., the number of samples is as large as the number of evaluation queries).

Therefore we need to choose a super-polynomial $Q/p = \lambda^{\omega(1)}$ if we want to rely on known LWE-to-LWR reductions [9]. Alternatively, one may prefer a more efficient choice of parameters with $Q/p = \text{poly}(\lambda)$ and rely on the plausible hardness of LWR in this parameter regime.¹ In this case, it is recommended in [9] to set Q/p as an integer larger than $\Omega(\sqrt{n})$ if n is the dimension of \mathbf{s} .

In [9, 12], LWR was conjectured to be exponentially hard when $Q/p = \Omega(\sqrt{n})$ and assuming uniform secret keys (i.e., $\chi_s = U(\mathbb{Z}_Q)$). In order to homomorphically evaluate the PRF using programmable bootstrapping, it is more convenient to sample the seed \mathbf{s} from a binary or ternary distribution. In practice, we use the lattice estimator [3]² to derive secure parameters. Although the lattice estimator is designed for LWE, we deploy the usual heuristic method of approximating the hardness of LWR by that of LWE with uniform noise in the interval $[-\frac{Q}{2p} + 1, \frac{Q}{2p}]$. Fortunately, even for the uniform binary distribution $\chi_s = U(\{0, 1\})$, we can find 128-bit secure parameters for the range of moduli Q and p that our constructions require.

2.3 TFHE Bootstrapping

An LWE ciphertext encrypting a message $m \in \mathbb{Z}_p$ with respect to secret key takes the form $(\mathbf{a}, b) = (\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle + e + m \lceil q/p \rceil) \in \mathbb{Z}_q^{n+1}$. Here, \mathbf{a} is sampled uniformly and e is sampled from an error distribution χ_e . As in [45], when (\mathbf{a}, b) is an LWE ciphertext with secret key \mathbf{s} , we denote by $\text{Dec}_s(\mathbf{a}, b) = b - \langle \mathbf{a}, \mathbf{s} \rangle \bmod q$ the decoding function (a.k.a. phase function) which outputs a noisy encoding $e + m \lceil q/p \rceil$ of the plaintext m . We can similarly define GLWE ciphertexts by taking $\mathbf{a} \in \mathcal{R}_q^k$, $\mathbf{s} \in \mathcal{R}_q^k$, error $e \in \mathcal{R}$ and message $m \in \mathcal{R}_p$.

Programmable Bootstrapping Subroutines. We rely on the following theorem, which is quoted from [45] but is implied by earlier results on the programmable bootstrapping of LWE ciphertexts for negacyclic functions. Note that in this theorem and throughout this paper, q is used to denote the TFHE modulus and Q will denote a smaller modulus dividing q . We will assume that $Q = 2N$ where N is the degree of the TFHE cyclotomic ring.

THEOREM 2.6 ([45, THEOREM 1]). *Take positive integers n , q and Q such that Q divides q and q is set to a power of 2. There is a bootstrapping procedure Boot with the following property: For any LWE ciphertext $(\mathbf{a}, b) \in \mathbb{Z}_Q^{n+1}$ and any function $f: \mathbb{Z}_Q \rightarrow \mathbb{Z}_q$ such that $f(x + Q/2) = -f(x) \bmod q$, the procedure $\text{Boot}[f](\mathbf{a}, b)$ outputs a ciphertext $(\mathbf{c}, d) \in \mathbb{Z}_q^{n+1}$ such that*

$$\text{Dec}_s(\mathbf{c}, d) = f(\text{Dec}_s(\mathbf{a}, b)) + e \pmod{q},$$

where $|e| < \beta$, for a noise bound β that only depends on the operations performed by Boot and not on the input ciphertext (\mathbf{a}, b) .

There are three subroutines in TFHE bootstrapping: blind rotation, sample extraction and key-switching. In what follows, we set $Q = 2N$ in the above theorem. The blind rotation

¹We note that, in the statement of [29, Theorem 3.1], the hypothesis $Q/p = \Omega(n^{\omega(1)})$ is only needed if a reduction from LWE is desired via the LWE-to-LWR reduction of [9]. The proof still works under the LWR assumption when Q/p is polynomial.

²<https://github.com/malb/lattice-estimator>

in TFHE takes as input an LWE ciphertext $(\mathbf{a}, b) \in \mathbb{Z}_{2N}^{n+1}$ under secret key $\mathbf{s} \in \{0, 1\}^n$ that “encrypts” a message $\tilde{\mu} \in \mathbb{Z}_{2N}$ and a test polynomial $v(X)$ whose coefficients encode the outputs of a negacyclic function f in a lookup table. It returns a GLWE ciphertext $\mathbf{c}' \in (\mathbb{Z}_q[X]/(X^N + 1))^{k+1}$, under secret key with bounded coefficients $\mathbf{s}' \in R^k$, which encrypts the polynomial $X^{-b+(\mathbf{a}, \mathbf{s}) \bmod 2N} \cdot v(X)$ whose degree-0 coefficient is $f(\text{Dec}_{\mathbf{s}}(\mathbf{a}, b))$ when f is negacyclic. This blind rotation operation requires a bootstrapping key in the form of generalized GSW (or GGSW) encryptions [25, 32] of the entries of \mathbf{s} . For completeness, we overview GGSW in Appendix B. The resulting GLWE ciphertext is then sample-extracted to obtain an LWE ciphertext encrypting the degree-0 coefficient of $X^{-b+(\mathbf{a}, \mathbf{s}) \bmod 2N} \cdot v(X)$. The resulting LWE ciphertext $(\mathbf{c}', d') \in \mathbb{Z}_q^{kN+1}$ is encrypted under the secret key \mathbf{s}' . A final key switch leads to an LWE ciphertext $(\mathbf{c}, d) \in \mathbb{Z}_q^{n+1}$ under the original secret key \mathbf{s} . To sum up, we have:

$$\begin{aligned}
 (\mathbf{c}, d) &\leftarrow \text{KeySwitch} \circ \text{SampleExtract} \circ \text{BlindRotate}(\mathbf{a}, b) \\
 &\qquad\qquad\qquad \underbrace{\qquad\qquad\qquad}_{=\text{GLWE}_{\mathbf{s}'}(X^{-\text{Dec}_{\mathbf{s}}(\mathbf{a}, b)} \cdot v(X))} \\
 &\qquad\qquad\qquad \underbrace{\qquad\qquad\qquad}_{=\text{LWE}_{\mathbf{s}'}(f(\text{Dec}_{\mathbf{s}}(\mathbf{a}, b)))} \\
 &\qquad\qquad\qquad \underbrace{\qquad\qquad\qquad}_{=\text{LWE}_{\mathbf{s}}(f(\text{Dec}_{\mathbf{s}}(\mathbf{a}, b)))}
 \end{aligned}$$

provided that test polynomial $v(X) = \sum_{i=0}^{N-1} v_i X^i$ is programmed as $v_i = f(i) \in \mathbb{Z}_q$ for some negacyclic function $f: \mathbb{Z}_{2N} \rightarrow \mathbb{Z}_q$. For a more detailed exposition, see [39]. Note that in order to get the output $(\mathbf{c}, d) \in \mathbb{Z}_q^{n+1}$ back into the domain of Boot, one can simply apply the mod-switching operation given by $\text{ModSwitch}_{q \rightarrow 2N}(\mathbf{c}, d) \triangleq (\lceil (2N/q) \cdot \mathbf{c} \rceil, \lceil (2N/q) \cdot d \rceil)$. This is useful when a sequence of multiple PBSes for negacyclic functions with different domains and ranges is required.

3 Homomorphic Evaluation of LWR-based PRF in Depth 2

We now describe the homomorphic evaluation of the standard LWR-based PRF in depth 2. Note that this construction is not as efficient as the depth-1 construction in Section 5 and may be skipped by the reader. Nonetheless, this section describes what can be achieved using known techniques and paves the way for the homomorphic BFV evaluation of the RLWR-based PRF discussed in Section 4. In this section, we set $Q = 2N$ and $\Delta = Q/p$ where p is a plaintext modulus dividing $2N$. Furthermore, the TFHE modulus $q > 2N$ is also assumed to be divisible by $2N$ (which is the case in practice).

In order to homomorphically evaluate the PRF in (2) for a public input x given an encryption of the seed $\mathbf{s} \in \mathbb{Z}^n$, the idea is to first compute an input-dependent $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ and view $(-\mathbf{a}, 0)$ as an LWE ciphertext with a very large noise. Namely, assuming that $\Delta \mid Q$, if we write

$$\begin{aligned}
 (-\mathbf{a}, 0) &= (-\mathbf{a}, -(\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) + \Delta \cdot \lceil (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) / \Delta \rceil + \underbrace{(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)}_{\in [-\Delta/2, \Delta/2]},
 \end{aligned}$$

we can view the term $(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)$ as a noise to clean up using bootstrapping. Using TFHE/FHEW-like schemes, the main difficulty is to do this using negacyclic functions.

A first solution is to use a technique proposed by Liu *et al.* [45, Section 4] which applies Theorem 2.6 to negacyclic functions described in Appendix C. This method is based on the

homomorphic floor function evaluation technique [45, Algorithm 2] that allows handling an arbitrarily large noise in the input ciphertext. For each slot of pseudorandomness, the resulting homomorphic evaluation algorithm (which we describe in Appendix C for completeness) requires 3 sequential PBSes.

To obtain better efficiency, we can actually use a technique from [47, Algorithm 1] so as to only call Boot twice for each plaintext slot. To do this, we use the negacyclic functions $f_C, f_{\text{eval}}: \mathbb{Z}_Q \rightarrow \mathbb{Z}_Q$ defined as

$$f_C(x) = \begin{cases} \frac{\Delta}{4} \cdot \left(2 \left\lfloor \frac{x}{\Delta} \right\rfloor + 1 \right) \bmod Q & \text{if } x \in \left[0, \frac{Q}{2} - 1 \right] \\ -\frac{\Delta}{4} \cdot \left(2 \left\lfloor \frac{x}{\Delta} \right\rfloor - p + 1 \right) \bmod Q & \text{if } x \in \left[\frac{Q}{2}, Q - 1 \right] \end{cases}$$

$$f_{\text{eval}}(x) = \begin{cases} \Delta \cdot \left(\left\lfloor \frac{2x}{\Delta} \right\rfloor \bmod p \right) & \text{if } x \in \left[0, \frac{Q}{4} - 1 \right] \\ \Delta \cdot \left(\left\lfloor \frac{2(Q-x)}{\Delta} \right\rfloor + \frac{p}{2} \bmod p \right) & \text{if } x \in \left[\frac{3Q}{4}, Q - 1 \right] \\ -f_{\text{eval}}\left(x - \frac{Q}{2}\right) \bmod Q & \text{if } x \in \left[\frac{Q}{4}, \frac{3Q}{4} - 1 \right] \end{cases}$$

where $\Delta = Q/p$ and the input $x \in \mathbb{Z}_Q$ is seen as a positive integer in $\{0, \dots, Q - 1\}$. We note that f_C is negacyclic since, for each $x \in [Q/2, Q - 1]$,

$$\begin{aligned} f_C(x - Q/2) &= \frac{Q}{4p} \cdot \left(2 \left\lfloor \frac{p}{Q} \cdot \left(x - \frac{Q}{2}\right) \right\rfloor + 1 \right) \\ &= \frac{Q}{4p} \cdot \left(2 \left\lfloor \frac{p}{Q} \cdot x \right\rfloor - p + 1 \right) = -f_C(x) . \end{aligned}$$

The negacyclic property of f_{eval} can also be checked in a similar way (with additional cases to consider):

$$\begin{aligned} f_{\text{eval}}(x + Q/2) &= \begin{cases} -f_{\text{eval}}\left(\left(x + Q/2\right) - \frac{Q}{2}\right) \bmod Q & \text{if } x \in \left[0, \frac{Q}{4} - 1 \right] \\ \Delta \cdot \left(\left\lfloor \frac{2(Q - (x + Q/2))}{\Delta} \right\rfloor + \frac{p}{2} \bmod p \right) & \text{if } x \in \left[\frac{Q}{4}, \frac{Q}{2} - 1 \right] \\ \Delta \cdot \left(\left\lfloor \frac{2(x + Q/2)}{\Delta} \right\rfloor \bmod p \right) & \text{if } x \in \left[\frac{Q}{2}, \frac{3Q}{4} - 1 \right] \\ -f_{\text{eval}}\left(\left(x + Q/2\right) - \frac{Q}{2}\right) \bmod Q & \text{if } x \in \left[\frac{3Q}{4}, Q - 1 \right] \end{cases} \\ &= \begin{cases} -f_{\text{eval}}(x) \bmod Q & \text{if } x \in \left[0, \frac{Q}{4} - 1 \right] \\ \Delta \cdot \left(\left\lfloor \frac{2(Q-x)}{\Delta} \right\rfloor + \frac{p}{2} \bmod p \right) & \text{if } x \in \left[\frac{Q}{4}, \frac{Q}{2} - 1 \right] \\ \Delta \cdot \left(\left\lfloor \frac{2x}{\Delta} \right\rfloor \bmod p \right) & \text{if } x \in \left[\frac{Q}{2}, \frac{3Q}{4} - 1 \right] \\ -f_{\text{eval}}(x) \bmod Q & \text{if } x \in \left[\frac{3Q}{4}, Q - 1 \right] \end{cases} \\ &= -f_{\text{eval}}(x) \bmod Q . \end{aligned}$$

Again, we assume that Q and p are both powers of 2. In the description hereunder, we also assume that the seed $\mathbf{s} \in \mathbb{Z}^n$ of the LWR-based PRF is encrypted in the same way as the bootstrapping key of an TFHE encryption scheme where the LWE secret key is \mathbf{s} . We thus assume a bootstrapping key $\text{bsk} = \text{GGSW}_{\mathbf{s}'}(\mathbf{s})$ consisting of a GGSW encryption [25, 32] of the seed \mathbf{s} (which is viewed as the decryption key of an LWE-based secret-key encryption scheme) under a GGSW secret key \mathbf{s}' . The evaluation algorithm then goes as follows. Here, as in [47, Section 3], the inner product $\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q$ is interpreted as an unsigned element of $\{0, \dots, Q - 1\}$ (rather than $\{-Q/2, \dots, Q/2 - 1\}$) and μ is viewed as an element of $\{0, \dots, p - 1\}$.

$\text{Eval}_{\text{pp}}(\text{bsk}, x)$: Given public parameters $\text{pp} = (n, Q, p, \beta)$, an evaluation key bsk and an input $x \in \{0, 1\}^\ell$, compute the input-dependent vector $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ and do the following:

1. $(\mathbf{c}, d) := \text{Boot}[f_C](-\mathbf{a}, \frac{\Delta}{2}) \pmod{Q}$
 2. $(\bar{\mathbf{a}}, \bar{b}) = \text{Boot}[f_{\text{eval}}](\mathbf{c}, d) \pmod{Q}$
- Output the ciphertext $\text{ct} = (\bar{\mathbf{a}}, \bar{b}) \in \mathbb{Z}_Q^{n+1}$.

The above homomorphic evaluation algorithm outputs an LWE encryption of $\text{PRF}_s(x) = \lceil \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \rceil \bmod p$ (interpreted as an element of $[0, p-1]$) under an LWE secret key which is the PRF seed \mathbf{s} itself. In order to obtain an LWE encryption of $\text{PRF}_s(x)$ under the LWE secret key \mathbf{s}' , we can remove the key-switching step in the second call to `Boot`. We further note that the final modulus switch may be removed.

The following lemma is adapted from [47, Lemma 3.1], with all details written out.

LEMMA 3.1. *Assume that p and Q are both powers of 2 and that $\Delta = Q/p > 4\beta$, where β is the bootstrapping error from Theorem 2.6. For any $x \in \{0, 1\}^\ell$, `Eval` outputs a ciphertext $(\bar{\mathbf{a}}, \bar{b}) \in \mathbb{Z}_Q^{n+1}$ such that $\text{Dec}_s(\bar{\mathbf{a}}, \bar{b}) = \Delta \cdot \mu + e \pmod{Q}$, where $|e| < \beta$,*

$$\mu = \left\lceil \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \right\rceil \bmod p \quad (3)$$

with $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$.

PROOF. We first note that

$$(-\mathbf{a}, 0) \in \mathbb{Z}_Q^{n+1} = (-\mathbf{a}, -(\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) + \Delta \cdot \underbrace{\lceil (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) / \Delta \rceil + (\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)}_{\in [-\Delta/2, \Delta/2]})$$

where $\Delta = Q/p$ and $(\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \in \{0, \dots, Q-1\}$. Then, before Line 1, we have

$$\begin{aligned} \text{Dec}_s(-\mathbf{a}, 0) &= \langle \mathbf{a}, \mathbf{s} \rangle \bmod Q \\ &= \Delta \cdot \underbrace{\lceil (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) / \Delta \rceil}_{\triangleq \mu} + (\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta) \pmod{Q}, \end{aligned}$$

where $\mu \in \{0, \dots, p-1\}$,³ and we can interpret $(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)$ as a very large noise. Initially, we have

$$\text{Dec}_s(-\mathbf{a}, \frac{\Delta}{2}) = \Delta \cdot \mu + \underbrace{((\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta) + \frac{\Delta}{2})}_{\triangleq \bar{e}} \pmod{Q} \quad (4)$$

where $\bar{e} \in [0, \Delta)$. Then, we distinguish two cases.

Case I: $\Delta \cdot \mu + \bar{e} \in [0, Q/2 - 1]$

We have

$$\begin{aligned} f_C(\Delta \cdot \mu + \bar{e}) &= \frac{\Delta}{4} \cdot \left(2 \left\lfloor \frac{\Delta \cdot \mu + \bar{e}}{\Delta} \right\rfloor + 1 \right) \bmod Q \\ &= \frac{\Delta}{4} \cdot (2\mu + 1) \bmod Q. \end{aligned}$$

After Line 1, we obtain

$$\begin{aligned} \text{Dec}_s(\mathbf{c}, d) &\equiv f_C(\text{Dec}_s(-\bar{\mathbf{a}}, \frac{\Delta}{2})) + e_\beta \\ &\equiv f_C(\Delta \cdot \mu + \bar{e}) + e_\beta \\ &\equiv \frac{\Delta}{4} \cdot (2\mu + 1) + e_\beta \pmod{Q} \end{aligned}$$

³Recall that, in this section, elements of \mathbb{Z}_Q are viewed as unsigned integers with a representative in $[0, Q-1]$.

for some $e_\beta \in (-\beta, \beta)$. Moreover, we know that

$$f_C(\Delta \cdot \mu + \bar{e}) \in \left[\frac{\Delta}{4}, \frac{Q}{4} - \frac{\Delta}{4} \right]$$

because $f_C(x) \in [0, Q/4 - \Delta/4]$ for any $x \in [0, Q/2 - 1]$ and we cannot have $\frac{\Delta}{4} \cdot (2\mu + 1) \in [0, \Delta/4]$ for $\mu \in \{0, \dots, p-1\}$. Since $|e_\beta| < \beta < \Delta/4$, this implies

$$\text{Dec}_s(\mathbf{c}, d) \bmod Q = \frac{\Delta}{4} \cdot (2\mu + 1) + e_\beta \in [0, Q/4 - 1] .$$

By the definition of f_{eval} , this in turn yields

$$\begin{aligned} f_{\text{eval}}(\text{Dec}_s(\mathbf{c}, d) \bmod Q) &\equiv \Delta \cdot \left[\frac{2}{\Delta} \cdot \left(\frac{\Delta}{4} \cdot (2\mu + 1) + e_\beta \right) \right] \\ &\equiv \Delta \cdot \left[\left(\mu + \frac{1}{2} + \frac{2}{\Delta} \cdot e_\beta \right) \right] \\ &\equiv \Delta \cdot \mu \pmod{Q} \end{aligned}$$

since $\left| \frac{2}{\Delta} \cdot e_\beta \right| < (2/\Delta) \cdot (\Delta/4) = 1/2$. By Theorem 2.6, after Line 2, we obtain

$$\begin{aligned} \text{Dec}_s(\bar{\mathbf{a}}, \bar{\mathbf{b}}) &\equiv f_{\text{eval}}(\text{Dec}_s(\mathbf{c}, d)) + e'_\beta \\ &\equiv \Delta \cdot \mu + e'_\beta \end{aligned}$$

for some $e'_\beta \in (-\beta, \beta)$.

Case II: $\Delta \cdot \mu + \bar{e} \in [Q/2, Q - 1]$

We have

$$\begin{aligned} f_C(\Delta \cdot \mu + \bar{e}) &= -\frac{\Delta}{4} \cdot \left(2 \left\lfloor \frac{\Delta \cdot \mu + \bar{e}}{\Delta} \right\rfloor - p + 1 \right) \bmod Q \\ &= -\frac{\Delta}{4} \cdot (2\mu - p + 1) \bmod Q \\ &= -\frac{\Delta}{2} \cdot \mu + \frac{Q}{4} - \frac{\Delta}{4} \bmod Q \end{aligned}$$

so that, after Line 1,

$$\begin{aligned} \text{Dec}_s(\mathbf{c}, d) &\equiv f_C(\Delta \cdot \mu + \bar{e}) + e_\beta \\ &\equiv -\frac{\Delta}{2} \cdot \mu + \frac{Q}{4} - \frac{\Delta}{4} + e_\beta \pmod{Q} \end{aligned} \tag{5}$$

for some $e_\beta \in (-\beta, \beta)$.

Also, for any $x \in [Q/2, Q - 1]$, we have $f_C(x) \in [(3Q + \Delta)/4, Q - \frac{\Delta}{4}]$, so that $f_C(x) + e_\beta \in [3Q/4, Q - 1]$ whenever $e_\beta \in (-\Delta/4, \Delta/4)$ and the rightmost side of (5) thus lives in $[3Q/4, Q - 1]$.

Since $f_C(\Delta \cdot \mu + \bar{e}) = -\frac{\Delta}{2} \cdot \mu + \frac{Q}{4} - \frac{\Delta}{4} \pmod{Q}$, we have (over \mathbb{Q})

$$\begin{aligned} \frac{2}{\Delta} \cdot (Q - f_C(\Delta \cdot \mu + \bar{e}) - e_\beta) &= \frac{2}{\Delta} \cdot \left(\frac{\Delta}{2} \cdot \mu + \frac{3Q}{4} + \frac{\Delta}{4} - e_\beta + k \cdot Q \right) \\ &= \mu + \frac{3p}{2} + \frac{1}{2} - \underbrace{\frac{2}{\Delta} \cdot e_\beta}_{\in (-\frac{1}{2}, \frac{1}{2})} + 2k \cdot p \end{aligned}$$

for some integer $k \in \mathbb{Z}$. By rounding, it comes that

$$\left\lfloor \frac{2}{\Delta} \cdot (Q - f_C(\Delta \cdot \mu + \bar{e}) - e_\beta) \right\rfloor = \mu + \frac{3p}{2} + 2k \cdot p$$

(still over \mathbb{Q}) and

$$\begin{aligned} \left\lfloor \frac{2}{\Delta} \cdot (Q - f_C(\Delta \cdot \mu + \bar{e}) - e_\beta) \right\rfloor + \frac{p}{2} &\equiv \mu + 2(k+1) \cdot p \\ &\equiv \mu \pmod{p} . \end{aligned}$$

Given that $\text{Dec}_s(c, d) \bmod Q \in [3Q/4, Q - 1]$ after Line 1, we have

$$\begin{aligned} f_{\text{eval}}(\text{Dec}_s(c, d) \bmod Q) &= f_{\text{eval}}(f_C(\Delta \cdot \mu + \bar{e}) + e_\beta) \\ &= \Delta \cdot \left(\left\lfloor \frac{2}{\Delta} \cdot (Q - f_C(\Delta \cdot \mu + \bar{e}) - e_\beta) \right\rfloor + \frac{p}{2} \bmod p \right) \bmod Q \\ &= \Delta \cdot \mu \bmod Q . \end{aligned}$$

Therefore, after Line 2, we get

$$\begin{aligned} \text{Dec}_s(\bar{a}, \bar{b}) &\equiv f_{\text{eval}}(\text{Dec}_s(c, d)) + e'_\beta \\ &\equiv \Delta \cdot \mu + e'_\beta \end{aligned}$$

for some $e'_\beta \in (-\beta, \beta)$, as claimed. \square

We need to assume that $Q = 2N$, where N is the ring dimension (i.e., the degree of the cyclotomic polynomial $X^N + 1$ used in the GGSW scheme encrypting the PRF seed s), in order to evaluate the negacyclic functions using look-up tables of reasonable size. As an example, suppose that $Q = 2 \times 2048$, $p = 2^5$ and $n = 761$ with a binary secret key (which is the case for the PARAM_MESSAGE_2_CARRY_2 parameters from TFHE-rs). In this case, the lattice estimator suggests around 200 bits of security for the $\text{LWR}_{n,m,2N,p,U(\{0,1\})}$ for unbounded number of samples m , meaning that the LWR problem of interest is concretely hard.

Remark 3.2. We note that Theorem 2.6 applies to the bootstrapping algorithm of [43], which does not require secret keys to be small. This allows homomorphically evaluating the PRF described in (2) when its secret key is sampled from a wide discrete Gaussian distribution (rather than a uniform binary/ternary distribution). The methodology is the same as above.

4 Extension to RLWR-based PRFs Using BFV

The approach of Section 3 extends to homomorphically evaluate the ring analogue of the PRF recalled Section 2.2. We assume a random oracle $H: \{0, 1\}^\ell \rightarrow \mathcal{R}_q$ that ranges over the ring \mathcal{R}_q .

Recall that, in the notations of Definition 2.3, the BFV FHE [26] involves ciphertexts of the form $(-a, -a \cdot s + \Delta \cdot m + \text{noise})$, where $a \sim U(\mathcal{R}_q)$ is a random ring element and $s \sim \chi_s$ is the secret key sampled from some distribution over \mathcal{R} . The standard bootstrapping of BFV can be seen as a restricted programmable bootstrapping for the function $f(x) = \Delta \cdot \lceil x/\Delta \rceil$ that only refreshes the input ciphertext. We can use the random oracle to encode the input x as a ring element $a = H(x) \in \mathcal{R}_q$ and interpret $(-a, 0) \in \mathcal{R}_q^2$ as a noisy BFV ciphertext of the form

$$(-a, 0) = \left(-a, -a \cdot s + \Delta \cdot Q + \underbrace{r}_{\in [-\Delta/2, \Delta/2]} \right) \in \mathcal{R}_q^2,$$

where Q is the quotient obtained by Euclidean division and r is the remainder. By applying non-programmable bootstrapping techniques for BFV, we can homomorphically compute a low-noise encryption $(c, d) \in \mathcal{R}_q^2$ of the same plaintext Q . Note that in order to enable the BFV bootstrapping of the high noise ciphertext, one should set q to be the intermediate modulus to avoid the modulus switching step [28]. Then, we obtain that

$$Q = \lceil ((p/q) \cdot a \cdot s \bmod q) \rceil = \lceil (p/q) \cdot (a \cdot s \bmod q) \rceil \bmod p,$$

so that the bootstrapping algorithm outputs a BFV encryption (c, d) of the RLWR-based PRF $\lceil (p/q) \cdot (a \cdot s \bmod q) \rceil$.

Using the bootstrapping techniques of BFV, we can thus obtain many pseudorandom slots in one bootstrap achieving very high *amortized* throughput. On the downside, it incurs a much higher latency (typically at least in the 10's of seconds) than TFHE. This latency may be prohibitive in applications such as transciphering.

5 Modified PRFs Supporting Homomorphic Evaluation Using Depth-1 Bootstrapping

In Section 3, the TFHE evaluator has to perform two programmable bootstrappings for each plaintext slot in order to evaluate non-negacyclic functions.

In this section, we modify the LWR-based PRF in such a way that it can be evaluated using only one PBS per plaintext slot. In this modified construction, we assume again that $Q = 2N$ where N is the ring dimension used in TFHE.

We start from the previous approach and view $(-\mathbf{a}, 0) \in \mathbb{Z}_{2N}^{n+1}$ as a highly noisy ciphertext that decrypts to $\lfloor (\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N) / \Delta \rfloor$ (in this case, we view the noise as a positive integer in $[0, \Delta)$). We now apply Theorem 2.6 to the negacyclic function $f: \mathbb{Z}_{2N} \rightarrow \mathbb{Z}_q$ defined by:

$$f(x) = \begin{cases} \Delta' \cdot \lfloor \frac{p}{N} \cdot x \rfloor \bmod q & \text{if } x \in [0, N-1] \\ -\Delta' \cdot \lfloor \frac{p}{N} \cdot (x-N) \rfloor \bmod q & \text{if } x \in [N, 2N-1] \end{cases} \quad (6)$$

$$= (-1)^{\text{msb}(x)} \cdot \Delta' \cdot \lfloor \frac{p}{N} \cdot (x \bmod N) \rfloor$$

where $\Delta' = q/2N \in \mathbb{Z}$. By Theorem 2.6, the Eval_{pp} evaluation algorithm outputs a ciphertext encrypting

$$\text{PRF}_{\mathbf{s}}(x) = (-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N)} \left\lfloor \frac{p}{N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) \right\rfloor \bmod p \quad (7)$$

where $\mathbf{a} \triangleq \mathbf{a}(x) = H(x) \in (\mathbb{Z}_{2N})^n$ and $\mathbf{s} \in \{0, 1\}^n$ is the LWE secret key. Note that we hash x onto \mathbb{Z}_{2N} even though the inner product inside the rounding function is reduced modulo N . Another difference with the original construction is that the underlying rounding function is the floor function $\lfloor \cdot \rfloor$ (whereas the initial PRF can use any rounding function like floor, ceiling or the nearest integer although its depth-2 evaluation works for the $\lceil \cdot \rceil$ rounding function). We note that this rounding function was used recently in [29] for similarly small moduli, and that there is no reason to suspect that the use of $\lfloor \cdot \rfloor$ affects hardness [9, Sect. 2]. We later show that one can use the “rounding to nearest integer” function if desired. Although the function (7) is not the standard PRF considered in [14, 29], we can still prove it pseudorandom via a reduction from the pseudorandomness of the *standard* LWR-based PRF.

In more detail, the Eval_{pp} algorithm takes as input the bootstrapping keys $\text{bsk}[j]$, $1 \leq j \leq n$, each of which is a GGSW encryption (with ring dimension N) of the secret key bit s_j under GLWE secret key \mathbf{s}' . We form the test polynomial $v \in \mathbb{Z}_p[X]/(X^N + 1)$ given by $v(X) = \sum_{i=0}^{N-1} v_i X^i$ with coefficients defined as $v_i = \lfloor i \cdot p/N \rfloor \bmod p$ for each $i \in [0, N-1]$.

Letting $\mathbf{a} = (a_1, \dots, a_n)$, define

$$w_j(X) = \begin{cases} v(X) & \text{if } j = 0 \\ X^{s_j a_j} \cdot w_{j-1}(X) & \text{for } 1 \leq j \leq n \end{cases}.$$

Then, due to the congruence $X^N = -1 \pmod{X^N + 1}$, we have

$$X^{-\sum_{j=1}^n s_j \cdot a_j \pmod{2N}} \cdot v(X) = (-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N)} \cdot X^{-\langle \mathbf{a}, \mathbf{s} \rangle \bmod N} \cdot v(X) \pmod{X^N + 1} \quad (8)$$

since

$$\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N = (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) + b \cdot N,$$

where $b = \text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N) \in \{0, 1\}$. In the right-hand-side member of (8), the degree-0 coefficient is thus

$$(-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N)} \cdot v_{\langle \mathbf{a}, \mathbf{s} \rangle \bmod N} = (-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N)} \cdot \lfloor \langle \mathbf{a}, \mathbf{s} \rangle \bmod N \rfloor \cdot p/N \bmod p$$

which is the correct evaluation for $\mathbf{a} = H(x) \in \mathbb{Z}_{2N}^n$. Note that the blind rotation and sample extraction subroutines on the ciphertext $(-\mathbf{a}, 0)$ exactly run this procedure in the encrypted domain. In particular, the *phase* of $(-\mathbf{a}, 0)$ (i.e., $\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N$) picks out the correct coefficient of the test polynomial. To summarize, we have:

$\text{Eval}_{\text{pp}}(\text{bsk}, x)$: Given public parameters $\text{pp} = (n, N, p, q, \beta)$, an evaluation key bsk and an input $x \in \{0, 1\}^\ell$, compute the input-dependent vector $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ and output $\text{SampleExtract} \circ \text{BlindRotate}[f](-\mathbf{a}, 0)$ i.e., $\text{Boot}[f]$ without the keyswitch.

Note that removing the keyswitch outputs an LWE ciphertext of the PRF encrypted under the key \mathbf{s}' interpreted as an LWE secret key. One can then keyswitch this ciphertext to the TFHE secret key in the wider application if required.

Remark 5.1. Although we have presented the case where the PRF key has the same form as the LWE secret key used in TFHE, other alternatives are available. Since the modified LWR PRF is as secure as an LWR PRF, the required LWR dimension n_{LWR} can be smaller than the LWE dimension n used in TFHE. This allows us to use a $\text{bsk}[j]$ for $1 \leq j \leq n_{\text{LWR}}$ in our construction which can significantly increase efficiency as fewer blind rotation steps are carried out. Furthermore, since we do not perform any modulus switch from q to $2N$ unlike in TFHE, there is a possibility to reduce N while increasing the module rank for efficiency. When doing this, the hardness of the corresponding LWR problem must be checked since N is the largest modulus of the LWR assumption, so that a smaller N induces a smaller deterministically generated noise (of magnitude $< N/p$). Note that it would be useful in terms of memory requirement to reuse the TFHE bootstrapping key for PRF evaluation, but this would violate the principle of key separation.

5.1 Pseudorandomness of the Modified PRF

We have shown that one can evaluate the function

$$\text{PRF}_s(x) \triangleq (-1)^{\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N)} \left\lfloor \frac{p}{N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) \right\rfloor \pmod{p}$$

with a depth-1 programmable bootstrapping. However, we must also show that PRF_s is indeed pseudorandom as it is a *modified version* of the standard LWR PRF described in Section 2.2. As a standard TFHE parameter choice, we assume that $N = 2^{\ell_N}$ and $p = 2^{\ell_p}$ where $k \triangleq \ell_N - \ell_p > 0$.

Consider any $y \in \mathbb{Z}$ bounded in absolute value by $2N \cdot B$ for appropriately large B . It is easy to see that $y \bmod 2N$ is simply the lowest $\ell_N + 1$ bits of $y + 2NB$. Further, $y \bmod N$ is the lowest ℓ_N bits of $y + 2NB$. We can additionally interpret the operation on y described by

$$\left\lfloor \frac{2p}{2N} \cdot (y \bmod 2N) \right\rfloor \bmod 2p \tag{9}$$

in terms of operations on bits too. In particular, one can think of the above operation as taking the $\ell_N + 1$ bottom bits of $y + 2NB$ and then dropping the k least significant bits. This leaves $\ell_N + 1 - k = \ell_p + 1$ bits which represents an integer modulo $2p$. This interpretation implies:

OBSERVATION 5.2. For $N > p$ both powers-of-two and any $y \in \mathbb{Z}$,

$$\text{msb}(y \bmod 2N) = \text{msb}\left(\left\lfloor \frac{2p}{2N} \cdot (y \bmod 2N) \right\rfloor \bmod 2p\right).$$

Changing $(2N, 2p)$ to (N, p) in Equation (9) changes the operation to “take ℓ_N bottom bits of $y + 2NB$ and then drop the k least significant bits”. Therefore, we have the following:

OBSERVATION 5.3. For $N > p$ both powers-of-two and any $y \in \mathbb{Z}$, one can compute $\left\lfloor \frac{p}{N} \cdot (y \bmod N) \right\rfloor \bmod p$ from the bits of the quantity in (9) by simply dropping the most significant bit.

With these two observations, we prove that the modified PRF_s is as secure as a standard LWR-based PRF (denoted as G_s), which is identical to the one recalled in Section 2.2 except that the rounding function $\lceil \cdot \rceil$ is replaced by $\lfloor \cdot \rfloor$.

LEMMA 5.4. Assume $p = 2^{\ell_p}$ and $N = 2^{\ell_N}$ are powers-of-two such that $k \triangleq \ell_N - \ell_p > 0$. Take a random function $H: \{0, 1\}^* \rightarrow (\mathbb{Z}_{2N})^n$ and assume that $G_s: \{0, 1\}^* \rightarrow \mathbb{Z}_{2p}$,

$$G_s(x) \triangleq \left\lfloor \frac{2p}{2N} \cdot (\langle H(x), s \rangle \bmod 2N) \right\rfloor \bmod 2p$$

is pseudorandom for s sampled uniformly from $\{0, 1\}^n$. Then, the function $\text{PRF}_s: \{0, 1\}^* \rightarrow \mathbb{Z}_p$,

$$\text{PRF}_s(x) \triangleq (-1)^{\text{msb}(\langle H(x), s \rangle \bmod 2N)} \cdot \left\lfloor \frac{p}{N} \cdot (\langle H(x), s \rangle \bmod N) \right\rfloor \bmod p$$

is pseudorandom where s is also sampled uniformly in $\{0, 1\}^n$.

PROOF. We describe a reduction \mathcal{A} that attempts to build a PPT PRF distinguisher for G from any PPT distinguisher \mathcal{D} for PRF. The reduction is as follows:

- When \mathcal{D} wishes to query its oracle on input x , \mathcal{A} forwards the request on to its challenger, receiving $g \in \mathbb{Z}_{2p}$ in response.
- Define $g_0 \triangleq \text{msb}(g)$ and g' to be the integer in $\{0, \dots, p-1\}$ resulting from dropping the MSB of g . \mathcal{A} sends $f \triangleq (-1)^{g_0} \cdot g' \bmod p$ back to \mathcal{D} in response to the query x .
- \mathcal{A} ultimately outputs whatever \mathcal{D} does.

When \mathcal{A} 's challenger is returning uniform values for g , \mathcal{A} 's response f that is sent to \mathcal{D} is clearly uniform. On the other hand, when \mathcal{A} 's challenger is using G_s to compute g , we can use Observation 5.2 to show that the exponent of (-1) is correct and Observation 5.3 to show that the remaining term (i.e., g') is correctly computed for PRF_s . Therefore, \mathcal{A} perfectly simulates \mathcal{D} 's PRF challenger which implies that \mathcal{D} 's advantage against the pseudorandomness of PRF_s is equal to that of \mathcal{A} 's advantage against G_s . By assumption on the pseudorandomness of G_s , PRF_s must also be pseudorandom. \square

5.2 Replacing Floor with Nearest Integer Rounding

It is of course possible to evaluate the modified PRF in depth one when the rounding function is $\lceil \cdot \rceil$ instead of $\lfloor \cdot \rfloor$. This only requires to modify the coefficients of the test polynomial $v(X)$ accordingly. In the interpretation of $(-a, 0)$ as a noisy ciphertext, we need to replace $\lfloor \cdot \rfloor$ by $\lceil \cdot \rceil$ in the negacyclic function of (6) and go back to our view of the “noise” as an integer in $[-\Delta/2, \Delta/2)$.

However, we need to slightly modify the argument proving the pseudorandomness of the resulting function. We may replace the definition of PRF_s by

$$\text{PRF}_s(x) \triangleq (-1)^{\text{msb}(\langle H(x), s \rangle \bmod 2N)} \cdot \left\lceil \frac{p}{N} \cdot (\langle H(x), s \rangle \bmod N) \right\rceil \bmod p \quad (10)$$

in Lemma 5.4 whilst changing the parameters (p, N) in the definition of G_s . To see how, note that we can write $\lceil y \rceil = \lfloor y \rfloor + b_y$ where b_y is the bit of y just below the fixed binary point. Essentially, the reduction in the proof needs an extra bit in order to convert the floor function to the rounding function (i.e., to simulate the function in Equation (10)). In order to gain access to this bit, one needs to increase p in the definition of G_s by a factor of 2. The conclusion is the following lemma.

LEMMA 5.5. *Assume $p = 2^{\ell_p}$ and $N = 2^{\ell_N}$ are powers-of-two such that $k \triangleq \ell_N - \ell_p - 1 > 0$. Take a random function $H: \{0, 1\}^* \rightarrow (\mathbb{Z}_{2N})^n$ and assume that $G_s: \{0, 1\}^* \rightarrow \mathbb{Z}_{4p}$,*

$$G_s(x) \triangleq \left\lfloor \frac{4p}{2N} \cdot (\langle H(x), s \rangle \bmod 2N) \right\rfloor \bmod 4p$$

is pseudorandom for s sampled uniformly from $\{0, 1\}^n$. Then, the function $\text{PRF}_s: \{0, 1\}^ \rightarrow \mathbb{Z}_p$,*

$$\text{PRF}_s(x) \triangleq (-1)^{\text{msb}(\langle H(x), s \rangle \bmod 2N)} \cdot \left\lfloor \frac{p}{N} \cdot (\langle H(x), s \rangle \bmod N) \right\rfloor \bmod p$$

is pseudorandom where s is also sampled uniformly in $\{0, 1\}^n$. \square

5.3 Reducing the Range for Padding

In TFHE and wider applications, it is often useful to have one or more padding bits in the plaintext space. A common practice is to always leave the most significant bit of a plaintext to 0 when this plaintext has to be involved in further homomorphic computations. So far, we have not considered this issue.⁴

In the depth-2 construction of Section 3, a simple solution is to modify the test polynomial of the second PBS and shift the bits of all coefficients to the right. In the depth-1 case, we cannot do this since it would unsuitably interfere with the $(-1)^{\text{msb}(\langle a, s \rangle \bmod 2N)}$ factor during the blind rotation.

In the depth-1 case, we can address this problem by considering yet another modified PRF. As before, we assume that the full plaintext space has a power-of-two modulus $p = 2^{\ell_p}$. However, we now introduce a usable plaintext modulus $p' = 2^{\ell_{p'}} < p$ meaning that we have $\ell_p - \ell_{p'}$ padding bits. The modified PRF can then be described as

$$\begin{aligned} \text{PRF}_s(x) &\triangleq (-1)^{\text{msb}(\langle a, s \rangle \bmod 2N)} \cdot \left(\left\lfloor \frac{p'}{2N} \cdot (\langle a, s \rangle \bmod N) \right\rfloor + \frac{1}{2} \right) \\ &\quad + \frac{p'-1}{2} \bmod p \end{aligned} \tag{11}$$

$$\begin{aligned} &= (-1)^{\text{msb}(\langle a, s \rangle \bmod 2N)} \cdot \left\lfloor \frac{p'}{2N} \cdot (\langle a, s \rangle \bmod N) \right\rfloor \\ &\quad + \frac{p'}{2} - \text{msb}(\langle a, s \rangle \bmod 2N) \end{aligned} \tag{12}$$

which ranges over $[0, p' - 1]$. To evaluate this function, we can first apply Theorem 2.6 with $Q = 2N$ to the function

$$f: \mathbb{Z}_{2N} \rightarrow \mathbb{Z}_q, x \mapsto f(x) \triangleq (-1)^{\text{msb}(x)} \cdot \frac{q}{p} \cdot \left(\left\lfloor \frac{p'}{2N} \cdot (x \bmod N) \right\rfloor + \frac{1}{2} \right)$$

which is negacyclic since

$$f(x + N) = (-1) \cdot f(x) = -f(x) \bmod 2N \quad \forall x \in [0, N - 1] .$$

⁴In the application to transpiling in Section 7, it is not necessary to keep the padding bit clear and we can use the full precision of the message space.

Then, after the PBS, we can add the term $\frac{q}{p} \cdot \frac{p'-1}{2}$ to the evaluated ciphertext.⁵ Due to the additive homomorphism, this yields an encryption of the correct PRF value (11).

We now argue the pseudorandomness of the function in (11). Note that for any $\mathbf{a}, \mathbf{s} \in \mathbb{Z}^n$,

$$\begin{aligned} \left\lfloor \frac{p'/2}{N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) \right\rfloor \bmod \frac{p'}{2} &= \left\lfloor \frac{p'/2}{N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) \right\rfloor \bmod p \\ &= \left\lfloor \frac{p'/2}{N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod N) \right\rfloor \end{aligned}$$

as the modular reduction $\bmod(p'/2)$ and p is inconsequential. By Observation 5.3, if $N > p'/2$, one can compute the above by dropping the most significant bit of

$$\left\lfloor \frac{p'}{2N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N) \right\rfloor. \quad (13)$$

Furthermore, by Observation 5.2, if $N > p'/2$,

$$\text{msb}(\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N) = \text{msb}\left(\left\lfloor \frac{p'}{2N} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod 2N) \right\rfloor\right). \quad (14)$$

To prove the pseudorandomness property, we rely on the following lemma whose proof is similar to that of the un-padded case.

LEMMA 5.6. *Assume $p' = 2^{\ell p'}$ and $N = 2^{\ell N}$ are powers-of-two such that $N > p'/2$. Take a random function $H: \{0, 1\}^* \rightarrow (\mathbb{Z}_{2N})^n$ and assume that $G_s: \{0, 1\}^* \rightarrow \mathbb{Z}_{p'}$,*

$$G_s(x) \triangleq \left\lfloor \frac{p'}{2N} \cdot (\langle H(x), \mathbf{s} \rangle \bmod 2N) \right\rfloor$$

is pseudorandom for \mathbf{s} sampled uniformly from $\{0, 1\}^n$. Then, the function

$$\text{PRF}'_s: \{0, 1\}^* \rightarrow \left\{ \pm \frac{1}{2}, \pm(1 + \frac{1}{2}), \pm(2 + \frac{1}{2}), \dots, \pm(\frac{p'}{2} - 1 + \frac{1}{2}) \right\}$$

$$\text{PRF}'_s(x) \triangleq (-1)^{\text{msb}(\langle H(x), \mathbf{s} \rangle \bmod 2N)} \cdot \left(\left\lfloor \frac{p'}{2N} \cdot (\langle H(x), \mathbf{s} \rangle \bmod N) \right\rfloor + \frac{1}{2} \right)$$

is pseudorandom where \mathbf{s} is also sampled uniformly in $\{0, 1\}^n$.

PROOF. We describe a reduction \mathcal{A} that attempts to build a PPT PRF distinguisher for G from any PPT distinguisher \mathcal{D} for PRF' . The reduction is as follows:

- When \mathcal{D} wishes to query its oracle on input x , \mathcal{A} forwards the request on to its challenger, receiving $g \in \mathbb{Z}_{p'}$ in response.
- Define $g_0 \triangleq \text{msb}(g) \in \{0, 1\}$ and let g' the integer in $\{0, \dots, p'/2 - 1\}$ resulting from dropping the MSB of g . \mathcal{A} sends $f \triangleq (-1)^{g_0} \cdot (g' + \frac{1}{2})$ back to \mathcal{D} in response to the query x .
- \mathcal{A} ultimately outputs whatever \mathcal{D} does.

When \mathcal{A} 's challenger is returning uniform values for g , \mathcal{A} 's response f that is sent to \mathcal{D} is clearly uniform in the appropriate range as \mathcal{A} 's operations are invertible. On the other hand, when \mathcal{A} 's challenger is using G_s to compute g , we can use Equation (14) to show that the exponent of (-1) is correct and Observation 5.3 with the parametrization in (13) to show that the remaining term (i.e., g') is correctly computed for PRF'_s . Therefore, \mathcal{A} perfectly simulates \mathcal{D} 's PRF challenger which implies that \mathcal{D} 's advantage against the pseudorandomness of PRF'_s is equal to that of \mathcal{A} 's advantage against G_s . By assumption on the pseudorandomness of G_s , PRF'_s must also be pseudorandom. \square

⁵This does not quite correspond to adding $\frac{p'}{2} - \frac{1}{2}$ to the plaintext since $1/2$ is not defined modulo p . However, it still provides an encryption of the correct PRF evaluation (11) after the final addition.

We finally complete the proof of pseudorandomness for the function defined in (11) by additively shifting the pseudorandom function PRF'_S from the above lemma.

6 Implementation and Performance

In order to test our depth-1 construction in practice, we use the TFHE-rs library using an AWS hpc7a.96xlarge instance with 4th Gen AMD EPYC processor, 768 GiB total RAM and using AVX512 on a single thread. We also provide benchmarks run on a 2022 Apple Macbook Pro with an Apple M2 chip and 8 GB RAM. In particular, we test the latency (specifically, the time required to perform the blind rotation step) for the homomorphic evaluation of the PRF in Section 5.1. To ensure the practical relevance of our results, we use practical TFHE parameter sets PARAM_MESSAGE_2_CARRY_2 (where the bootstrapping key has size 23.9MB) and PARAM_MESSAGE_1_CARRY_1 (with bootstrapping key size 11 MB) as opposed to some bespoke parameter set. Our benchmarks do not include the computation of the hash value $H(x)$ and therefore mimic cases where the hash values have been precomputed e.g. transciphering where the server stores the hashes. In any case, the time taken to hash will be negligible compared to the runtime of homomorphic PRF evaluations, especially for short 128-bit or 256-bit inputs. Note that the plaintext space for PARAM_MESSAGE_2_CARRY_2 is effectively \mathbb{Z}_p for $p = 2^5$ as there are 2 “carry” bits, 2 “message” bits and a padding bit (equaling 5 bits in total). Note that the nomenclature arises from the fact that TFHE-rs is designed to implement large integer arithmetic. Further, the ring dimension used is $N = 2048 = 2^{11}$. Using the optimization outlined in Remark 5.1 and Lemma 5.4, we require that $\text{LWR}_{n_{\text{LWR}}, m, Q=2N, 2p, U(\{0,1\})}$ is hard. Using the lattice estimator for unbounded m and modeling LWR as LWE with uniform rounding noise, we conclude that we may choose $n_{\text{LWR}} = 445$ for an estimated 128 bits of security. Taking n_{LWE} to be the TFHE LWE dimension leads to a PRF evaluation key that has a size of approximately $n_{\text{LWR}}/n_{\text{LWE}} = 445/742 \approx 60\%$ of the size of a regular bootstrapping key. Making this choice leads to a latency of 6.0328 ms (averaged over 60 seconds). Stated differently, we obtain a throughput of around 829 encrypted pseudorandom bits on a single thread. Naturally, one can increase this throughput by using multiple threads.

For the PARAM_MESSAGE_1_CARRY_1 parameter set, the plaintext space is effectively set to $p = 2^3$ and the ring dimension is $N = 512$ leading to $n_{\text{LWR}} = 409$. The PRF evaluation key size in this case is also approximately $n_{\text{LWR}}/n_{\text{LWE}} = 409/702 \approx 60\%$ of the bootstrapping key. The latency in the depth-1 construction is around 2.8029 ms (averaged over 60 seconds) leading to around 1070 pseudorandom bits per second on a single thread. The results are summarized in Table 1.

Table 1. Single threaded experimental results. The first reported result is on a hpc7a.96xlarge instance whereas the second is on an Apple Macbook Pro.

Parameter set	MESSAGE_1_CARRY_1	MESSAGE_2_CARRY_2
Plaintext bits	3	5
Latency (ms)	2.803 / 3.714	6.033 / 8.187
Throughput (bits/s)	1070 / 808	829 / 611
Bootstrap Key	11.0 MB	23.9 MB
PRF Eval Key	6.4 MB	13.9 MB

Further Optimization. As mentioned in Remark 5.1, our conditions on N are different to those in TFHE parameter selection and reducing the size of N can offer improvements in efficiency. In particular, we do not have to worry about any modulus switching error from q to $2N$ in our construction. We do however have to worry about choosing N and n_{LWR} so that LWR holds with respect to moduli $2N$ and $2p$. We describe our strategy for optimizing parameters next, defining k to be the module rank (as defined in Definition 2.5) in the TFHE GLWE assumption. Suppose that we are using parameters $(N, k, n_{\text{LWR}}$ with a GLWE/GGSW. Then, we can map to some $k', N' = N/2^\kappa$ (for some integer $\kappa < \log_2(N)$) and set n'_{LWR} such that the LWR assumption in Lemma 5.4 holds.

Next, we must ensure that the new parameters are chosen so that the output of our construction looks like a bootstrapped TFHE PARAM_MESSAGE_X_CARRY_X ciphertext. A detail is that these parameter sets use the more efficient “keyswitch then blind rotate” [10] pattern: i.e., the output to the PBS is a $(k \cdot N + 1)$ -dimensional LWE ciphertext. On the other hand, the output to our optimized scheme would be a $(k' \cdot N' + 1)$ -dimensional LWE ciphertext. We will always assume that whenever $(k', N') \neq (k, N)$, we use distinct secret keys. In other words, we do not share a secret key between a GLWE in dimensions (k', N') and (k, N) . This choice may be overly conservative, but may potentially allow for a less heuristic security guarantee. Then, to summarize we pick parameters to ensure that a blind rotation followed by a “ $k'N'$ to kN ” keyswitch results in a ciphertext under the correct key, with the noise level of a bootstrapped ciphertext. We note that when using the “blind rotate then keyswitch” pattern, the $k'N'$ to kN keyswitch is unnecessary as one would simply keyswitch directly down to the LWE dimension to obtain a bootstrapped ciphertext. However, this pattern generally leads to an overall less efficient FHE application.

Using the optimization techniques from [10], we obtain new parameters and run benchmarks. Unfortunately, the optimized PARAM_MESSAGE_1_CARRY_1 setting did not improve throughput, even when ignoring the $k'N'$ to kN keyswitch. Therefore, we do not report on this parameter set. However, the optimization and improved performance of the PARAM_MESSAGE_2_CARRY_2 is reported in Table 2. A good choice for the dimension N' in terms of latency and key size appears to be 512. At this dimension, the throughput increases by 16% (or 44%) on the hpc7a.96xlarge (respectively, laptop) whereas the evaluation key shrinks by around 5 MB or 36% compared to results in Table 1. As can be seen in the latter table, when N' becomes small, the LWR dimension increases dramatically hindering performance. Note that optimizing without the $k'N'$ to kN keyswitch does not appear to change the throughput here either. In particular, removing the keyswitch still does not allow us to pick $k'N' < kN$ or improve the blind rotation decomposition parameters (which are already optimal in PARAM_MESSAGE_2_CARRY_2). We note that the key compression techniques of [43] could also be applied to compress the size of the PRF evaluation key. However, due to the tight noise constraints considered when optimizing the parameters used in our benchmarks, we do not expect their techniques to be directly applicable to the parameter sets used in this work, since it would add additional noise in the PRF evaluation key. Instead, to utilize these techniques, it would be necessary to use larger parameters, as in [2, 43]

7 Application to Transciphering

We now describe an application of our homomorphic PRF evaluation algorithms to transciphering. We consider a client wishing to send large amounts of data to the cloud that will eventually be used in an FHE application—think for example of image processing over

Table 2. Single threaded experimental results with “optimized” parameters for PARAM_MESSAGE_2_CARRY_2. The first reported result is on a hpc7a.96xlarge instance whereas the second is on an Apple Macbook Pro.

Parameter set	Opt-2-2-256	Opt-2-2-512	Opt-2-2-1024
$(N', k', n'_{\text{LWR}})$	(256, 8, 980)	(512, 4, 455)	(1024, 2, 455)
Plaintext bits	5	5	5
Latency (ms)	14.267 / 13.388	5.205 / 5.675	5.156 / 6.379
Throughput (bits/s)	350 / 373	961 / 881	970 / 784
PRF Eval Key	17.2 MB	8.9 MB	10.7 MB

encrypted data. Instead of sending FHE encryptions of the data on the cloud, transcribing allows the transmission of symmetric key encryptions to dramatically reduce cloud bandwidth requirements. Specifically, the symmetric encryptions sent on the cloud do not need to be any larger than the original data, whereas FHE ciphertexts would be much larger. As an example of this application, a symmetric encryption of a message M can take the form

$$(x, M \oplus \text{PRF}_k(x))$$

where x is a random string of sufficient length chosen by the sender.⁶ Then, in order to obtain an FHE encryption of M , the cloud can use an FHE encryption of k to compute an FHE encryption of $\text{PRF}_k(x)$. With this, the cloud can *homomorphically* subtract $\text{PRF}_k(x)$ to obtain an FHE encryption of M that can be computed on further. Note that other forms of symmetric encryption (e.g. certain block cipher modes of operation) may also be used provided that PRF inputs remain public.

As a concrete example, we may consider our depth-1 construction of Section 5.1 (or Section 5.2). This yields a symmetric encryption scheme where the sender chooses a random $x \leftarrow U(\{0, 1\}^\lambda)$ and encrypts $\mathbf{M} \in \mathbb{Z}_p^m$ using PRF secret key $k \triangleq \mathbf{s} \leftarrow U(\{0, 1\}^n)$ by computing $(x, \mathbf{c}) = (x, \mathbf{M} + \text{PRF}_s(x) \bmod p)$, where

$$\forall i \in [m] : (\text{PRF}_s(x))_i \triangleq (-1)^{\text{msb}(\langle H(x, i), \mathbf{s} \rangle \bmod 2N)} \cdot \left\lfloor \frac{p}{N} \cdot (\langle H(x, i), \mathbf{s} \rangle \bmod N) \right\rfloor, \quad (15)$$

for a hash function $H: \{0, 1\}^* \rightarrow \mathbb{Z}_p^n$ modeled as a random oracle. Then, the encryptor sends the ciphertext (x, \mathbf{c}) for storage. We assume that the evaluator has a copy of the corresponding public bootstrapping key $\text{bsk} = \text{GGSW.Enc}_s(\mathbf{s})$ for some secret key \mathbf{s}' . From (x, \mathbf{c}) and bsk , the evaluator can compute

$$(\mathbf{A}, \mathbf{b} = \mathbf{A}^\top \mathbf{s}' + \mathbf{e} + \Delta \cdot \text{PRF}_s(x)) = \text{Eval}_{\text{pp}}(\text{bsk}, x) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m,$$

where $\|\mathbf{e}\|_\infty \leq \beta$ for the bound β of Theorem 2.6. Then, the evaluator obtains

$$(-\mathbf{A}, -\mathbf{b} + \Delta \cdot \mathbf{c} \bmod q)$$

which is an encryption of $\mathbf{M} \in \mathbb{Z}_p^m$ under the secret key \mathbf{s}' since $(\mathbf{0}, \Delta \cdot \mathbf{c})$ is a (trivial) encryption of \mathbf{c} .

The main advantage of the proposed solution is that the encryption of m plaintexts in \mathbb{Z}_p only requires sending a random seed x (typically, a 256-bit value) along with m values modulo p . This is much better than the plain solution that would send a matrix \mathbf{A} of $m \times n$ entries

⁶This construction satisfies the definition of CPA security (see, e.g., [41, Chapter 3]) for secret-key encryption schemes assuming that the underlying PRF is pseudorandom.

modulo q plus m values modulo q (typically, n is of the order of 1000). This is even better than the folklore improved solution consisting in sending a seed σ for building matrix A along with m values modulo q . Our solution trades modulo- q values against modulo- p values, where $p \ll q$; typically, $p = 2^4$ and $q = 2^{64}$. Compared to the improved solution, this saves $m \cdot \log_2 q/p$ bits of transmission. For example, for (1024×1024) 8-bit gray-scale images, with $p = 16$ and $q = 2^{64}$ (and thus $m = 2^{21}$), this results in saving more than 10^8 bits per encrypted image.

We also observe that, as in stream-cipher-based solutions, we can dynamically decide to change the number m of plaintext blocks if the need arises without changing anything to the parameters or the key material. It only requires to update the number of indexes i when the encryptor computes (15).

References

- [1] Shweta Agrawal, Shafi Goldwasser, and Saleet Mossel. 2021. Deniable fully homomorphic encryption from learning with errors. In *Advances in Cryptology – CRYPTO 2021, Part II (Lecture Notes in Computer Science, Vol. 12826)*, T. Malkin and C. Peikert (Eds.). Springer, 641–670. https://doi.org/10.1007/978-3-030-84245-1_22
- [2] Martin R. Albrecht, Alex Davidson, Amit Deo, and Daniel Gardham. 2024. Crypto dark matter on the torus: Oblivious PRFs from shallow PRFs and FHE. In *Advances in Cryptology – EUROCRYPT 2024 (Lecture Notes in Computer Science)*, M. Joye and G. Leander (Eds.). Springer. To appear (Preprint available at <https://ia.cr/2023/232>).
- [3] Martin R. Albrecht, Rachel Player, and Sam Scott. 2015. On the concrete hardness of learning with errors. *J. Math. Cryptol.* 9, 3 (2015), 169–203. <https://doi.org/10.1515/jmc-2015-0016>
- [4] Martin R. Albrecht, Christian Rechberger, Thomas Schneider, Tyge Tiessen, and Michael Zohner. 2015. Ciphers for MPC and FHE. In *Advances in Cryptology – EUROCRYPT 2015, Part I (Lecture Notes in Computer Science, Vol. 9056)*, E. Oswald and M. Fischlin (Eds.). Springer, 430–454. https://doi.org/10.1007/978-3-662-46800-5_17
- [5] Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. 2013. Learning with rounding, revisited. In *Advances in Cryptology – CRYPTO 2013, Part I (Lecture Notes in Computer Science, Vol. 8042)*, R. Canetti and J. A. Garay (Eds.). Springer, 57–74. https://doi.org/10.1007/978-3-642-40041-4_4
- [6] Tomer Ashur, Mohammad Mahzoun, and Dilara Toprakhisar. 2022. Chaghri: A FHE-friendly block cipher. In *2022 ACM SIGSAC Conference on Computer and Communications Security*, H. Yin et al. (Eds.). ACM, 139–150. <https://doi.org/10.1145/3548606.3559364>
- [7] Youngjin Bae, Jung Hee Cheon, Jaehyung Kim, Jai Hyun Park, and Damien Stehlé. 2023. HERMES: Efficient ring packing using MLWE ciphertexts and application to transciphering. In *Advances in Cryptology – CRYPTO 2023, Part IV (Lecture Notes in Computer Science, Vol. 14084)*, H. Handschuh and A. Lysyanskaya (Eds.). Springer, 37–69. https://doi.org/10.1007/978-3-031-38551-3_2
- [8] Thibault Balenbois, Jean-Baptiste Orfila, and Nigel P. Smart. 2023. Trivial transciphering with Trivium and TFHE. In *11th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC 2023)*, M. Brenner, A. Costache, and K. Rohloff (Eds.). ACM Press, 69–78. <https://doi.org/10.1145/3605759.3625255>
- [9] Abhishek Banerjee, Chris Peikert, and Alon Rosen. 2012. Pseudorandom functions and lattices. In *Advances in Cryptology – EUROCRYPT 2012 (Lecture Notes in Computer Science, Vol. 7237)*, D. Pointcheval and T. Johansson (Eds.). Springer, 719–737. https://doi.org/10.1007/978-3-642-29011-4_42
- [10] Loris Bergerat, Anas Boudi, Quentin Bourgerie, Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. 2023. Parameter optimization and larger precision for (T)FHE. *J. Cryptol.* 36, 3 (2023), 28. <https://doi.org/10.1007/S00145-023-09463-5>
- [11] Andrej Bogdanov, Siyao Guo, Daniel Masny, Silas Richelson, and Alon Rosen. 2016. On the hardness of learning with rounding over small modulus. In *Theory of Cryptography (TCC 2016), Part I (Lecture Notes in Computer Science, Vol. 9562)*, E. Kushilevitz and T. Malkin (Eds.). Springer, 209–224. https://doi.org/10.1007/978-3-662-49096-9_9
- [12] Andrej Bogdanov and Alon Rosen. 2017. Pseudorandom functions: Three decades later. In *Tutorials on the Foundations of Cryptography*, Y. Lindell (Ed.). Springer, Chapter 3, 79–158. https://doi.org/10.1007/978-3-319-57048-8_3
- [13] Nicolas Bon, David Pointcheval, and Matthieu Rivain. 2024. Optimized homomorphic evaluation of Boolean functions. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2024, 3 (2024). To appear (Preprint available at <https://ia.cr/2024/15>).

- [//ia.cr/2023/1589](https://ia.cr/2023/1589)).
- [14] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. 2013. Key homomorphic PRFs and their applications. In *Advances in Cryptology – CRYPTO 2013, Part I (Lecture Notes in Computer Science, Vol. 8042)*, R. Canetti and J. A. Garay (Eds.). Springer, 410–428. https://doi.org/10.1007/978-3-642-40041-4_23
 - [15] Zvika Brakerski. 2012. Fully homomorphic encryption without modulus switching from classical GapSVP. In *Advances in Cryptology – CRYPTO 2012 (Lecture Notes in Computer Science, Vol. 7417)*, R. Safavi-Naini and R. Canetti (Eds.), Springer, 868–886. https://doi.org/10.1007/978-3-642-32009-5_50
 - [16] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. 2019. Leveraging linear decryption: Rate-1 Fully-homomorphic encryption and time-lock puzzles. In *Theory of Cryptography (TCC 2019), Part II (Lecture Notes in Computer Science, Vol. 11892)*, D. Hofheinz and A. Rosen (Eds.). Springer, 407–437. https://doi.org/10.1007/978-3-030-36033-7_16
 - [17] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. 2012. (Leveled) fully homomorphic encryption without bootstrapping. In *3rd Innovations in Theoretical Computer Science (ITCS 2012)*, S. Goldwasser (Ed.). ACM Press, 309–325. <https://doi.org/10.1145/2090236.2090262>
 - [18] Anne Canteaut, Sergiu Carpov, Caroline Fontaine, Tancrede Lepoint, Maria Naya-Plasencia, Pascal Paillier, and Renaud Sirdey. 2016. Stream ciphers: A practical solution for efficient homomorphic-ciphertext compression. In *Fast Software Encryption (FSE 2016) (Lecture Notes in Computer Science, Vol. 9783)*, T. Peyrin (Ed.). Springer, 313–333. https://doi.org/10.1007/978-3-662-52993-5_16
 - [19] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. 2020. TFHE: Fast fully homomorphic encryption over the torus. *J. Cryptol.* 33, 1 (2020), 34–91. <https://doi.org/10.1007/s00145-019-09319-x>
 - [20] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. 1998. Private Information Retrieval. *J. ACM* 45, 6 (1998), 965–981. <https://doi.org/10.1145/293347.293350>
 - [21] Carlos Cid, John Petter Indrøy, and Håvard Raddum. 2022. FASTA: A stream cipher for fast FHE evaluation. In *Topics in Cryptology – CT-RSA 2022 (Lecture Notes in Computer Science, Vol. 13161)*, S. D. Galbraith (Ed.). Springer, 451–483. https://doi.org/10.1007/978-3-030-95312-6_19
 - [22] Orel Cosserson, Clément Hoffmann, Pierrick Méaux, and François-Xavier Standaert. 2022. Towards case-optimized hybrid homomorphic encryption: Featuring the Elisabeth stream cipher. In *Advances in Cryptology – ASIACRYPT 2022, Part III (Lecture Notes in Computer Science, Vol. 13793)*, S. Agrawal and D. Lin (Eds.). Springer, 32–67. https://doi.org/10.1007/978-3-031-22969-5_2
 - [23] Morten Dahl, Clément Danjou, Daniel Demmler, Tore Frederiksen, Peter Ivanov, Marc Joye, Dragos Rotaru, Nigel Smart, and Louis Tremblay Thibault. 2023. *fhEVM: Confidential EVM smart contracts using fully homomorphic encryption*. White Paper. Zama. <https://github.com/zama-ai/fhevml/raw/main/fhevml-whitepaper.pdf>
 - [24] Christoph Dobraunig, Lorenzo Grassi, Lukas Helming, Christian Rechberger, Markus Schofnegger, and Roman Walch. 2023. Pasta: A case for hybrid homomorphic encryption. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2023, 3 (2023), 30–73. <https://doi.org/10.46586/TCHES.V2023.I3.30-73>
 - [25] Léo Ducas and Daniele Micciancio. 2015. FHEW: Bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology – EUROCRYPT 2015, Part I (Lecture Notes in Computer Science, Vol. 9056)*, E. Oswald and M. Fischlin (Eds.). Springer, 617–640. https://doi.org/10.1007/978-3-662-46800-5_24
 - [26] Junfeng Fan and Frederik Vercauteren. 2012. Somewhat practical fully homomorphic encryption. *Cryptology ePrint Archive*, Paper 2012/144. <https://ia.cr/2012/144>
 - [27] Ben Fisch, Arthur Lazzaretti, Zeyu Liu, and Charalampos Papamanthou. 2024. Single server PIR via homomorphic Thorp shuffles. *Cryptology ePrint Archive*, Paper 2024/482. <https://ia.cr/2024/482>
 - [28] Robin Geelen and Frederik Vercauteren. 2023. Bootstrapping for BGV and BFV Revisited. *J. Cryptol.* 36, 2 (2023), 12. <https://doi.org/10.1007/S00145-023-09454-6>
 - [29] Matthias Geihs and Hart Montgomery. 2024. LaKey: Efficient lattice-based distributed PRFs enable scalable distributed key management. In *33rd USENIX Security Symposium*, D. Balzarotti and W. Xu (Eds.). USENIX. To appear (Preprint available at <https://ia.cr/2023/1254>).
 - [30] Craig Gentry. 2009. *A Fully Homomorphic Encryption Scheme*. Ph. D. Dissertation. Stanford University. <https://crypto.stanford.edu/craig>
 - [31] Craig Gentry, Shai Halevi, and Nigel P. Smart. 2012. Homomorphic evaluation of the AES circuit. In *Advances in Cryptology – CRYPTO 2012 (Lecture Notes in Computer Science, Vol. 7417)*, R. Safavi-Naini and R. Canetti (Eds.). Springer, 850–867. https://doi.org/10.1007/978-3-642-32009-5_49
 - [32] Craig Gentry, Amit Sahai, and Brent Waters. 2013. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Advances in Cryptology – CRYPTO 2013, Part I (Lecture Notes in Computer Science, Vol. 8042)*, R. Canetti and J. A. Garay (Eds.). Springer, 75–92. https://doi.org/10.1007/978-3-642-40041-4_5

- [33] Henri Gilbert, Rachele Heim Boissier, Jérémy Jean, and Jean-René Reinhard. 2023. Cryptanalysis of Elisabeth-4. In *Advances in Cryptology – ASIACRYPT 2023, Part III (Lecture Notes in Computer Science, Vol. 14440)*, J. Guo and R. Steinfield (Eds.). Springer, 256–284. https://doi.org/10.1007/978-981-99-8727-6_9
- [34] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. 1986. How to construct random functions. *J. ACM* 33, 4 (1986), 792–807. <https://doi.org/10.1145/6490.6503>
- [35] Lorenzo Grassi, Irati Manterola Ayala, Martha Norberg Hovd, Morten Øyngarden, Håvard Raddum, and Qingju Wang. 2023. Cryptanalysis of symmetric primitives over rings and a key recovery attack on Rubato. In *Advances in Cryptology – CRYPTO 2023, Part III (Lecture Notes in Computer Science, Vol. 14083)*, H. Handschuh and A. Lysyanskaya (Eds.). Springer, 305–339. https://doi.org/10.1007/978-3-031-38548-3_11
- [36] Jincheol Ha, Seongkwang Kim, Byeonghak Lee, Jooyoung Lee, and Mincheol Son. 2022. Rubato: Noisy ciphers for approximate homomorphic encryption. In *Advances in Cryptology – EUROCRYPT 2022, Part I (Lecture Notes in Computer Science, Vol. 13275)*, O. Dunkelman and S. Dziembowski (Eds.). Springer, 581–610. https://doi.org/10.1007/978-3-031-06944-4_20
- [37] Jiahui He, Kai Hu, Hao Lei, and Meiqin Wang. 2024. Massive superpoly recovery with a meet-in-the-middle framework – Improved cube attacks on Trivium and Kreyvium. In *Advances in Cryptology – EUROCRYPT 2024 (Lecture Notes in Computer Science)*, M. Joye and G. Leander (Eds.). Springer. To appear (Preprint available at <https://ia.cr/2024/342>).
- [38] ISO/IEC 29192-3:2012. 2012. Information Technology – Security Techniques – Lightweight Cryptography, Part 3: Stream Ciphers. <https://www.iso.org/standard/56426.html>
- [39] Marc Joye. 2022. SoK: Fully Homomorphic encryption over the [discretized] torus. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2022, 4 (2022), 661–692. <https://doi.org/10.46586/TCHES.V2022.I4.661-692>
- [40] Marc Joye. 2024. TFHE public-key encryption revisited. In *Topics in Cryptology – CT-RSA 2024 (Lecture Notes in Computer Science)*, E. Oswald (Ed.). Springer. To appear (Preprint available at <https://ia.cr/2023/603>).
- [41] Jonathan Katz and Yehuda Lindell. 2020. *Introduction to Modern Cryptography* (3rd ed.). Chapman and Hall/CRC. <https://doi.org/10.1201/9781351133036>
- [42] Adeline Langlois and Damien Stehlé. 2015. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.* 75, 3 (2015), 565–599. <https://doi.org/10.1007/S10623-014-9938-4>
- [43] Yongwoo Lee, Daniele Micciancio, Andrey Kim, Rakyong Choi, Maxim Deryabin, Jieun Eom, and Donghoon Yoo. 2023. Efficient FHEW bootstrapping with small evaluation keys, and applications to threshold homomorphic encryption. In *Advances in Cryptology – EUROCRYPT 2023, Part III (Lecture Notes in Computer Science, Vol. 14006)*, C. Hazay and M. Stam (Eds.). Springer, 227–256. https://doi.org/10.1007/978-3-031-30620-4_8
- [44] Fukang Liu, Ravi Anand, Libo Wang, Willi Meier, and Takanori Isobe. 2023. Coefficient Grouping: Breaking Chaghri and More. In *Advances in Cryptology – EUROCRYPT 2023, Part IV (Lecture Notes in Computer Science, Vol. 14007)*, C. Hazay and M. Stam (Eds.). Springer, 287–317. https://doi.org/10.1007/978-3-031-30634-1_10
- [45] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. 2022. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In *Advances in Cryptology – ASIACRYPT 2022, Part II (Lecture Notes in Computer Science, Vol. 13792)*, S. Agrawal and D. Lin (Eds.). Springer, 130–160. https://doi.org/10.1007/978-3-031-22966-4_5
- [46] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. 2013. On ideal lattices and learning with errors over rings. *J. ACM* 60, 6 (2013), 43:1–43:35. <https://doi.org/10.1145/2535925>
- [47] Shihe Ma, Tairong Huang, Anyu Wang, Qixian Zhou, and Xiaoyun Wang. 2024. Fast and accurate: Efficient full-domain functional bootstrap and digit decomposition for homomorphic computation. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* 2024, 1 (2024), 592–616. <https://doi.org/10.46586/tches.v2024.i1.592-616>
- [48] Pierrick Méaux, Claude Carlet, Anthony Journault, and François-Xavier Standaert. 2019. Improved filter permutators for efficient FHE: Better instances and implementations. In *Progress in Cryptology – INDOCRYPT 2019 (Lecture Notes in Computer Science, Vol. 11898)*, F. Hao, S. Ruj, and S. Sen Gupta (Eds.). Springer, 68–91. https://doi.org/10.1007/978-3-030-35423-7_4
- [49] Pierrick Méaux, Anthony Journault, François-Xavier Standaert, and Claude Carlet. 2016. Towards stream ciphers for efficient FHE with low-noise ciphertexts. In *Advances in Cryptology – EUROCRYPT 2016, Part I (Lecture Notes in Computer Science, Vol. 9665)*, M. Fischlin and J.-S. Coron (Eds.). Springer, 311–343. https://doi.org/10.1007/978-3-662-49890-3_13
- [50] Parker Newton and Silas Richelson. 2023. A lower bound for proving hardness of learning with rounding with polynomial modulus. In *Advances in Cryptology – CRYPTO 2023, Part V (Lecture Notes in Computer Science, Vol. 14085)*, H. Handschuh and A. Lysyanskaya (Eds.). Springer, 805–835. https://doi.org/10.1007/978-3-031-38554-4_26
- [51] Oded Regev. 2009. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* 56, 6 (2009), 34:1–34:40. <https://doi.org/10.1145/1568318.1568324>

- [52] Daphné Trama, Pierre-Emmanuel Clet, Aymen Boudguiga, and Renaud Sirdey. 2023. A homomorphic AES evaluation in less than 30 seconds by means of TFHE. In *11th Workshop on Encrypted Computing & Applied Homomorphic Cryptography (WAHC 2023)*, M. Brenner, A. Costache, and K. Rohloff (Eds.). ACM Press, 79–90. <https://doi.org/10.1145/3605759.3625260>
- [53] Zama. 2022. TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data. <https://github.com/zama-ai/tfhe-rs>.

A Pseudorandom Functions

In this section, we recall the standard definition of pseudorandom function families.

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

$$\text{Adv}_Q^{\mathcal{D}, \text{prf}}(\lambda) := \left| \Pr[\mathcal{D}^{\text{PRF}(K, \cdot)}(1^\lambda) = 1] - \Pr[\mathcal{D}^{F(\cdot)}(1^\lambda) = 1] \right|,$$

where the probability is taken over the random choice of $K \leftarrow U(\{0, 1\}^\lambda)$ and $F \leftarrow U(\Omega)$ and the coin tosses of \mathcal{D} .

B GGSW Ciphertexts

Generalized GSW (GGSW) encryption is a natural extension of the scheme by Gentry, Sahai, and Waters [32] (in its ring version) to higher ranks. GGSW ciphertexts play a central role in the programmable bootstrapping of TFHE as they enable the external product of certain ciphertexts (as defined below). In particular, the bootstrapping keys are GGSW encryptions of private-key components.

The simplest way to view GGSW ciphertexts is through gadget decomposition of GLWE ciphertexts. Given a gadget vector $\mathbf{g} = (g_1, \dots, g_\ell) \in \mathcal{R}_q^\ell$ and a GLWE ciphertext $\mathbf{c} \leftarrow \text{GLWE}_\delta(\mu) \in \mathcal{R}_q^{k+1}$ under private key $\delta = (\delta_1, \dots, \delta_k) \in \mathcal{R}^k$, the corresponding gadget GLWE ciphertext (usually indicated with a $'$) is defined as

$$\text{GLWE}'_\delta(\mu) \leftarrow (\text{GLWE}_\delta(g_1 \cdot \mu), \dots, \text{GLWE}_\delta(g_\ell \cdot \mu)) .$$

This leveled encryption gives rise to a GGSW ciphertext; i.e.,

$$\text{GGSW}_\delta(\mu) \leftarrow (\text{GLWE}'_\delta(-\delta_1 \cdot \mu), \dots, \text{GLWE}'_\delta(-\delta_k \cdot \mu), \text{GLWE}'_\delta(\mu)) .$$

Importantly, the external product of a GLWE ciphertext and a GGSW ciphertext, denoted with \otimes , satisfies

$$\text{GLWE}_\delta(\mu_1) \otimes \text{GGSW}_\delta(\mu_2) = \text{GLWE}_\delta(\mu_1 \cdot \mu_2 + e_1 \cdot \mu_2)$$

where e_1 is the noise error present in $\text{GLWE}_\delta(\mu_1)$. The output of the external product is therefore a GLWE encryption of $\mu_1 \cdot \mu_2$ provided that message μ_2 is “small” so that $\|e_1 \cdot \mu_2\|_\infty \approx \|e_1\|_\infty$. This is the case for the TFHE bootstrapping keys, which are GGSW encryptions of key bits (i.e., values in $\{0, 1\}$).

C Evaluation Using Depth-3 Bootstrapping

Following Liu, Micciancio and Polyakov [45], we apply Theorem 2.6 with the following negacyclic functions $f_2: \mathbb{Z}_\Delta \rightarrow \mathbb{Z}_Q$, $f_0: \mathbb{Z}_{2\Delta} \rightarrow \mathbb{Z}_{2\Delta}$ and $f_1: \mathbb{Z}_{2\Delta} \rightarrow \mathbb{Z}_Q$, where we define

$\Delta = Q/p \in \mathbb{Z}$:

$$f_2(x) = \begin{cases} -\Delta/4 & \text{if } 0 \leq x < \Delta/4 \\ +\Delta/4 & \text{if } \Delta/2 \leq x < 3\Delta/4 \\ 0 & \text{otherwise} \end{cases}$$

$$f_0(x) = \left(\Delta \cdot \left\lfloor \frac{x}{\Delta} \right\rfloor - \frac{\Delta}{2} \right) \bmod 2\Delta$$

$$f_1(x) = \begin{cases} x & \text{if } x < \Delta \\ Q - x & \text{if } x \geq \Delta \end{cases}.$$

In the above function definitions, the value $f_1(x) = Q - x$ if $x \geq \Delta$ is not actually used and only defined to satisfy the negacyclic constraint. Note that although TFHE natively handles negacyclic functions $f: \mathbb{Z}_Q \rightarrow \mathbb{Z}_q$ for $Q = 2N$, we can adapt it to handle negacyclic functions $f': \mathbb{Z}_{Q'} \rightarrow \mathbb{Z}_{q'}$ where Q' divides $2N$ and q' divides q . Essentially, from f' , we may derive a negacyclic function $f: \mathbb{Z}_{2N} \rightarrow \mathbb{Z}_q$ by setting $f(x) = (q/q') \cdot f'(\lfloor x \cdot 2N/Q' \rfloor)$ where the multiplication is performed over the integers.

The homomorphic evaluation algorithm is based on the technique of [45, Section 4], which is adapted using their homomorphic floor function [45, Algorithm 2] that allows handling an arbitrarily large noise in the input ciphertext.

In the description hereunder, we assume that the seed $s \in \mathbb{Z}^n$ of the LWE-based PRF is encrypted in the same way as the bootstrapping key of an LWE-based encryption scheme where the secret key is s . We thus assume a bootstrapping key $\text{bsk} = \text{GGSW}_{s'}(s)$ consisting of a generalized-GSW encryption of the seed s (viewed as the decryption key of a LWE-based encryption scheme) under a GGSW secret key s' .

$\text{Eval}_{\text{pp}}(\text{bsk}, x)$: Given public parameters $\text{pp} = (n, Q, p, \beta)$, an evaluation key bsk and an input $x \in \{0, 1\}^\ell$, compute the input-dependent vector $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$ and do the following:

1. $(\bar{\mathbf{a}}, 0) := (-\mathbf{a}, 0) \bmod \Delta$
2. $(\mathbf{c}, d) := (-\mathbf{a}, 0) - \text{Boot}[f_2](\bar{\mathbf{a}}, 0) \pmod{Q}$
3. $d := d + \beta - \frac{\Delta}{4}$
4. $(\bar{\mathbf{c}}, \bar{d}) := (\mathbf{c}, d) \bmod \Delta$
5. $(\bar{\mathbf{c}}, \bar{d}) := (\bar{\mathbf{c}}, \bar{d}) \bmod 2\Delta$
6. $(\bar{\mathbf{c}}, \bar{d}) := (\bar{\mathbf{c}}, \bar{d}) - \text{Boot}[f_0](\bar{\mathbf{c}}, \bar{d}) \pmod{2\Delta}$
7. $\bar{d} := \bar{d} + \beta - \frac{\Delta}{2}$
8. $(\mathbf{c}, d) = (\bar{\mathbf{c}}, \bar{d}) - \text{Boot}[f_1](\bar{\mathbf{c}}, \bar{d}) + \beta \pmod{Q}$

Output the ciphertext $\text{ct} = (\mathbf{c}, d) \in \mathbb{Z}_Q^{n+1}$.

At the end of the evaluation process, ct is an LWE encryption of the PRF evaluation $\text{PRF}_s(x) = \left\lceil \frac{p}{Q} \cdot (\langle H(x), \mathbf{s} \rangle \bmod Q) \right\rceil \bmod p$ under an LWE secret key which is the PRF secret key \mathbf{s} itself. In order to obtain an LWE encryption of $\text{PRF}_s(x)$ under the LWE secret key \mathbf{s}' (which is the one used to encrypt \mathbf{s} in the bootstrapping key), a simple solution is to remove the key-switching step in the last call to the Boot procedure in Line 8.

We now prove the correctness property.

LEMMA C.1. *Assume that p divides Q and that $\Delta = Q/p > 16\beta$, where β is the bootstrapping error from Theorem 2.6. For any $x \in \{0, 1\}^\ell$, Eval outputs a ciphertext $(\mathbf{c}, d) \in \mathbb{Z}_Q^{n+1}$ such that*

$\text{Dec}_s(\mathbf{c}, d) = \Delta \cdot \mu + e \pmod{Q}$, where $|e| < 2\beta$,

$$\mu = \left\lfloor \frac{p}{Q} \cdot (\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) \right\rfloor \bmod p$$

with $\mathbf{a} = H(x) \in \mathbb{Z}_Q^n$.

PROOF. Since

$$(-\mathbf{a}, 0) \in \mathbb{Z}_Q^{n+1} = \left(-\mathbf{a}, -(\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q) + \Delta \cdot \left\lceil \frac{\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q}{\Delta} \right\rceil + \underbrace{(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)}_{\in [-\Delta/2, \Delta/2]} \right)$$

where $\Delta = Q/p$, we have

$$\begin{aligned} \text{Dec}_s(-\mathbf{a}, 0) &= \langle \mathbf{a}, \mathbf{s} \rangle \bmod Q \\ &= \Delta \cdot \underbrace{\left\lceil \frac{\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q}{\Delta} \right\rceil}_{\triangleq \mu} + (\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta) \pmod{Q} \end{aligned}$$

before Line 1. After Line 1, we have $(\bar{\mathbf{a}}, 0) = (-\mathbf{a}, 0) \bmod \Delta$. By the same arguments as in [45, Lemma 3], we have that

$$m' \triangleq \text{Dec}_s(\bar{\mathbf{a}}, 0) \bmod \Delta = \langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta$$

are the $\log_2 \Delta$ least significant bits of $\langle \mathbf{a}, \mathbf{s} \rangle \bmod Q$.

Then, $f_2(m') = f_2(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta) \in \{-\Delta/4, 0, \Delta/4\}$. Also, the action of f_2 only depends on the two most significant bits of the input and maps $00 \rightarrow 11$, $10 \rightarrow 01$, and $01, 11 \rightarrow 00$ (so that the second MSB is always flipped). By Theorem 2.6, at Line 2, we then have

$$\begin{aligned} \text{Dec}_s(\mathbf{c}, d) &= \text{Dec}_s(-\mathbf{a}, 0) - \text{Dec}_s(\text{Boot}[f_2](\bar{\mathbf{a}}, 0)) \bmod Q \\ &= \text{Dec}_s(-\mathbf{a}, 0) - f_2(m') + \underbrace{e_\beta}_{\in (-\beta, \beta)} \bmod Q \\ &= \Delta \cdot \mu + \underbrace{(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta) - f_2(\langle \mathbf{a}, \mathbf{s} \rangle \bmod \Delta)}_{\triangleq \mu'} + e_\beta \pmod{Q} \end{aligned}$$

where the second MSB of μ' is always 1 and subtracting $\Delta/4$ at Line 3 always sets this bit to 0. At the end of Line 3, we thus have

$$\text{Dec}_s(\mathbf{c}, d) = \mu \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e \pmod{Q} \quad (16)$$

for some $\tilde{b} \in \{0, 1\}$, $x \in [0, \Delta/4)$ and $e \in [0, 2\beta)$. At Line 4, this yields

$$\text{Dec}_s(\bar{\mathbf{c}}, \bar{d}) = \tilde{b} \cdot (\Delta/2) + x + e \pmod{\Delta} .$$

At Line 5, it becomes

$$\text{Dec}_s(\bar{\mathbf{c}}, \bar{d}) = \theta \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e \pmod{2\Delta}$$

for some $\theta \in \{0, 1\}$. Before starting Line 6, we thus have

$$\begin{aligned}
& \text{Dec}_s(\bar{c}, \bar{d}) - \text{Dec}_s(\text{Boot}[f_0](\bar{c}, \bar{d})) \\
& \equiv \text{Dec}_s(\bar{c}, \bar{d}) - f_0(\text{Dec}_s(\bar{c}, \bar{d})) + e'_\beta \\
& \equiv (\theta \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e) \\
& \quad + f_0(\theta \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e) + e'_\beta \\
& \equiv (\theta \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e) \\
& \quad - \left(\Delta \cdot \left\lfloor \frac{\tilde{b} \cdot (\Delta/2) + x + e}{\Delta} \right\rfloor + \theta \cdot \Delta \right) + \frac{\Delta}{2} + e'_\beta \\
& \equiv (\tilde{b} \cdot (\Delta/2) + x + e) \\
& \quad - \left(\Delta \cdot \left\lfloor \frac{\tilde{b} \cdot (\Delta/2) + x + e}{\Delta} \right\rfloor \right) + \frac{\Delta}{2} + e'_\beta \\
& \equiv (\tilde{b} \cdot (\Delta/2) + x + e) + \frac{\Delta}{2} + e'_\beta \pmod{2\Delta}
\end{aligned}$$

where $e'_\beta \in (-\beta, \beta)$ and the last equality holds because

$$\left\lfloor \frac{\tilde{b} \cdot (\Delta/2) + x + e}{\Delta} \right\rfloor = 0$$

since $x \in [0, \Delta/4)$ and $e \in [0, 2\beta]$ with $2\beta < \Delta/4$ (due to the hypothesis $\Delta > 16\beta$). At the end of Line 7, we then have

$$\text{Dec}_s(\bar{c}, \bar{d}) \equiv \tilde{b} \cdot (\Delta/2) + x + e + e' \pmod{2\Delta}$$

with $e \in [0, 2\beta)$ and $e' \triangleq e'_\beta + \beta \in (0, 2\beta)$. Since $x \in [0, \Delta/4)$ and $4\beta < \Delta/4$ (due to the hypothesis $\Delta > 16\beta$), we then have $(\text{Dec}_s(\bar{c}, \bar{d}) \bmod 2\Delta) \in [0, \Delta/2)$ if $\tilde{b} = 0$ and $(\text{Dec}_s(\bar{c}, \bar{d}) \bmod 2\Delta) \in [\Delta/2, \Delta)$ if $\tilde{b} = 1$.

Therefore $f_1(\text{Dec}_s(\bar{c}, \bar{d}) \bmod 2\Delta) = \tilde{b} \cdot (\Delta/2) + x + e + e'$ by the definition of $f_1(x)$. By combining this with (16) and applying Theorem 2.6 again, we obtain

$$\begin{aligned}
& \text{Dec}_s(c, d) - \text{Dec}_s(\text{Boot}[f_1](\bar{c}, \bar{d})) \\
& \equiv (\mu \cdot \Delta + \tilde{b} \cdot (\Delta/2) + x + e) - (\tilde{b} \cdot (\Delta/2) + x + e + e') + e''_\beta \\
& \equiv \mu \cdot \Delta - e' + e''_\beta \pmod{Q}
\end{aligned}$$

with $e''_\beta \in (-\beta, \beta)$. Since $-e' + e''_\beta \in (-3\beta, \beta)$, after Line 8, we obtain

$$\text{Dec}_s(c, d) = \mu \cdot \Delta + \bar{e} \pmod{Q}$$

for some $\bar{e} \in (-2\beta, 2\beta)$, as claimed. \square