



Circuit Privacy for FHEW/TFHE-Style Fully Homomorphic Encryption in Practice

Kamil Kluczniak^a

Independent Researcher, Munich, Germany

Abstract. A fully homomorphic encryption (FHE) scheme allows a client to encrypt and delegate its data to a server that performs computation on the encrypted data that the client can then decrypt. While FHE gives confidentiality to clients' data, it does not protect the server's input and computation. Nevertheless, FHE schemes are still helpful in building delegation protocols that reduce communication complexity, as the ciphertext's size is independent of the size of the computation performed on them.

We can further extend FHE by a property called circuit privacy, which guarantees that the result of computing on ciphertexts reveals no information on the computed function and the inputs of the server. Thereby, circuit private FHE gives rise to round optimal and communication efficient secure two-party computation protocols. In this work, we design a randomized FHEW/TFHE-style bootstrapping algorithm whose single invocation sanitizes a ciphertext and, consequently, serves as a tool to provide circuit privacy. We give an extensive analysis, propose parameters, and provide a C++ implementation of our scheme. Our bootstrapping can sanitize a ciphertext to achieve circuit privacy at an 80-bit statistical security level in between 1.3 and 0.9 seconds, depending which Gaussian sampling algorithm is used, and whether the parameter set targets a fast Fourier or a number theoretic transform-based implementation. In addition, we can perform non-sanitized bootstrapping in around 0.27 or 0.14 seconds. Crucially, we do not need to increase the parameters to perform computation before or after sanitization takes place. For comparison's sake, we revisit the Ducas-Stehlé washing machine method. In particular, we give a tight analysis, estimate efficiency, review old, and provide new parameters.

Keywords: Fully Homomorphic Encryption · Circuit Privacy

1 Introduction

Fully homomorphic encryption (FHE) is an encryption scheme that allows performing arbitrary computation on encrypted data. A client encrypts a message m and sends the ciphertext to a server which, given a function F , returns another ciphertext that decrypts to $F(m)$. The concept of FHE was first introduced by Rivest, Adleman and Dertouzos [RAD78], and the first theoretical realization of that concept is due to Gentry [Gen09b].

A critical property for FHE is circuit privacy (also called function privacy). Roughly speaking, the ciphertext that is the product of the server computing a function F on encrypted data should not reveal any information about F except that the ciphertext decrypts to $F(m)$. To prove circuit privacy, we need to show a simulator that, on input $F(m)$ and a public key, outputs a fresh encryption of $F(m)$, which is indistinguishable

E-mail: kamil.kluczniak@gmail.com (Kamil Kluczniak)

^aAuthor was partially supported by the German Ministry for Education and Research through funding for the project CISP-Stanford Center for Cybersecurity (Funding number: 16KIS0927).



from the servers' computed ciphertext. In particular, the distribution of an evaluated ciphertext should be close to or the same as the distribution of a fresh encryption.

We can easily see that circuit private FHE gives us semi-honest two-party computation with optimal communication. Namely, we only need one round of communication. The first message can be reused, and the communication complexity is independent of the size of the computation. Furthermore, we can reuse the ciphertexts output from the evaluation process and keep computing on them. Since circuit private FHE gives us two-party computation, all applications for two-party computation protocols apply here as well. Among other these are private set intersection [HFH99, Mea86, CLR17], oblivious pseudorandom functions [BIP⁺18, ADDG24] neural network inference [DGBL⁺16, CdWM⁺17, LJLA17, JKLS18, JVC18, BGGJ20, ABSdV19, CDKS19, RSC⁺19, BGPG20] or analysis on genomic data [KSK⁺18, KSK⁺20, BGPG20]. Recently, Akavia, Gentry, Halevi, and Vald [AGHV22] showed that circuit private IND-CPA secure homomorphic encryption satisfies a relaxed notion of CCA2 security. Note that circuit privacy is not always needed. Without circuit privacy FHE reduces to secure delegation. For example, in (single-server) private information retrieval we are only interested in protecting the user's query, but not in the confidentiality of a potentially large database of the server. On the other hand, we believe that for neural network inference, as an example, confidentiality of the neural network is essential. In contrast to PIR, it is difficult to make an argument for compressing the communication, as current FHE schemes require sending public keys and ciphertexts that are an order of magnitude larger than the size of deep neural networks.

Despite over a decade of advances in constructing scalable fully homomorphic encryption schemes [GH11, BV11, BGV12, GHS12, AP13, GSW13, BV14, AP14, HS15, DM15, CGGI16a, CH18, CGGI20, HS21], and numerous implementations [PAL21, CGGI16b, CJL⁺20, Lat22] there are few constructions and implementation that we are aware of that natively provide circuit privacy.

Current approaches to Circuit Privacy. In this paper, we are interested in fully homomorphic encryption as in [Gen09b]. Namely, ciphertexts do not grow with the size of computation, and evaluation results are reusable. A trivial way to re-randomize a ciphertext is to create a fresh encryption of zero using the public key and add it to the ciphertext resulting from the computation. Unfortunately, such an approach is insufficient to provide circuit privacy in current FHE schemes because all secure FHE schemes we know are based on noisy encryptions. This noise may depend on the computed circuit and is the main obstacle to overcome when re-randomizing (or sanitizing) a ciphertext to provide circuit privacy. Below we summarize current approaches.

Noise Flooding: The technique requires adding a fresh ciphertext of zero and an exponentially larger noise term to the sanitized ciphertext. Unfortunately, in practice, this additional noise term is substantially large and requires us to choose very big parameters. We note that it is required to take the noise exponentially larger than the noise in the sanitized ciphertext due to the smudging lemma [AJL⁺12]. Hence, if the noise in this ciphertext is already large due to some previous computation, then the the magnitude of the additional noise must be chosen accordingly. Nevertheless, the method has found some applications in leveled homomorphic computation [CLR17] where we do not bootstrap the ciphertexts and can tolerate larger parameters.

Ducas-Stehlé Washing Machine: Introduced by Ducas and Stehlé [DS16], requires to run a sequence of re-randomization steps (flooding cycles), each with a smaller noise flooding error, followed by invocations of a bootstrapping algorithm. The paper only roughly estimates the number of times re-randomization and bootstrapping must be invoked. However, as the authors admit, the estimates should be taken with great caution

and defer a concrete analysis to future work. For example, they suggest running the FHEW [DM15] bootstrapping algorithm between 8 and 16 times. It is not entirely clear what security level they are able to achieve and whether the parameter set of the FHEW algorithm proposed at the time satisfies the given correctness constraints. To the best of our knowledge, no concrete analysis or implementations have been done so far.

Secure Two-Party Computation: A few works [OPP14, GHV10, CO17] proposed to use garbled circuit-based techniques to provide circuit privacy. For example, Gentry, Halevi, Vaikuntanathan [GHV10] give a non-compact homomorphic encryption scheme that can be thought of as a re-randomizable version of garbled circuits, which can be instantiated from the DDH assumption. [GHV10] shows how to make the scheme compact, by performing the evaluation over a compact FHE scheme and using the circuit-private scheme to re-encrypt the ciphertext. Ostrovsky, Paskin-Cherniavsky and Paskin-Cherniavsky [OPP14], and Chongchitmate and Ostrovsky [CO17] use the same technique as [GHV10] to build a compact FHE by sanitizing the ciphertext using a non-compact 2PC scheme, but to achieve malicious security, they prove the well-formedness of the ciphertext using zero-knowledge proofs.

Rerandomizing Computation: Bourse, del Pino, Minelli and Wee [BdPMW16] exploits properties of the Gentry, Sahai, Waters (GSW) cryptosystem [GSW13] to build a circuit private homomorphic encryption scheme. Specifically, when multiplying GSW ciphertexts, they use a randomized version of gadget decomposition instead of a deterministic one. In [BdPMW16] the authors show that when gadget decomposition is implemented via Gaussian sampling [MP12, GM18] with appropriate parameters, then we can build an FHE scheme for circuits in NC^1 (circuit of depth logarithmic in the number of inputs). The results are asymptotic, without concrete parameter proposals nor implementation.

1.1 Our Contributions

We design a randomized FHEW/TFHE-style [DM15, CGGI16a] bootstrapping algorithm that can sanitize a given ciphertext. In contrast to the Ducas-Stehlé washing machine method [DS16], which we shortly refer to as DS-WM, we need to run our bootstrapping algorithm only once. Our results solve an open problem posted in [BdPMW16], in that we use their randomization concept in an FHEW-style bootstrapping scheme. We note that porting the ideas from [BdPMW16] to the ring setting is non-trivial since there is no analogue of the Gaussian leftover hash lemma [AGHS13, AR13, BdPMW16] for the ring setting. Moreover, we can argue that designing a “leaky” analogue of the regularity lemmas [Mic02, SS11, LPR13] as in [DSGKS21], may result in practically inefficient bootstrapping. While we use some techniques from [BdPMW16], the overall method departs from [BdPMW16]. In this work we show how to bypass the need to port [BdPMW16] into the ring setting using simple techniques in the right places by exploiting structural properties of FHEW/TFHE [DM15, CGGI16a] instantiated over the ring $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$ where N is a power of two. Along the way, we generalize the technical lemmas from [BdPMW16] to support any modulus $Q \in \mathbb{N}$, instead of moduli of the form $Q = L^\ell$ for some $L, \ell \in \mathbb{N}$.

We compare our method with DS-WM that is the most competitive. While [DS16] gives a heuristic instantiation based on FHEW [DM15], they left a serious analysis as an open problem. We resolve the problem and give a tight error analysis, and provide scripts that automate noise and security estimations of our randomized bootstrapping and DS-WM. We show that the parameters proposed in [DS16] cannot be circuit private due to correctness issues. In general, we show that instantiations over a ring of dimension 2^{10} , or smaller, cannot give more than 30-bits of statistical security. Note that many

efficient bootstrapping schemes [DM15, CGGI16a, CGGI20] are instantiated over rings of dimension 2^{10} .

Finally, we give an efficient C++ implementation¹ of our bootstrapping algorithms. To the best of our knowledge, this is the first practical realization of a circuit private FHEW/TFHE-style FHE scheme. Our implementation supports number theoretic transform-based (NTT) and fast Fourier transform-based (FFT) multiplication of ring elements. Due to our versatile implementation, we can experiment with different moduli choices that lead to different algorithm variants. In particular, we can instantiate different Gaussian samplers that are optimized toward a specific choice of modulus. We choose different parameters targeting the different representations. Nevertheless, we identify some drawbacks to the FFT-based implementation due to the relatively low precision of the floating point arithmetic. Our algorithm sanitizes a ciphertext in about 1.3 seconds for both the NTT and FFT-based implementations when using Karney’s Gaussian sampling algorithm [Kar16]. When sampling from the rounded continuous Gaussian distribution via Box-Muller transform [BM58] we can reduce the times to 1.0 and 0.9 seconds for NTT and FFT, respectively. Furthermore, without increasing the size of the public key, we can compute standard FHEW/TFHE-style bootstrapping in around 0.14 and 0.27 seconds for NTT and FFT, respectively. We show that the DS-WM method is between $1.56\times$ to $7.88\times$ slower than ours, depending on the parameter sets and the implementation of the Gaussian sampling algorithm. We compared parameters with the same key sizes, non-sanitized (deterministic) computation time, and correctness level. We also stress that Gaussian sampling in our method constitutes about 78% and 40% of the entire computation. Our implementation is linear and does not take advantage of special vector instructions or parallelism. Based on the speedups when using different Gaussian samplers, we believe that an optimized implementation could significantly improve the execution times, while we do not see much room for improvement in DS-WM as it simply repeats the base bootstrapping algorithm.

Note, that our and [DS16] methods follow the theoretical blueprint from [GHV10] in that we compute on a non-circuit-private FHE, but run homomorphic decryption with a circuit-private scheme before returning the ciphertext. The difference is that [GHV10, OPP14, CO17] run garbled circuits that introduce new assumptions, and rounds and increase the output size in practice since we need to send wire encodings and oblivious transfer labels for each gate in the decryption circuit. To the best of our knowledge [GHV10, OPP14, CO17] was never implemented. In our scheme and [DS16], the output is an Learning with Errors (LWE) ciphertext, just like in a non-circuit-private FHE. Finally, note that we can achieve malicious circuit-privacy as in [OPP14, CO17] by proving well-formedness of the public key and ciphertexts, but efficient realization of such proofs is currently under extensive research [LNP22, BN20, BCGZ24, Bot24] and out of scope of this paper.

1.2 Our Techniques

In this section, we first give a high-level overview of FHEW/TFHE-style bootstrapping and our circuit private version. Then we discuss the technical problems to realize the idea and our solutions.

FHEW/TFHE-Style Bootstrapping. First, let us recall the symmetric key version of Regev’s encryption [Reg09]. The encryption algorithm chooses a uniform random vector $\mathbf{a} \in \mathbb{Z}_q^n$ and a secret key $\mathbf{s} \in \mathbb{Z}_q^n$, and computes an encryption of a message $m \in \mathbb{Z}_t$ as $[b, \mathbf{a}] \in \mathbb{Z}_q^{n+1}$, where $b = \langle \mathbf{a}, \mathbf{s} \rangle + \tilde{m} + e \pmod{q}$, $\tilde{m} = \frac{q}{t} \cdot m$, and $e < \frac{q}{2t}$ is a small discrete

¹Available at at Github <https://github.com/FHE-Deck>.

Gaussian noise. For simplicity, we assume that $t \mid q$. The decryption algorithm calculates $\left\lfloor \frac{t}{q} \cdot (b - \langle \mathbf{a}, \mathbf{s} \rangle) \right\rfloor = \left\lfloor \frac{t}{q} \cdot \left(\frac{q}{t} \cdot m + e \right) \right\rfloor = m$.

The scheme can also be instantiated over the cyclotomic ring \mathcal{R}_Q , where $\mathcal{R} = \mathbb{Z}[X]/(X^N + 1)$ and $\mathcal{R}_Q = \mathcal{R}/Q\mathcal{R}$. To encrypt a message \mathbf{m} from the ring \mathcal{R}_t , the scheme selects a uniform random $\mathbf{a} \in \mathcal{R}_Q$ and a secret key $\mathbf{s} \in \mathcal{R}_Q$ and computes the encrypted message $[\mathbf{b}, \mathbf{a}]$ where $\mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \frac{q}{t} \cdot \mathbf{m} + \boldsymbol{\epsilon}$, and $\boldsymbol{\epsilon}$ is a small error term in \mathcal{R}_Q with coefficients from the discrete Gaussian Distribution.

Now let us proceed to the ideas underlying the FHEW-style bootstrapping scheme introduced by Ducas and Micciancio in [DM15]. We want to re-encrypt an LWE ciphertext $[b, \mathbf{a}] \in \mathbb{Z}_q^{n+1}$. Assuming that the LWE modulus is $q = 2 \cdot N$, the scheme sets up a homomorphic accumulator acc as an Ring LWE (RLWE) encryption of $\mathbf{a}_{\text{rot}} \cdot X^b \in \mathcal{R}_Q$. We refer to [DM15] on how to choose \mathbf{a}_{rot} .

Next, the scheme multiplies acc with encryptions of $X^{-\mathbf{a}^{[i]} \cdot \mathbf{s}^{[i]}} \in \mathcal{R}_Q$ for each $i \in [1, n]$. Finally, the message in the accumulator will be:

$$\mathbf{a}_{\text{rot}} \cdot X^{b - \sum_{i=1}^n \mathbf{a}^{[i]} \cdot \mathbf{s}^{[i]}} = \mathbf{a}_{\text{rot}} \cdot X^{k \cdot q + \tilde{m} + e} = \mathbf{a}_{\text{rot}} \cdot X^{\tilde{m} + e \bmod 2 \cdot N} \in \mathcal{R}_Q.$$

The last step is to extract an LWE encryption of the constant term from the rotated accumulator. Denote this LWE ciphertext as $[\mathbf{a}_{\text{out}}, b_{\text{out}}]$, where $b = \langle \mathbf{a}_{\text{out}}, \mathbf{s} \rangle + m_{\text{out}} + e_{\text{out}}$.

How to Get Circuit Privacy. Note that the output ciphertext is entirely determined by the input ciphertext and the evaluation key. One solution is to use the noise flooding technique [Gen09a]. Essentially, the technique involves adding a fresh ciphertext and some noise to the output ciphertext. Although this makes the noise in the output ciphertext uniform in an interval and independent, the magnitude of the noise term is exponential in the security parameter. This means that a large modulus (over 110 bits) must be chosen and the dimension must be increased to compensate for the security loss, resulting in a large evaluation key and reduced efficiency. We give sample estimates for this technique in Appendix C.

Another solution by Ducas and Stehlé [DS16] is to apply smaller flooding noise and repeat the bootstrapping process $O(\lambda)$ times to achieve the desired level of circuit privacy security. However, this approach has the immediate downside of having to repeat the expensive bootstrapping operation multiple times. Nevertheless, in this paper we show new parameters and optimizations of [DS16] applied to FHEW/TFHE-style bootstrapping, in order to give better insights into the state of the art and a proper comparison with our method.

Our idea is to re-randomize the blind rotation algorithm so that the distribution of the extracted LWE ciphertext will already be independent of the input ciphertext. To do this, we first need to construct a randomized algorithm to multiply RLWE ciphertexts. Our starting point is the algorithm introduced by Bourse et al. in [BdPMW16], who showed a randomized product of GSW [GSW13] ciphertexts. It turns out, however, that there are multiple problems that we need to overcome to apply this high level idea for the bootstrapping algorithm.

Brief Overview of the Randomized GSW Product. First we define a gadget vector $\mathbf{g} = [1, 2, \dots, 2^\ell]$ and the matrix $\mathbf{G} = \mathbf{g} \otimes \mathbf{I}_n \in \mathbb{Z}_Q^{n \times \ell n}$, where \mathbf{I}_n is the n dimensional identity matrix. A GSW encryption of a message $m \in \mathbb{Z}_Q$ is given as

$$\mathbf{C} = \begin{bmatrix} \mathbf{A} \\ \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top \end{bmatrix} + m \cdot \mathbf{G},$$

where $\mathbf{s} \in \mathbb{Z}_Q^n$ is the secret key, $\mathbf{A} \in \mathbb{Z}_Q^{(n-1) \times \ell n}$ is public, and $\mathbf{e} \in \mathbb{Z}_Q^{\ell n}$ is a noise vector whose entries are from the discrete Gaussian distribution.

Now we are ready to recall the method from [BdPMW16]. Define a randomized gadget decomposition algorithm $\mathbf{X} \leftarrow \mathbf{G}_{\text{rand}}^{-1}(a \cdot \mathbf{G})$, where $a \in \mathbb{Z}_Q$ and the matrix $\mathbf{X} \in \mathbb{Z}_Q^{\ell n \times \ell n}$ has entries from the discrete Gaussian distribution such that $\mathbf{G} \cdot \mathbf{X} = a \cdot \mathbf{G}$. We can use $\mathbf{G}_{\text{rand}}^{-1}$ to multiply the GSW ciphertext $\mathbf{C} \in \mathbb{Z}_Q^{n \times \ell n}$ by a and randomize the outcome as follows.

$$\mathbf{C} \cdot \mathbf{X} + \begin{bmatrix} \mathbf{0} \\ \mathbf{y}^\top \end{bmatrix} = \begin{bmatrix} \mathbf{A} \cdot \mathbf{X} \\ \mathbf{s}^\top \mathbf{A} \cdot \mathbf{X} + \mathbf{e}^\top \mathbf{X} + \mathbf{y}^\top \end{bmatrix} + m \cdot a \cdot \mathbf{G},$$

where $\mathbf{y} \in \mathbb{Z}_Q^{\ell n}$ is chosen from the discrete Gaussian distribution. The main idea and technical contribution in [BdPMW16] is to show that such a product already gives us a ciphertext that is statistically independent of the input ciphertext. At the heart of their proof is the core randomization lemma that states that a tuple $(\mathbf{A} \cdot \mathbf{X}, \mathbf{e}^\top \mathbf{X} + \mathbf{y}^\top)$ is statistically indistinguishable from $(\bar{\mathbf{A}}, \bar{\mathbf{e}}^\top)$, where $\bar{\mathbf{A}} \in \mathbb{Z}_Q^{(n-1) \times \ell n}$ is from the uniform distribution and $\bar{\mathbf{e}} \in \mathbb{Z}_Q^{\ell n}$ is an independent random variable from the discrete Gaussian distribution with a slightly higher standard deviation, given that \mathbf{X} and \mathbf{y} have sufficiently high standard deviations.

The first step to prove the core randomization lemma is to show that $\mathbf{A} \cdot \mathbf{X}$ is close to uniform from the generalized leftover hash lemma [DRS04]. To this end, we need to analyze the entropy of \mathbf{X} given $\mathbf{e}^\top \mathbf{X} + \mathbf{y}^\top$, and \mathbf{e} . To show that $\mathbf{e}^\top \mathbf{X} + \mathbf{y}^\top$ is close to an independent discrete Gaussian random variable, [BdPMW16] use an adaptation of the Gaussian leftover hash lemma [AGHS13, AR13].

In FHEW/TFHE-style bootstrapping, we need to perform external products of ring LWE ciphertexts. Hence the immediate problem is to translate the randomization technique into the ring setting. But as we discuss later, such translation may still be impractical.

Problems with Translating [BdPMW16] Into the Ring Setting. At the heart of FHEW/TFHE-style bootstrapping algorithms is an algorithm called blind rotation that outputs an RLWE ciphertext $\mathbf{c} = (\mathbf{a}, \mathbf{b}) \in \mathcal{R}_Q^2$ whose constant coefficient of the encrypted message encodes the decryption of an input ciphertext. We can show that the ciphertext can be represented as $(\mathbf{a}^\top \mathbf{x}, \mathbf{e}^\top \mathbf{x} + \eta)$ as above but where $\eta \in \mathcal{R}_Q$, $\mathbf{a}, \mathbf{x}, \mathbf{e} \in \mathcal{R}_Q^m$ and elements in \mathbf{x} have coefficients from the discrete Gaussian distribution. Recall that $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$.

If we want to follow the technique from [BdPMW16], we would need to show that $\mathbf{a} = \mathbf{a}^\top \mathbf{x}$ is close to uniform given $\mathbf{e}^\top \mathbf{x} + \eta$ and \mathbf{e} . We may try to define a “leaky” version of the regularity lemma [LPR13] as in [DSGKS21]. But we argue that even if we would ignore the security loss due to the leak, the regularity lemma from [LPR13] requires us to choose a small decomposition basis resulting in high ℓ and, consequently, in relatively slow (Ring) GSW products. Concretely, we must choose the standard deviation σ_x of \mathbf{x} to be larger than $N \cdot Q^{1/\ell+2/N\ell}$ to achieve a negligibly small statistical distance. Naturally, σ_x cannot be too large to guarantee correctness. Thus, ℓ must be sufficiently high to minimize $Q^{1/\ell+2/N\ell}$. For example, let us fix $\sigma_x = N$, where $N = 2^{11}$, which is a typical ring dimension used in practice. In this case, we must choose ℓ such that $Q^{1/\ell+2/N\ell} < 1$. Typically, Q is a 32- to 54-bit number. This implies that $\ell > 32$. The ℓ parameter is critical as it affects the most time-consuming operation in the bootstrapping scheme. Hence it is imperative to keep ℓ small in practical implementations. Another problem is that we do not have a ring analogue of the Gaussian leftover hash lemma.

Our Solution. To bypass these problems, we exploit that in FHEW/TFHE-style schemes, we extract an LWE ciphertext from \mathbf{c} . In particular, observe that $(\mathbf{a}', b') \in \mathbb{Z}_Q^{N+1}$, where $b' = \mathbf{b}[1] \in \mathbb{Z}_Q$ is \mathbf{b} 's constant coefficient, and $\mathbf{a}' \in \mathbb{Z}_Q^N$ is such that $\mathbf{a}'[1] = \mathbf{a}[1]$ and $\mathbf{a}'[i] = \mathbf{a}[N-i]$ for $i = 0 \dots N-2$, is a correct LWE ciphertext with respect to the secret key $\mathbf{s}' = \mathbf{s}$ (the coefficient vector of \mathbf{s}) encrypting the constant coefficient of \mathbf{c} 's

message. Note that we still cannot claim that \mathbf{a}' is close to uniform, but what we can do is sample a fresh LWE ciphertext of 0 denoted as $(\mathbf{a}_{\text{rand}}, b_{\text{rand}})$, add it to (\mathbf{a}', b') obtaining a ciphertext $(\bar{\mathbf{a}}, \bar{b})$ of the same message where $\bar{\mathbf{a}}$ is statistically close to uniform. Note that if \mathbf{a}' is fixed and \mathbf{a}_{rand} is from the uniform distribution, then $\mathbf{a}' + \mathbf{a}_{\text{rand}} = \bar{\mathbf{a}}$ is a fixed shift of uniformly distributed random variables. Finally, we can show that the error term of (\mathbf{a}', b') is already in the form required by the Gaussian leftover hash lemma [AGHS13, AR13, BdPMW16] (see Lemma 15). To show this, we exploit the specific structure of the ring $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$ where the product of two ring elements is a negacyclic convolution of two polynomials. Our analysis applies only to \mathcal{R}_Q which is used in FHEW/TFHE-style bootstrapping schemes. To summarize, our construction completely bypasses the need to adapt [BdPMW16] to the ring setting.

Concurrent Work. Concurrently and independently, Bourse and Izabachène [BI22] gave a circuit private FHEW/TFHE-style algorithm. However, the techniques and practical efficiency differ significantly. Roughly speaking, [BI22] build a multiplication algorithm between RLWE and RGSW ciphertexts that outputs an RLWE ciphertext of the product, which is statistically close to a “fresh” ciphertext. To randomize the bootstrapping algorithm, [BI22] needs to publish a sanitization key that consists of $\geq 2^{17}$ RLWE ciphertexts alongside the bootstrapping key. Such key requires over 677 MB memory². In contrast, our algorithm requires less than 186 or 69 MB (depending on the parameter set) of additional sanitization key. Furthermore, [BI22] needs to sample a “fresh” RLWE ciphertext for every external product in the blind rotation step. As reported in [BI22] sampling a single RLWE ciphertext takes approximately 45 seconds, and they need to sample 612 to perform a single bootstrapping operation. That gives us 7.6 hours to perform the sampling. To overcome the timing issue, the authors assume that the RLWE ciphertexts and Gaussian are precomputed, and the machine has unrestricted memory and precomputation. However, to make the computation feasible on a laptop they test the algorithm by reusing a single RLWE ciphertext. Even with the precomputed values, deterministic bootstrapping takes 3.15 seconds, and sanitization bootstrapping takes between 21 and 4.68 seconds depending on how many Gaussian samples were precomputed. In comparison, our deterministic bootstrapping takes 0.14 or 0.27 seconds, sanitization takes between 0.9 and 1.3 seconds. Most importantly, we do not need any precomputation for our algorithms to be efficient, and all random variables are generated on-the-fly.

2 Background and Notation

We denote as \mathcal{R}_Q the ring of polynomials $\mathbb{Z}_Q[X]/(X^N + 1)$ where N is a power of two. We only use Q and N in the context of the ring \mathcal{R}_Q . We denote vectors with bold lowercase letters, e.g., \mathbf{v} , and matrices with uppercase letters \mathbf{V} . We denote an n dimensional column vector as $[f(\cdot, i)]_{i=1}^n$, where $f(\cdot, i)$ defines the i -th coordinate. For brevity, we will also denote as $[n]$ the vector $[i]_{i=1}^n$, and more generally $[i]_{i=n}^m$ the vector $[n, \dots, m]$. We address the i th entry of a vector \mathbf{v} by $\mathbf{v}[i]$. For matrices we address the i th row and j th column as $\mathbf{A}[i, j]$. Sometimes we view ring elements $\mathbf{a} \in \mathcal{R}_Q$ as vectors of coefficients and we address the coefficients as vector coordinates. For a random variable $x \in \mathbb{Z}$ we denote as $\text{Var}(x)$ the variance of x , as $\text{stddev}(x)$ its standard deviation and as $\text{E}(x)$ the upper bound on absolute value of the expectation. For $a \in \mathcal{R}_Q$, we define $\text{Var}(a)$, $\text{stddev}(a)$ and $\text{E}(a)$ to be the largest variance, standard deviation and absolute value of the expectation respectively among the coefficients of the polynomial a . By $\text{Hw}(\mathbf{a})$ we denote the hamming weight of the vector \mathbf{a} , i.e., the number of non-zero coordinates of \mathbf{a} . We represent numbers in \mathbb{Z}_Q as integers in $[-Q/2, Q/2)$.

²Assuming each integer is stored in a byte array. If integers are stored in 64-bit registers, then the key size grows to ≈ 9000 MB of memory.

We say that an algorithm is **PPT** if it is a probabilistic polynomial-time algorithm. We denote any polynomial as $\text{poly}(\cdot)$. We denote as $\text{negl}(\lambda)$ a negligible function in $\lambda \in \mathbb{N}$. That is, for any positive polynomial $\text{poly}(\cdot)$ there exists $c \in \mathbb{N}$ such that for all $\lambda \geq c$ we have $\text{negl}(\lambda) \leq \frac{1}{\text{poly}(\lambda)}$. Given two distributions X, Y over a finite domain D , their statistical distance is defined as $\Delta(X, Y) = \frac{1}{2} \sum_{v \in D} |X(v) - Y(v)|$. We say that two distributions are statistically close if their statistical distance is negligible.

Probability Theory

Lemma 1 (Smudging Lemma [AJL⁺12]). *Let B_1 and B_2 be two positive integers and let $e_1 \in [-B_1, B_1]$ be a fixed integer. Let $e_2 \leftarrow_{\$} [-B_2, B_2]$ be chosen uniformly at random. Then the statistical distance between e_2 and $e_2 + e_1$ is*

$$\Delta(e_2, e_2 + e_1) = B_1/B_2.$$

Lemma 2 (Lemma 2.3 from [DS16]). *Let $\delta \in [0, 1]$ and $f : \mathcal{S} \rightarrow \mathcal{S}$ be a randomized function such that $\Delta(f(a), f(b)) \leq \delta$ holds for all $a, b \in \mathcal{S}$. Then*

$$\forall k \geq 0, \forall a, b \in \mathcal{S}, \quad \Delta(f^k(a), f^k(b)) \leq \delta^k,$$

where f^k denotes composing the function f k -times.

Definition 1. The min-entropy of a random variable X is defined as

$$\mathbf{H}_{\infty}(X) = -\log \left(\max_x \Pr[X = x] \right)$$

Furthermore we recall the definition of average min-entropy A given B as

$$\tilde{\mathbf{H}}_{\infty}(A|B) = -\log \left(\mathbb{E}_{b \leftarrow B} \left[\max_a \Pr[A = a | B = b] \right] \right) = -\log \left(\mathbb{E}_{b \leftarrow B} \left[2^{-\mathbf{H}_{\infty}(A|B=b)} \right] \right).$$

Lemma 3 (Generalized Leftover Hash Lemma [DRS04, DORS08]). *Assume $\{\mathbf{H}_X \{0, 1\}^n \mapsto \{0, 1\}^{\ell}\}_{x \in \mathcal{X}}$ is a family of universal hash functions. Then, for any random variables W and I ,*

$$\Delta((\mathbf{H}_X(W), X, I), (U_{\ell}, X, I)) \leq \frac{1}{2} \cdot \sqrt{2^{\ell} \cdot 2^{-\tilde{\mathbf{H}}_{\infty}(W|I)}}$$

Lemma 4 (Lemma 2.2. in [DRS04, DORS08]). *Let A, B and C be random variables. Then*

1. *For any $\delta > 0$, the conditional entropy $\mathbf{H}_{\infty}(A|B = b)$ is at least $\tilde{\mathbf{H}}_{\infty}(A|B) - \log(1/\delta)$ with probability at least $1 - \delta$ over the choice of b .*
2. *If B has at most 2^{λ} possible values, then*

$$\tilde{\mathbf{H}}_{\infty}(A|(B, C)) \geq \tilde{\mathbf{H}}_{\infty}((A, B)|C) - \lambda \geq \tilde{\mathbf{H}}_{\infty}(A|C) - \lambda.$$

In particular, $\tilde{\mathbf{H}}_{\infty}(A|B) \geq \mathbf{H}_{\infty}((A, B)) - \lambda \geq \mathbf{H}_{\infty}(A) - \lambda$.

Lattices. An m -dimensional lattice Λ is a discrete additive subgroup of \mathbb{R}^m . For an integer $k < m$ and a rank matrix $\mathbf{B} \in \mathbb{R}^{m \times k}$, $\Lambda(\mathbf{B}) = \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^k\}$ is the lattice generated by the columns of \mathbf{B} . We denote $\Lambda_q^{\perp}(\mathbf{B}) = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{B}^{\top} \mathbf{v} = \mathbf{0} \pmod{q}\}$.

Gaussian distribution. For any $\sigma > 0$, the spherical Gaussian function with parameter σ is defined as $\rho_\sigma(\mathbf{x}) = \exp\left(\frac{-\pi\|\mathbf{x}\|^2}{\sigma^2}\right)$, for any $\mathbf{x} \in \mathbb{R}^m$. Given a lattice $\Lambda \subseteq \mathbb{R}^m$, a parameter $\sigma \in \mathbb{R}$ and a vector $\mathbf{c} \in \mathbb{R}^m$ the spherical Gaussian distribution with parameter σ and support $\Lambda + \mathbf{c}$ is defined as

$$\mathcal{D}_{\Lambda+\mathbf{c},\sigma}(\mathbf{x}) = \frac{\rho_\sigma(\mathbf{x})}{\rho_\sigma(\Lambda + \mathbf{c})}, \forall \mathbf{x} \in \Lambda + \mathbf{c}$$

where $\rho_\sigma(\Lambda + \mathbf{c})$ denotes $\sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho_\sigma(\mathbf{x})$.

We write $x \leftarrow_{\S} \mathcal{D}_{\Lambda+\mathbf{c},\sigma}$ to denote that x is sampled from the discrete Gaussian distribution with support $\Lambda + \mathbf{c}$ and parameter σ . We write $\boldsymbol{\eta} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}^N,\sigma}$ or $\mathbf{y} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}^n,\sigma}$ when sampling the coefficients of $\boldsymbol{\eta} \in \mathcal{R}$ or components of $\mathbf{y} \in \mathbb{Z}^N$ from $\mathcal{D}_{\mathbb{Z},\sigma}$. For a set \mathcal{S} we write $x \leftarrow_{\S} \mathcal{S}$ to denote the uniform distribution over \mathcal{S} unless said otherwise.

Throughout the paper we denote $C_{\delta,m} = \sqrt{\frac{\ln(2m(1+1/\delta))}{\pi}}$.

Definition 2 (Smoothing Parameter). For a lattice $\Lambda \subseteq \mathbb{Z}^m$ and positive real $\delta > 0$, the smoothing parameter η_δ is the smallest real $r > 0$ such that $\rho_{1/r}(\Lambda^* \setminus \{\mathbf{0}\}) \leq \delta$, where $\Lambda^* = \{\mathbf{x} \in \mathbb{R}^m \mid \mathbf{x}^\top \Lambda \subseteq \mathbb{Z}\}$.

Lemma 5 ([MR04], Lemma 3.3). *Let Λ be any rank- m lattice, and $\delta \in \mathbb{R}^+$. Then*

$$\eta_\delta \leq \lambda_m(\Lambda) \cdot C_{\delta,m},$$

where $\lambda_m(\Lambda)$ is the smallest R such that the ball \mathcal{B}_R centered in the origin and with radius R contains m linearly independent vectors of Λ . Remind that we denote $C_{\delta,m} = \sqrt{\frac{\ln(2m(1+1/\delta))}{\pi}}$.

Lemma 6 (Claim 3.8 in [Reg09]). *For any lattice Λ , $\mathbf{c} \in \mathbb{R}^n$, $\epsilon > 0$ and $\sigma \geq \eta_\epsilon$,*

$$\rho(\Lambda + \mathbf{c}) \in \frac{\sigma^n}{\det(\Lambda)} (1 \pm \epsilon)$$

Lemma 7 (Corollary 2.8 in [GPV08]). *Let $\Lambda \subseteq \mathbb{Z}^m$ be a lattice, $0 < \epsilon < 1$, $\sigma > 0$. For any vector $\mathbf{c} \in \mathbb{R}^m$, if $\sigma \geq \eta_\epsilon(\Lambda)$, then we have*

$$\rho(\Lambda + \mathbf{c}) \in \left[\frac{1-\epsilon}{1+\epsilon}, 1 \right] \cdot \rho_\sigma(\Lambda)$$

Learning With Errors. We recall the learning with errors assumption by Regev [Reg05]. Our description is a generalized version due to Brakerski, Gentry, and Vaikuntanathan [BGV12].

Definition 3 (Generalized Learning With Errors). Let \mathcal{D}_{sk} be a (not necessarily uniform) distribution over \mathcal{R}_Q , and $\sigma > 0$, $n \in \mathbb{N}$ and $N \in \mathbb{N}$ be a power of two, that are chosen according to a security parameter λ . We define a Generalized Learning With Errors (GLWE) sample of a message $\mathbf{m} \in \mathcal{R}_Q$ with respect to a secret key $\mathbf{s} \in \mathcal{D}_{\text{sk}}^n$, as

$$\text{GLWE}_{\sigma,n,N,Q}(\mathbf{s}, \mathbf{m}) = \begin{bmatrix} \mathbf{a}^\top \\ \mathbf{b} = \mathbf{a}^\top \cdot \mathbf{s} + \boldsymbol{\epsilon} \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{m} \end{bmatrix} \in \mathcal{R}_Q^{(n+1)},$$

where $\mathbf{a} \leftarrow_{\S} \mathcal{R}_Q^n$ and $\boldsymbol{\epsilon} \leftarrow_{\S} \mathcal{D}_{\mathcal{R},\sigma}$. We say that the $\text{GLWE}_{\sigma,n,N,Q}$ -assumption holds if for any PPT adversary A we have

$$\left| \Pr[A(\text{GLWE}_{\sigma,n,N,Q}(\mathbf{s}, 0)) =] - \Pr[A(\mathcal{U}_Q^{(n+1) \times 1}) =] \right| \leq \text{negl}(\lambda)$$

where $\mathcal{U}_Q^{(n+1) \times 1}$ is the uniform distribution over $\mathcal{R}_Q^{(n+1)}$.

We denote a Learning With Errors (LWE) sample as $\text{LWE}_{\sigma,n,Q}(\mathbf{s}, \mathbf{m}) = \text{GLWE}_{\sigma,n,1,Q}$, which is a special case of a GLWE sample where the ring is $\mathbb{Z}_q[X]$. Similarly we denote a Ring-Learning with Errors (RLWE) sample as $\text{RLWE}_{\sigma}(\mathbf{s}, \mathbf{m}) = \text{GLWE}_{\sigma,1,N,Q}$ which is the special case of an GLWE sample with $n = 1$. For simplicity, we omit to state the modulus and ring dimension for RLWE samples because we always use $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$ where N is a power of two. For LWE samples, we will be switching between different moduli and different dimensions; hence we will indicate the current modulus in the notation. Sometimes we use the notation $\mathbf{c} \in \text{GLWE}_{\sigma,n,1,Q}(\mathbf{s}, \mathbf{m})$ (resp. LWE and RLWE) to indicate that a vector \mathbf{c} is a GLWE (resp. LWE and RLWE) sample of the corresponding parameters and inputs. Sometimes we leave the inputs unspecified and substitute them with “” when it is not necessary to refer to them within the scope of a function. We define the phase of $\mathbf{c} = \text{GLWE}_{\sigma,n,N,Q}(\mathbf{s}, \mathbf{m})$, as $\text{Phase}(\mathbf{c}) = [1, -\mathbf{s}] \cdot \mathbf{c}$. We define the error of \mathbf{c} as $\text{Error}(\mathbf{c}) = \text{Phase}(\mathbf{c}) - \mathbf{m}$.

Fully Homomorphic Encryption. Below we recall the definition of fully homomorphic encryption [RAD78, Gen09b].

Definition 4 (Fully Homomorphic Encryption). A fully homomorphic encryption FHE consists of algorithms (Setup, Enc, Eval, Dec) with the following syntax.

Setup(λ): This **PPT** algorithm takes as input a security parameter λ and outputs an evaluation key ek and a secret key sk .

Enc(sk, \mathbf{m}): This **PPT** algorithm takes as input a secret key sk , and a message \mathbf{m} , and returns a ciphertext ct .

Eval($\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C}$): Given as input an evaluation key ek , a set of ciphertexts $[\text{ct}_i]_{i=1}^n$, and a circuit \mathcal{C} , this (non-)deterministic algorithm outputs a ciphertext ct .

Dec(sk, ct): Given a secret key sk and a ciphertext ct , this deterministic algorithm outputs a message \mathbf{m} .

Correctness: We say that $\text{FHE} = (\text{Setup}, \text{Enc}, \text{Eval}, \text{Dec})$ is correct, if for all security parameters $\lambda \in \mathbb{N}$, circuits $\mathcal{C} : \mathcal{M}^n \mapsto \mathcal{M}$ over the message space \mathcal{M} of depth $\text{poly}(\lambda)$, and messages $[\mathbf{m}_i \in \mathcal{M}]_{i=1}^n$ we have

$$\Pr [\text{Dec}(\text{sk}, \text{ct}_{\text{out}}) = \mathcal{C}([\mathbf{m}_i]_{i=1}^n)] = 1 - \text{negl}(\lambda),$$

where $(\text{ek}, \text{sk}) \leftarrow \text{Setup}(\lambda)$, $[\text{Dec}(\text{sk}, \text{ct}_i) = \mathbf{m}_i]_{i=1}^n$ and $\text{ct}_{\text{out}} \leftarrow \text{Eval}(\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C})$.

Efficiency: We require that Setup, Enc and Dec run in $\text{poly}(\lambda)$ time, and Eval runs in $\text{poly}(\lambda, |\mathcal{C}|)$ time. Finally, we say that fully homomorphic encryption is compact if the size of the output of Eval is independent of \mathcal{C} . Namely, if $|\text{Eval}(\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C})|$ is $\text{poly}(\lambda, |\mathcal{M}|)$.

Definition 5 (Indistinguishability Under Chosen Plaintext Attack). Let $\lambda \in \mathbb{N}$ be a security parameter and $\mathbf{A} = (\mathbf{A}_0, \mathbf{A}_1)$ be a **PPT** adversary. We define the advantage $\text{Adv}_{\mathbf{A}, \text{FHE}}^{\text{INDCPA}}(\lambda)$. We say that a FHE scheme is INDCPA-secure if for all **PPT** adversaries \mathbf{A} the following probability

$$\Pr \left[\mathbf{A}_1(\text{ct}_b, \text{st}) = b : \begin{array}{l} \text{sk} \leftarrow \text{Setup}(\lambda), \\ (\text{st}, \mathbf{m}_0, \mathbf{m}_1) \leftarrow \mathbf{A}_0^{\mathcal{O}(\text{sk}, \cdot)}(\lambda), \\ b \leftarrow_{\mathcal{S}} \{0, 1\}, \\ \text{ct}_b \leftarrow \text{Enc}(\lambda, \text{sk}, \mathbf{m}_b) \end{array} \right],$$

is at most $\text{negl}(\lambda)$, where the oracle \mathcal{O} on input a message \mathbf{m} outputs $\text{ct} \leftarrow \text{Enc}(\lambda, \text{sk}, \mathbf{m})$.

Circuit Privacy: Let $\mathcal{C} : \mathcal{M}^n \mapsto \mathcal{M}$ be a polynomial size circuit. A fully homomorphic encryption scheme FHE is said to be circuit private if there exists a **PPT** simulator for all fixed $(\text{ek}, \text{sk}) \leftarrow \text{Setup}(\lambda)$ and for all c_1, \dots, c_n such that $[m_i \leftarrow \text{Dec}(\text{sk}, c_i)]_{i=1}^n$ and $m_{\text{out}} \leftarrow \mathcal{C}(m_1, \dots, m_n)$ and

$$\Delta((\text{sk}, \text{Sim}(\text{ek}, m_{\text{out}})), (\text{sk}, \text{Eval}(\text{ek}, c_1, \dots, c_n, \mathcal{C}))) \leq \text{negl}(\lambda).$$

Our simulation-based definition of circuit privacy is stronger than [IP07, BdPMW16] in two aspects. First, our simulator does not require us to know the size of the circuit as in [IP07]. In fact, our simulator only needs to know the outcome of the circuit and nothing else. Second, we only assume that the ciphertexts input to the evaluator decrypt to the messages input to the circuit.

3 Sanitization Bootstrapping

We describe all algorithms necessary to build the sanitization bootstrapping. For algorithms that are part of the sanitization bootstrapping but are not crucial as for the circuit privacy analysis in Section 4, we only define the interfaces and state their correctness and functionality. In Appendix B we give the full specification and correctness proofs of these algorithms.

Gadgets and Gaussian Sampling. Let us first denote $\ell = \lceil \log_L Q \rceil$ for some radix $L \in \mathbb{N}$. In particular we denote \mathbf{L}_{br} for the blind rotation key defined by Figure 2. We also use \mathbf{L}_{ksK} as a decomposition base of the key-switching procedure, the interface of which we recall in Lemma 13 but defer the full specification of this algorithm to Appendix B.

Let $\mathbf{g}_{L,Q} = [1, L, \dots, L^{\ell-1}]$ be the gadget vector parameterized by L and Q . We use different decomposition algorithms but refer to all with the same interface. In particular, we have the decomposition algorithm $\mathbf{x} = \mathbf{G}_{\text{ver}}^{-1}(\mathbf{c}, L; \sigma) \in \mathcal{R}^\ell$ that takes as input a ring element $\mathbf{c} \in \mathcal{R}_Q$, a radix L , and optionally a Gaussian parameter σ , and outputs a low norm vector $\mathbf{x} \in \mathcal{R}^\ell$ such that $\mathbf{c} = \mathbf{g}_{L,Q}^\top \cdot \mathbf{x} \in \mathcal{R}_Q$. Note that $\mathbf{G}_{\text{ver}}^{-1}$ also takes the modulus Q implicitly as input. A special case of the above is when $\mathbf{G}_{\text{ver}}^{-1}$ takes as input a single element from \mathbb{Z}_Q instead of a polynomial from \mathcal{R}_Q . We use a parameter ver that takes a value from $\{\text{simul}, \text{det}\}$. If $\text{ver} = \text{simul}$ then we apply the algorithm from Lemma 8 coefficient wise. In particular, in our implementation we implement two algorithms from [MP12, GM18]. One for a general $Q < L^\ell$ and one specialized for $Q = L^\ell$. We recall both in Appendix A, but in Lemma 8 we refer only to the case with $Q < L^\ell$ since, for the other case, we found it hard to find efficient parameters (despite the Gaussian sampling for $Q = L^\ell$ being more efficient). We give more details on the parameters in Section 5. Note that for $\text{ver} = \text{simul}$, we take the additional $\sigma_{\mathbf{x}}$ as input. For $\text{ver} = \text{det}$, we take the deterministic decomposition algorithm like binary decomposition but generalized to any radix $L \geq 2$ and apply it coefficient-wise to elements from \mathcal{R}_Q . We note that for the det -mode, we may also use randomized algorithms, e.g., the subgaussian sampling algorithms [GMP19]. We can generalize the gadget vector for some $w \in \mathbb{N}$ to a matrix $\mathbf{G}_{L,Q,w} = \mathbf{g}_{L,Q} \otimes \mathbf{I}_w \in \mathbb{Z}_Q^{w \cdot \ell \times w}$. Then the decomposition algorithm takes as input vectors $\mathbf{a} \in \mathcal{R}_Q^w$ and outputs $\mathbf{x} \in \mathcal{R}_L^{w \cdot \ell}$ such that $\mathbf{a} = \mathbf{x}^\top \cdot \mathbf{G}_{L,Q,w}$.

Lemma 8 (Gaussian Sampling [GM18]). *There exists a sampling algorithm $\mathbf{G}_{\text{simul}}^{-1}(a, L, \sigma_{\mathbf{x}})$ that on input $a \in \mathbb{Z}_Q$, $L \in \mathbb{N}$ and a Gaussian parameter $\sigma_{\mathbf{x}}$, outputs $\mathbf{y} \in \mathbb{Z}^\ell$ such that*

$$\Delta(\mathbf{y}, \mathbf{x}[i] \in \mathcal{D}_{\Lambda_Q^\perp(\mathbf{g}_{L,Q}) + \mathbf{G}_{\text{det}}^{-1}(a[i], L), \sigma_{\mathbf{x}}}) \leq \ell \cdot \delta,$$

if $\sigma_{\mathbf{x}} \leq \sqrt{2L} \cdot (2L + 1) \cdot C_{\delta, \ell}$.

RGSW($\mathfrak{s}, \mathbf{m}_G$): <hr/> Input: Secret key $\mathfrak{s} \in \mathcal{R}_Q$. Message $\mathbf{m}_G \in \mathcal{R}_Q$. <hr/> 1: For $i \in [\ell_{\text{br}}]$: 2: $\mathbf{C}_G[* , i] \leftarrow \text{RLWE}_{\sigma_G}(\mathfrak{s}, \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{i-1})$. 3: For $i \in [\ell_{\text{br}} + 1, 2\ell_{\text{br}}]$: 4: $\mathbf{C}_G[* , i] \leftarrow \text{RLWE}_{\sigma_G}(\mathfrak{s}, -\mathfrak{s} \cdot \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{i-1-\ell_{\text{br}}})$. 5: Return $\mathbf{C}_G \in \mathcal{R}_Q^{2 \times 2\ell_{\text{br}}}$.	extProd_{ver}($\mathbf{c}, \mathbf{C}_G; \sigma_{\mathbf{x}}$): <hr/> Input: Ciphertext $\mathbf{c} \in \text{RLWE}_{\sigma}(\mathfrak{s}, \mathbf{m})$. Ciphertext $\mathbf{C}_G \in \text{RGSW}_{\sigma_G}(\mathfrak{s}, \mathbf{m}_G)$. [If simul] A Gaussian param. $\sigma_{\mathbf{x}}$. <hr/> 1: $\mathbf{c}_{\text{out}} \leftarrow \mathbf{C}_G \cdot \mathbf{G}_{\text{ver}}^{-1}(\mathbf{c}, \mathbf{L}_{\text{br}}; \sigma_{\mathbf{x}})$. 2: Return $\mathbf{c}_{\text{out}} \in \mathcal{R}_Q^2$.
--	---

Figure 1: RGSW Encryption and External Product.

Depending on the `ver` parameter, the distribution of the image of $\mathbf{G}_{\text{ver}}^{-1}$ may greatly differ. In the correctness analysis we denote the noise of $\mathbf{G}_{\text{ver}}^{-1}$'s output as $\mathbf{B}(\mathbf{G}^{-1}(\cdot, \mathbf{L}))$. For example, for deterministic base- \mathbf{L} decomposition, we take \mathbf{L}^2 , or when the decomposition returns a discrete Gaussian, we take its variance. We concretize this quantity when estimating correctness in Section 5.

Ring GSW Encryption. We recall the ring-version of the RGSW cryptosystem [GSW13] on Figure 1. We also recall the external product [CGGI16a, CGGI20], that multiplies an RGSW ciphertexts with an RLWE ciphertext. Below we state the functionality of the external product, but we limit our exposition to the case of binary plaintexts, which is the relevant case in our application.

Lemma 9 (The External Product). *Let \mathbf{c} and \mathbf{C}_G be RLWE and GSW encryptions of \mathbf{m} and \mathbf{m}_G , respectively, as in Figure 1. If $\mathbf{c}_{\text{out}} \leftarrow \text{extProd}_{\text{ver}}(\mathbf{C}_G, \mathbf{c}; \sigma_{\mathbf{x}})$ and $\mathbf{m}_G \in \{0, 1\}$ then $\mathbf{c}_{\text{out}} \in \text{RLWE}_{\sigma_{\text{out}}}(\mathfrak{s}, \mathbf{m}_{\text{out}})$, where $\mathbf{m}_{\text{out}} = \mathbf{m} \cdot \mathbf{m}_G$ and*

$$\sigma_{\text{out}} \leq \sqrt{2\ell_{\text{br}} \cdot N \cdot \sigma_{\text{br}}^2 \cdot \mathbf{B}(\mathbf{G}_{\text{ver}}^{-1}(\cdot, \mathbf{L}_{\text{br}}; \sigma_{\mathbf{x}})) + \mathbf{m}_G \sigma^2}.$$

Mux Gate. Informally, the Mux gate takes as input a control RGSW sample \mathbf{C} and two RLWE samples \mathbf{d} and \mathbf{h} . The gate outputs an RLWE encoding one of the message from \mathbf{d} or \mathbf{h} depending on the bit encoded in \mathbf{C} .

Lemma 10 (Homomorphic Mux Gate). *The Mux algorithm takes as input $\mathbf{C} \in \text{RGSW}_{\sigma_{\mathbf{C}}}(\mathfrak{s}, \mathbf{m}_{\mathbf{C}})$, $\mathbf{d} \in \text{RLWE}_{\sigma}(\mathfrak{s}, \mathbf{m}_{\mathbf{d}})$ and $\mathbf{h} \in \text{RLWE}_{\sigma}(\mathfrak{s}, \mathbf{m}_{\mathbf{h}})$, where $\mathbf{m}_{\mathbf{C}} \in \{0, 1\}$ and $\mathbf{m}_{\mathbf{d}}, \mathbf{m}_{\mathbf{h}} \in \mathcal{R}_Q$. Optionally it also takes a Gaussian parameter $\sigma_{\mathbf{x}}$. In particular, the gate computes and outputs $\text{extProd}_{\text{ver}}(\mathbf{C}_G, \mathbf{d} - \mathbf{h}; \sigma_{\mathbf{x}}) + \mathbf{h}$. If $\mathbf{c}_{\text{out}} \leftarrow \text{Mux}(\mathbf{C}, \mathbf{d}, \mathbf{h}; \sigma_{\mathbf{x}})$, then $\mathbf{c}_{\text{out}} \in \text{RLWE}_{\sigma_{\text{out}}}(\mathfrak{s}, \mathbf{m}_{\text{out}})$, where $\mathbf{m}_{\text{out}} = \mathbf{m}_{\mathbf{h}}$ for $\mathbf{m}_{\mathbf{C}} = 0$ and $\mathbf{m}_{\text{out}} = \mathbf{m}_{\mathbf{d}}$ for $\mathbf{m}_{\mathbf{C}} = 1$, and*

$$\sigma_{\text{out}} \leq \sqrt{2\ell_{\text{br}} \cdot N \cdot \sigma_{\text{br}}^2 \cdot \mathbf{B}(\mathbf{G}_{\text{ver}}^{-1}(\cdot, \mathbf{L}_{\text{br}}; \sigma_{\mathbf{x}})) + \sigma^2}$$

Modulus Switching and Sample Extraction. The modulus switching technique [BV11] allows us to change the modulus of a given ciphertext without the knowledge of the secret key.

Lemma 11 (Modulus Switching). *Let $\mathbf{c} = \text{LWE}_{\sigma, n, Q}(\mathfrak{s}, \mathbf{m})$. The modulus switching algorithm is defined as $\text{ModSwitch}(\mathbf{c}, q) = \left\lfloor \frac{q \cdot \mathbf{c}[i]}{Q} \right\rfloor$. If $\mathbf{c}_{\text{out}} \leftarrow \text{ModSwitch}(\mathbf{c}, q)$, then \mathbf{c}_{out}*

$\in \text{LWE}_{\sigma_{\text{out}}, n, q}(\mathbf{s}, \mathbf{m} \cdot \frac{q}{Q})$, where

$$\sigma_{\text{out}} \leq \sqrt{\left(\frac{q}{Q} \cdot \sigma\right)^2 + \frac{1}{4} \cdot \text{Hw}(\mathbf{s}) \cdot \text{Var}(\mathbf{s})}.$$

Furthermore, the expectation of $\text{Error}(\mathbf{c}_{\text{out}})$ satisfies

$$|\mathbb{E}(\text{Error}(\mathbf{c}_{\text{out}}))| \leq \left|\frac{q}{Q} \cdot \mathbb{E}(\text{Error}(\mathbf{c}))\right| + \frac{1}{2}(1 + \text{Hw}(\mathbf{s}) \cdot |\mathbb{E}(\mathbf{s})|)$$

Finally, if $\mathbf{m} = \mathbf{m}' \cdot \frac{Q}{t}$, then $\mathbf{m} \cdot \frac{q}{Q} = \mathbf{m}' \cdot \frac{q}{t}$.

Sample extraction, allows extracting an LWE from an RLWE sample that encodes the constant coefficient of the RLWE sample's message.

Lemma 12 (Sample Extraction). *Let $\text{KeyExtract}(\mathbf{s})$ be an algorithm that on input a key $\mathbf{s} \in \mathcal{R}_Q$ outputs its coefficient vector. The sample extraction algorithm $\text{LWE-Ext}(\mathbf{c})$ takes as input $\mathbf{c} \in \text{RLWE}_{\sigma}(\mathbf{s}, \mathbf{m})$ and outputs $\mathbf{c}_{\text{out}} = [\mathbf{a}, b] \in \mathbb{Z}_Q^{N+1}$ where $b = \mathbf{b}[1]$, and for all $i \in [N-1]$ we set $\mathbf{a}[i] \leftarrow -\mathbf{a}[N-i+1]$ and set $\mathbf{a}[1] \leftarrow \mathbf{a}[1]$.*

Denote the message encoded in \mathbf{c} as $\mathbf{m} = \sum_{i=1}^N \mathbf{m}[i] \cdot X^{i-1}$. If $\mathbf{s}' \leftarrow \text{KeyExtract}(\mathbf{s})$ and $\mathbf{c}_{\text{out}} \leftarrow \text{LWE-Ext}(\mathbf{c})$, then $\mathbf{c}_{\text{out}} \in \text{LWE}_{\sigma_{\text{out}}, N, Q}(\mathbf{s}', \mathbf{m}[k])$, where $\sigma_{\text{out}} = \sigma$.

Key Switching. By having a key switching key, the evaluator can map a given LWE sample to an LWE sample of a different key and dimension. We recall the interface for key switching and state its functionality by Lemma 13. We recall the full specification in Appendix B.

Lemma 13 (Key Switching). *We define the key-switching key generation procedure $\text{ksK} \leftarrow \text{KeySwitchSetup}(\sigma_{\text{ksK}}, \mathbf{s}, \mathbf{s}')$, to take as input a noise parameter σ_{ksK} , and LWE secret keys $\mathbf{s} \in \mathbb{Z}_Q^n$ and $\mathbf{s}' \in \mathbb{Z}_Q^N$ of (possibly) distinct dimensions $n, N \in \mathbb{N}$. The key-switch procedure $\mathbf{c}_{\text{out}} \leftarrow \text{KeySwitch}(\mathbf{c}, \text{ksK})$ takes as input a LWE ciphertext $\mathbf{c} \in \text{LWE}_{\sigma, N, Q}(\mathbf{s}', m)$ and the key-switching key ksK , and outputs a LWE sample $\mathbf{c}_{\text{out}} \in \text{LWE}_{\sigma_{\text{out}}, n, Q}(\mathbf{s}, m)$, where*

$$\sigma_{\text{out}} \leq \sqrt{\ell_{\text{ksK}} \cdot N \cdot \mathbf{B}(\mathbf{G}_{\text{det}}^{-1}(\cdot, \text{L}_{\text{ksK}})) \cdot \sigma_{\text{ksK}}^2 + \sigma^2}.$$

Randomized Blind Rotation and Sanitization Bootstrapping. In Figure 3 we show our sanitization bootstrapping. We give the sanitizing blind rotation and its key generation algorithm in Figure 2. In short the algorithm is given a LWE sample which phase is $m + e$ (e is the noise term) and outputs a RLWE sample of $\mathbf{a}_{\text{rot}} \cdot X^{m+e} \in \mathcal{R}_Q$. For any function f that is negacyclic, i.e., satisfies $f(x + N \bmod 2N) = -f(x) \bmod Q$ we choose the rotation polynomial \mathbf{a}_{rot} such that the constant coefficient of $\mathbf{a}_{\text{rot}} \cdot X^{m+e}$ is set to $f(m + e) \in \mathbb{Z}_Q$. We stress that the restriction on f is imposed by structural properties of the ring $\mathcal{R}_Q = \mathbb{Z}_Q[X]/(X^N + 1)$. In Appendix B we recall a version of the algorithm that applies a trick from [YXS⁺21, LMP22], which resolves the negacyclicity restriction on the functions that we can compute on the input plaintext at the cost of two blind rotation operations. Namely, we can program the polynomial \mathbf{a}_{rot} such that $F(m+e) = \mathbf{a}_{\text{rot}} \cdot X^{m+e}[1] \in \mathbb{Z}_Q$, where F is any function in \mathbb{Z}_N . Such full domain functional bootstrapping got recently much attention [KS22, YXS⁺21, CLOT21, LMP22, Klu22], as it allows to compute any function on finite fields, conveniently switch from finite field plaintexts to binary and back, etc.

Lemma 14 (Correctness of Bootstrapping). *Let br , \mathbf{c} and all other parameters be as in Figure 3, ksK be generated as described by Lemma 13, where $\mathbf{s}' \leftarrow \text{KeyExtract}(\mathbf{s})$.*

BRKeyGen($\sigma_{br}, \mathfrak{s}, \mathfrak{s}$):	BlindRotate _{ver} ($br, a_{rot}, c; \sigma_{\mathbf{x}}$):
<p>Input:</p> <p>An error distribution σ_{br}.</p> <p>A RLWE secret key $\mathfrak{s} \in \mathcal{R}_Q$.</p> <p>A LWE secret key $\mathfrak{s} \in \mathbb{Z}_t^n$.</p>	<p>Input:</p> <p>A blind rotation key $br = \text{RGSW}_{\sigma_{br}}(\mathfrak{s}, \cdot)^n$.</p> <p>An rotation polynomial $a_{rot} \in \mathcal{R}_Q$.</p> <p>A ciphertext $c \in \text{LWE}_{\sigma, n, 2N}(\mathfrak{s}, \cdot)$.</p> <p>[If simul] A Gaussian param. $\sigma_{\mathbf{x}}$.</p>
<p>1 : For $i \in [n]$</p> <p>2 : Set $br[i] = \text{RGSW}_{\sigma_{br}}(\mathfrak{s}, \mathfrak{s}[i])$.</p> <p>3 : Output $br \in \text{RGSW}_{\sigma_{br}}(\mathfrak{s}, \cdot)^n$.</p>	<p>1 : Let $\mathbf{c} = [\mathbf{a}, b] \in \mathbb{Z}_{2N}^{n+1}$.</p> <p>2 : Set $\mathbf{c}_{acc,0} \leftarrow [0, a_{rot} \cdot X^b] \in R_Q^2$.</p> <p>3 : For $i \in [n]$:</p> <p>4 : $\mathbf{c}_{acc,i} \leftarrow \text{Mux}_{ver}(br[i],$ <div style="text-align: right; margin-left: 20px;">$\mathbf{c}_{acc,i-1} \cdot X^{-a[i]},$ $\mathbf{c}_{acc,i-1};$ $\sigma_{\mathbf{x}})$.</div> </p> <p>5 : Output $\mathbf{c}_{acc,n} \in R_Q^2$.</p>

Figure 2: TFHE-style Blind Rotation and its Setup.

Let a_{rot} be such that $f(m + e) = a_{rot} \cdot X^{m+e}[1] \in \mathcal{R}_Q$, where $m + e = \text{Phase}(\mathbf{c})$ and $f : \mathbb{Z}_{2N} \mapsto \mathbb{Z}_Q$. If $\mathbf{c}_{out} = \text{Bootstrap}_{ver}(br, ksK, \mathbf{c}, a_{rot})$, then $\mathbf{c}_{out} \in \text{LWE}_{\sigma_{out}, N, Q}(\mathfrak{s}', f(m+e))$, with

$$\begin{cases} \sigma_{out} \leq \sqrt{2n \cdot \ell_{br} \cdot N \cdot \sigma_{br}^2 \cdot B(G_{det}^{-1}(\cdot, L_{br}))} & \text{if ver = det} \\ \sigma_{out} \leq \sqrt{2n \cdot \ell_{br} \cdot N \cdot \sigma_{br}^2 \cdot B(G_{simul}^{-1}(\cdot, L_{br}; \sigma_{\mathbf{x}})) + h \cdot \sigma_R^2 \cdot \sigma_{rand}^2} & \text{if ver = simul} \end{cases}$$

Additionally, we have that \mathbf{c}_{in} from step 2 on Figure 3 is such that $\mathbf{c}_{in} = \text{LWE}_{\sigma_{in}, n, 2N}(\mathfrak{s}, \cdot)$, where

$$\sigma_{in} \leq \sqrt{\left(\frac{q}{Q} \cdot \sigma_1\right)^2 + \frac{1}{4} \cdot \text{Hw}(\mathfrak{s}) \cdot \text{Var}(\mathfrak{s})}$$

with $\sigma_1 \leq \sqrt{N \cdot \ell_{ksK} \cdot B(G_{ver}^{-1}(\cdot, L_{ksK})) \cdot \sigma_{ksK}^2 + \sigma_{out}^2}$.

The full cryptosystem. Below we briefly describe how the complete cryptosystem fits into Definition 4.

Setup: We choose the modulus Q , a power-of-two dimension N of the ring \mathcal{R}_Q and LWE dimension $n \in \mathbb{N}$. Then we choose $\mathfrak{s} \in \mathcal{R}_Q$ for the RLWE key, set $\mathfrak{s}' \leftarrow \text{KeyExtract}(\mathfrak{s})$, and $\mathfrak{s} \in \{0, 1\}^n$ for the LWE key³. Choose the radices $L_{br}, L_{ksK} \in \mathbb{N}$ and the Gaussian parameters $\sigma, \sigma_{ksK}, \sigma_{br}, \sigma_R, \sigma_{rand}, \sigma_{\mathbf{x}} > 0$. Run $br \leftarrow \text{BRKeyGen}(\sigma_{br}, \mathfrak{s}, \mathfrak{s})$, $ksK \leftarrow \text{KeySwitchSetup}(\sigma_{ksK}, \mathfrak{s}, \mathfrak{s}')$, and $\mathbf{v} \leftarrow \text{LWE}_{\sigma_R, N, Q}(\mathfrak{s}', 0)^h$. Finally, set the evaluation key $ek = (br, ksK, \mathbf{v})$ and the secret key $sk = (\mathfrak{s}, \mathfrak{s}', \mathfrak{s})$.

Encryption: To encrypt a message $m' \in \mathbb{Z}_t$ we compute $\mathbf{c} = \text{LWE}_{\sigma, N, Q}(\mathfrak{s}', m) \in \mathbb{Z}_Q^{N+1}$, where $m = \frac{Q}{t} \cdot m' \in \mathbb{Z}_Q$. Note that we can also use the vector \mathbf{v} to obtain a LWE sample that is close to a “fresh” one, and then we simply add m .

³Other distributions for the LWE secret key are possible. See [MP21] for an excelent summary.

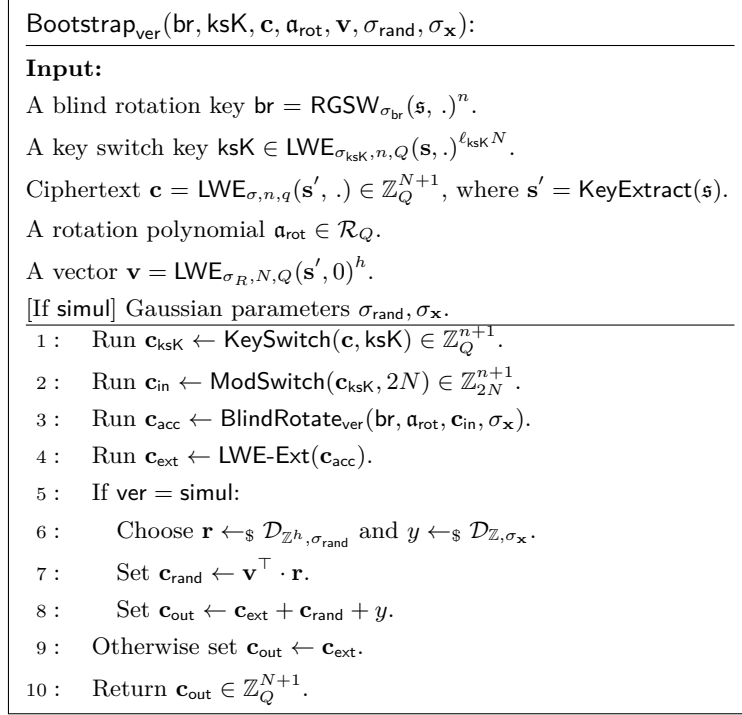


Figure 3: Bootstrapping.

Eval: We can represent homomorphic computation as a circuit with gates of the form $f(b + \sum_{i=1}^k x_i \cdot a_i \in \mathbb{Z}_{t_1}) \in \mathbb{Z}_{t_2}$ where the a_i 's and b are scalars known by the evaluator and the x_i 's are the encrypted plaintexts. We compute the affine function using the additive homomorphism of the LWE samples, and the function $f : \mathbb{Z}_{t_1} \mapsto \mathbb{Z}_{t_2}$ by applying the bootstrapping algorithm from Figure 3. We compute all gates with ver = det except for the output gates, where we run the sanitization bootstrap with ver = simul. Crucially, the evaluator should finish the computation with a sanitization bootstrap to achieve circuit privacy.

Decryption: Do decrypt a LWE sample $\mathbf{c}_{\text{out}} = [\mathbf{a}_{\text{out}}, b_{\text{out}}]$ we run $\text{Phase}(\mathbf{c}_{\text{out}}) = \mathbf{c}_{\text{out}}^\top [1, -\mathbf{s}] = b - \mathbf{a}_{\text{out}}^\top \mathbf{s} = \frac{Q}{t} m'_{\text{out}} + e \in \mathbb{Z}_t$, and round the result $\left\lfloor \frac{t}{Q} \left(\frac{Q}{t} m'_{\text{out}} + e \right) \right\rfloor = m'_{\text{out}}$ if $|e| \leq \frac{Q}{2t}$.

4 Analysis of Circuit Privacy

This section contains our core analytical contribution. First, in Section 4.1, we state a few technical lemmas needed for the circuit-privacy analysis in Section 4.2.

4.1 Generalized (Gaussian) Leftover Hash Lemma

Below we give our fixed and generalized version of the Gaussian Leftover hash lemma from [BdPMW16].

Lemma 15 (Gaussian Leftover Hash Lemma (Generalized Lemma 3.6 from [BdPMW16])). *Let $\delta, \sigma_{\mathbf{x}} > 0$. Let $\mathcal{L} = \{\mathbf{v} \in \hat{\Lambda} : \hat{\mathbf{e}}^\top \cdot \mathbf{v} = 0\}$, where $\hat{\mathbf{e}} = [\mathbf{e}, 1] \in \mathbb{Z}^{m+1}$ is fixed, $\hat{\Lambda} = \hat{\Lambda}_Q^\perp(\mathbf{G}_{L, Q, w}) \times \mathbb{Z}$, $m = w \cdot \ell$ and $L^\ell \leq Q$. Let $\mathbf{q}_{L, Q} = [q_i]_{i=1}^\ell$ be the base-L decomposition*

of Q for $Q < L^\ell$, and $\mathbf{q}_{L,Q} = [\mathbf{0}, L]$ for $Q = L^\ell$. For any $\mathbf{e} \in \mathbb{Z}^m$ and $\mathbf{c} \in \mathbb{R}^m$, if the Gaussian parameter $\sigma_{\mathbf{x}}$ satisfies

$$\sigma_{\mathbf{x}} \geq \sqrt{1 + \|\hat{\mathbf{e}}\|_\infty} \cdot \max(\|\mathbf{q}_{L,Q}\|, \sqrt{L^2 + 1}) \cdot C_{\delta,m}, \text{ then}$$

$$\Delta(\mathbf{e}^\top \mathbf{x} + y, e') < 2\delta,$$

where $\mathbf{x} \leftarrow_{\S} \mathcal{D}_{\Lambda_Q^\perp(\mathbf{G}_{L,Q,w}) + \mathbf{c}, \sigma_{\mathbf{x}}}$, $y \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}}}$, $e' \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}} \cdot \sqrt{1 + \|\mathbf{e}\|_2^2}}$, and $m = w \cdot \ell$ with $\ell \in \mathbb{N}$. Note that the distribution of e' is independent of the coset \mathbf{c} , and if $\mathbf{e} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\text{br}}}$, then $e' \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}} \cdot \sqrt{1 + m\sigma_{\text{br}}^2}}$.

The proof of the lemma follows from a technical lemma (Corollary 2.8 in [GPV08]), and a lemma (Lemma 3.7 in [BdPMW16]) that bounds the smoothing parameter η_δ for the lattice $\mathcal{L} = \{\mathbf{v} \in \hat{\Lambda} : \hat{\mathbf{e}}^\top \cdot \mathbf{v} = 0\}$, where $\hat{\mathbf{e}} = [\mathbf{e}, 1] \in \mathbb{Z}^{m+1}$. In [BdPMW16], the authors prove the lemma for a modulus Q that is a power of two. We generalize the lemma to modulus of form $Q \leq L^\ell$.

Proof. Let $\hat{\mathbf{c}} = [\mathbf{c}, 0] \in \mathbb{Z}^{m+1}$ and $\hat{\Lambda} = \Lambda_Q^\perp(\mathbf{G}_{L,Q,w}) \times \mathbb{Z}$. We want to show that

$$\Delta(\hat{\mathbf{e}}^\top \mathcal{D}_{\hat{\Lambda} + \hat{\mathbf{c}}, \sigma_{\mathbf{x}}}, \mathcal{D}_{\mathbb{Z}, \|\hat{\mathbf{e}}\| \sigma_{\mathbf{x}}}) \leq 2\delta.$$

The support of $\hat{\mathbf{e}}^\top \mathcal{D}_{\hat{\Lambda} + \hat{\mathbf{c}}}$ is $\hat{\mathbf{e}}^\top \hat{\Lambda} + \hat{\mathbf{e}}^\top \hat{\mathbf{c}} = \mathbf{e}^\top \Lambda_Q^\perp(\mathbf{G}_{L,Q,w}) + \mathbb{Z} + \mathbf{e}^\top \mathbf{c} = \mathbb{Z}$. Fix some $z \in \mathbb{Z}$. The probability mass assigned to z by $\hat{\mathbf{e}}^\top \mathcal{D}_{\hat{\Lambda} + \hat{\mathbf{c}}, \sigma_{\mathbf{x}}}$ is proportional to $\rho_{\sigma_{\mathbf{x}}}(\mathcal{L}_z)$, where

$$\mathcal{L}_z = \{\mathbf{v} \in \hat{\Lambda} + \hat{\mathbf{c}} : \hat{\mathbf{e}}^\top \mathbf{v} = z\}.$$

We define the lattice $\mathcal{L} = \{\mathbf{v} \in \hat{\Lambda} : \hat{\mathbf{e}}^\top \mathbf{v} = 0\}$; note that $\mathcal{L}_z = \mathcal{L} + \mathbf{w}_z$ for any $\mathbf{w}_z \in \mathcal{L}_z$. Let $\mathbf{u}_z = \frac{z}{\|\hat{\mathbf{e}}\|^2 \sigma_{\mathbf{x}}} \hat{\mathbf{e}}$. Then \mathbf{u}_z is clearly proportional to $\hat{\mathbf{e}}$. Observe that \mathbf{u}_z is orthogonal to $\sigma_{\mathbf{x}}^{-1} \mathcal{L}_z - \mathbf{u}_z$. Indeed for any $\mathbf{t} \in \sigma_{\mathbf{x}}^{-1} \mathcal{L}_z$ we have $\hat{\mathbf{e}}^\top (\mathbf{t} - \mathbf{u}_z) = 0$. From this we have $\rho(\mathbf{t}) = \rho(\mathbf{u}_z) \cdot \rho(\mathbf{t} - \mathbf{u}_z)$, and by summing for $\mathbf{t} \in \sigma_{\mathbf{x}}^{-1} \mathcal{L}_z$:

$$\rho(\sigma_{\mathbf{x}}^{-1} \mathcal{L}_z) = \rho(\mathbf{u}_z) \cdot \rho(\sigma_{\mathbf{x}}^{-1} \mathcal{L}_z - \mathbf{u}_z).$$

Observe that we have $\sigma_{\mathbf{x}}^{-1} \mathcal{L}_z - \mathbf{u}_z = \sigma_{\mathbf{x}}^{-1} (\mathcal{L} - \mathbf{c}')$ for some \mathbf{c}' in the vector span of the lattice \mathcal{L} (because $\mathcal{L}_z - \sigma_{\mathbf{x}}^{-1} \mathbf{u}_z = \mathcal{L} + \mathbf{w}_z - \sigma_{\mathbf{x}} \mathbf{u}_z$ and $\hat{\mathbf{e}}^\top (\mathbf{w}_z - \sigma_{\mathbf{x}} \mathbf{u}_z) = 0$). Then using⁴ Lemma 7 and Lemma 16 that bounds $\sigma_{\mathbf{x}}$ as in the theorem statement, we obtain

$$\begin{aligned} \rho(\sigma_{\mathbf{x}}^{-1} \mathcal{L}_z) &= \rho(\mathbf{u}_z) \cdot \rho_{\sigma_{\mathbf{x}}}(\mathcal{L} - \mathbf{c}') \\ &\in \left[\frac{1 - \delta}{1 + \delta}, 1 \right] \cdot \rho_{\sigma_{\mathbf{x}}}(\mathcal{L}) \cdot \rho(\mathbf{u}_z) \\ &\in \left[\frac{1 - \delta}{1 + \delta}, 1 \right] \cdot \rho_{\sigma_{\mathbf{x}}}(\mathcal{L}) \cdot \rho\left(\frac{z}{\|\hat{\mathbf{e}}\|^2 \sigma_{\mathbf{x}}} \hat{\mathbf{e}}\right) \\ &\in \left[\frac{1 - \delta}{1 + \delta}, 1 \right] \cdot \rho_{\sigma_{\mathbf{x}}}(\mathcal{L}) \cdot \rho_{\|\hat{\mathbf{e}}\| \sigma_{\mathbf{x}}}. \end{aligned}$$

This implies that the statistical distance between $\hat{\mathbf{e}}^\top \mathcal{D}_{\hat{\Lambda} + \hat{\mathbf{c}}, \sigma_{\mathbf{x}}}$ and $\mathcal{D}_{\mathbb{Z}, \|\hat{\mathbf{e}}\| \sigma_{\mathbf{x}}}$ is at most $1 - \frac{1 - \delta}{1 + \delta} \leq 2\delta$. □

Below we state our generalized version of Lemma 3.7 from [BdPMW16].

⁴This is the place where our proof differs from [BdPMW16]. Namely, we use our Lemma 16 to give a different bound on $\sigma_{\mathbf{x}}$.

Lemma 16 (Generalized Lemma 3.7 from [BdPMW16]). *Let $\delta > 0$. Let $\mathcal{L} = \{\mathbf{v} = \hat{\Lambda} : \hat{\mathbf{e}}^\top \cdot \mathbf{v} = 0\}$, where $\hat{\mathbf{e}} = [\mathbf{e}, 1] \in \mathbb{Z}^{m+1}$ is fixed, $\hat{\Lambda} = \hat{\Lambda}_Q^\perp(\mathbf{G}_{L,Q,w}) \times \mathbb{Z}$, $m = w \cdot \ell$ and $L^\ell \leq Q$. Furthermore, let $\mathbf{q}_{L,Q} = [q_i]_{i=1}^\ell$ be the base- L decomposition of Q for $Q < L^\ell$, and $\mathbf{q}_{L,Q} = [\mathbf{0}, L]$ for $Q = L^\ell$. Then we have*

$$\eta_\delta \geq \sqrt{1 + \|\hat{\mathbf{e}}\|_\infty} \cdot \max(\|\mathbf{q}_{L,Q}\|, \sqrt{L^2 + 1}) \cdot C_{\delta,m}.$$

Proof. We use Lemma 5 to bound the smoothing parameter of \mathcal{L} . Since $\hat{\Lambda} = \hat{\Lambda}_Q^\perp(\mathbf{G}_{L,Q,w}^\top) \times \mathbb{Z}$ is of dimension $m + 1$ and \mathcal{L} is a sub-lattice of $\hat{\Lambda}$ made of the vectors that are orthogonal to \mathbf{e} , we have that \mathcal{L} is of dimension m . We thus exhibit m independent short vectors of \mathcal{L} to obtain an upper bound on $\lambda_m(\mathcal{L})$. We first define the matrix

$$\bar{\mathbf{B}} = \begin{bmatrix} L & & & q_1 \\ -1 & L & & q_2 \\ & -1 & \ddots & \vdots \\ & & \ddots & L & q_{\ell-1} \\ & & & -1 & q_\ell \end{bmatrix} \in \mathbb{Z}^{\ell \times \ell},$$

where $\mathbf{q}_{L,Q} = [q_i]_{i=1}^\ell$ is the base- L decomposition of the modulus Q if $Q < L^\ell$, and $q_\ell = L$ and $q_i = 0$ for $i < \ell$ if $Q = L^\ell$. Note that $\bar{\mathbf{B}}$ is a basis for the lattice $\Lambda_Q^\perp(\mathbf{g}_{L,Q})$. The lattice $\hat{\Lambda}$ is then generated by the columns of the matrix:

$$\mathbf{B} = [\mathbf{b}_1 | \dots | \mathbf{b}_{m+1}] = \begin{bmatrix} \mathbf{I}_w \otimes \bar{\mathbf{B}} & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix} \in \mathbb{Z}^{(m+1) \times (m+1)}$$

For $k \leq m$ let $\mathbf{u}_k = \mathbf{b}_k - \mathbf{b}_{m+1} \cdot \hat{\mathbf{e}}^\top \cdot \mathbf{b}_k$. Since $\hat{\mathbf{e}}^\top \cdot \mathbf{b}_{m+1} = 1$ we directly have $\mathbf{e}^\top \cdot \mathbf{u}_k = 0$ and thus $\mathbf{u}_k \in \mathcal{L}$. The vectors $\mathbf{u}_1, \dots, \mathbf{u}_m$ are linearly independent since $\text{span}(\mathbf{u}_1, \dots, \mathbf{u}_m, \mathbf{b}_{m+1}) = \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_m, \mathbf{b}_{m+1}) = \mathbb{R}^{m+1}$ (which comes from the fact that $\bar{\mathbf{B}}$ is a basis of an $(m + 1)$ -dimensional lattice).

We now bound the norm of \mathbf{u}_k . Note that $\mathbf{b}_{m+1} \cdot \hat{\mathbf{e}}^\top \leq \|\hat{\mathbf{e}}\|_\infty$, because $\mathbf{b}_{m+1} = [\mathbf{0}, 1]$. Then we have

$$\begin{aligned} \|\mathbf{u}_k\| &= \|\mathbf{b}_k - \mathbf{b}_{m+1} \cdot \hat{\mathbf{e}}^\top \cdot \mathbf{b}_k\| \\ &\leq \|\mathbf{b}_k + \|\hat{\mathbf{e}}\|_\infty \cdot \mathbf{b}_k\| \\ &= \|(1 + \|\hat{\mathbf{e}}\|_\infty) \mathbf{b}_k\| \\ &= \sqrt{1 + \|\hat{\mathbf{e}}\|_\infty} \cdot \|\mathbf{b}_k\|. \end{aligned}$$

What is left to do is to bound the norm of \mathbf{b}_k . Note that for $k < m + 1$ the vector \mathbf{b}_k has L in its k th position, -1 in position $k + 1$, and 0 in all other positions. Furthermore, the vectors \mathbf{b}_k for $k \equiv 0 \pmod{\log_L(Q)}$ contain the vector $\mathbf{q}_{L,Q}$ and are zero at all other positions. Hence, we can bound the norm by $\|\mathbf{b}_k\| \leq \max(\|\mathbf{q}_{L,Q}\|, \sqrt{L^2 + 1})$. In particular, for $L^{\ell+1} = Q$ the norm of \mathbf{b}_k is bounded by $\sqrt{L^2 + 1}$, while for $L^{\ell+1} > Q$ the bound depends on the decomposition of Q .

To summarize we obtain

$$\lambda_m(\mathcal{L}) \leq \max_{k \leq m} \|\mathbf{u}_k\| \leq \sqrt{1 + \|\hat{\mathbf{e}}\|_\infty} \cdot \max(\|\mathbf{q}_{L,Q}\|, \sqrt{L^2 + 1}).$$

□

Below we give our version of the leftover hash lemma, which is an instantiation of the lemma from [DRS04, DORS08].

Lemma 17 (Leftover Hash Lemma). *Let $\epsilon > 0$ and Q be a odd prime. For any $\mathbf{e} \in \mathbb{Z}_Q^m$ and $\mathbf{C} \in \mathcal{R}^m$, if $\sigma_{\text{rand}} \geq C_{\epsilon, m}$ then*

$$\Delta((\mathbf{A}\mathbf{r}, \mathbf{A}, \mathbf{e}^\top \mathbf{r}), (\mathbf{u}, \mathbf{A}, \mathbf{e}^\top \mathbf{r})) \leq \frac{1}{2} \sqrt{\frac{2^{(n+1)\log(Q)}}{2^{\log(1-\epsilon)+m\log(\sigma_{\text{rand}})}}}$$

where $\mathbf{r} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}^m + \mathbf{c}, \sigma_{\text{rand}}}$, $\mathbf{A} \leftarrow_{\S} \mathbb{Z}_Q^{n \times m}$ and $\mathbf{u} \leftarrow_{\S} \mathbb{Z}_Q^n$.

In the proof we utilize the Generalized Leftover Hash lemma [DRS04, DORS08] but calculate the min-entropy of $\mathbf{r}|\mathbf{e}^\top \mathbf{r}$.

Proof. From Lemma 3 we have that

$$\Delta((\mathbf{A}\mathbf{r}, \mathbf{A}, \mathbf{e}^\top \mathbf{r}), (\mathbf{u}, \mathbf{A}, \mathbf{e}^\top \mathbf{r})) \leq \frac{1}{2} \sqrt{2^{n\log(Q)} \cdot 2^{-\tilde{\mathbf{H}}_\infty(\mathbf{r}|\mathbf{e}^\top \mathbf{r})}},$$

and from Lemma 4 we have that

$$\frac{1}{2} \sqrt{2^{n\log(Q)} \cdot 2^{-\tilde{\mathbf{H}}_\infty(\mathbf{r}|\mathbf{e}^\top \mathbf{r})}} \leq \frac{1}{2} \sqrt{2^{(n+1)\log(Q)} \cdot 2^{-\mathbf{H}_\infty(\mathbf{r})}}$$

because $\tilde{\mathbf{H}}_\infty(\mathbf{r}|\mathbf{e}^\top \mathbf{r}) \geq \mathbf{H}_\infty(\mathbf{r}) - \log(Q)$ since $\mathbf{e}^\top \mathbf{r}$ takes values in \mathbb{Z}_Q (see Lemma 4).

What is left is to analyse the min entropy $\mathbf{H}_\infty(\mathbf{r})$. Note that for any $\mathbf{x} \in \mathbb{Z}^m$ we have that $\rho_{\sigma_{\text{rand}}}(\mathbf{x}) \leq \rho_{\sigma_{\text{rand}}}(\mathbf{0}) = 1$. Furthermore, from Lemma 6 we have that $\rho_{\sigma_{\text{rand}}}(\mathbb{Z}^m + \mathbf{c}) \geq (1 - \epsilon) \frac{\sigma_{\text{rand}}^m}{\det(\mathbb{Z}^m)}$, where $\epsilon > 0$ and assuming that $\sigma_{\text{rand}}^m \geq \eta_\epsilon(\mathbb{Z}^m)$. From the fact that \mathbf{I}_m is the basis of \mathbb{Z}^m we have that $\det(\mathbb{Z}^m) = \det(\mathbf{I}_m) = 1$, and from Lemma 5 we have that $\eta_\epsilon(\mathbb{Z}^m) \leq C_{\epsilon, m}$.

Putting the above together we have that for all $\mathbf{x} \in \mathbb{Z}^m$

$$\mathcal{D}_{\mathbb{Z}^m + \mathbf{c}, \sigma_{\text{rand}}}(\mathbf{x}) \leq \frac{\rho_{\sigma_{\text{rand}}}(\mathbf{x})}{\rho_{\sigma_{\text{rand}}}(\mathbb{Z}^m + \mathbf{c})} \leq \frac{1}{\rho_{\sigma_{\text{rand}}}(\mathbb{Z}^m + \mathbf{c})} \leq \frac{1}{(1 - \epsilon) \cdot \sigma_{\text{rand}}^m}$$

Then from Definition 1 we have

$$\mathbf{H}_\infty(\mathbf{r}) \geq -\log\left(\frac{1}{(1 - \epsilon) \cdot \sigma_{\text{rand}}^m}\right) = \log(1 - \epsilon) + m \log(\sigma_{\text{rand}}).$$

□

4.2 Distribution of Our Randomized Bootstrapping and Circuit Privacy

Below we state and prove the core theorem on the distribution of bootstrapped ciphertexts. Circuit privacy, that we prove at the end of this section, follows nearly immediately from the theorem below.

Theorem 1 (Distribution of the Bootstrap). *Let br be the blind rotation key, $\mathbf{a}_{\text{rot}} \in \mathcal{R}_Q$ a rotation polynomial, and $\mathbf{c} \in \text{LWE}_\sigma(\mathbf{s}, m)$ a LWE sample as defined in the Bootstrap algorithm in Figure 2. Assume that \mathbf{a}_{rot} is such that $f(m) = (\mathbf{a}_{\text{rot}} \cdot X^{\text{Phase}(\mathbf{c}_{\text{in}})})[1]$ where \mathbf{c}_{in} is the LWE sample obtained at Step 2 of the Bootstrap algorithm. Let \mathbf{c}_{out} be the LWE sample returned by the Bootstrap algorithm for $\text{ver} = \text{simul}$ and Gaussian parameters σ_{rand} and $\sigma_{\mathbf{x}}$ where the Gaussian sampling algorithm $\mathbf{G}_{\text{simul}}^{-1}$ is as in Lemma 8. Assume that $\sigma_{\text{rand}} \geq C_{\epsilon, h}$ and*

$$\sigma_{\mathbf{x}} \geq \sqrt{1 + B_{\text{br}}} \cdot \max(\|\mathbf{q}_{\text{L-br}, Q}\|, \sqrt{L_{\text{br}}^2 + 1}) \cdot C_{\delta, 2 \cdot n \cdot N \cdot \ell_{\text{br}}},$$

where B_{br} is a bound on the infinity norm of the noise terms in the blind rotation key br . Then we have

$$\Delta(\mathbf{c}_{\text{out}}, \mathbf{c}_{\text{fresh}}) \leq \max\left(2\delta, \frac{1}{2} \sqrt{\frac{2^{(N+1)\log(Q)}}{2^{\log(1-\epsilon)+h\log(\sigma_{\text{rand}})}}}\right),$$

where $\mathbf{c}_{\text{fresh}} = [\mathbf{a}_{\text{fresh}}, \mathbf{b}_{\text{fresh}}]$, $\mathbf{b}_{\text{fresh}} = \langle \mathbf{a}_{\text{fresh}}, \mathbf{s}' \rangle + f(m) + e_{\text{rand}} + e_{\text{out}}$, $e_{\text{out}} \leftarrow_{\$} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}} \cdot \sqrt{1+|\mathbf{e}|}}$, $e_{\text{rand}} \leftarrow_{\$} \tilde{\mathbf{e}}^{\top} \cdot \mathbf{r}$, $\mathbf{r} \leftarrow_{\$} \mathcal{D}_{\mathbb{Z}^h, \sigma_{\text{rand}}}$ and where $\mathbf{e} \in \mathbb{Z}^{2 \cdot n \cdot N \cdot \ell_{\text{br}}}$ is the vector of error coefficients in the blind rotation key, and $\tilde{\mathbf{e}} \in \mathbb{Z}^h$ are the error terms in the vector of LWE samples \mathbf{v} .

Proof. The proof consists of two parts. First we analyze the LWE sample that is extracted after blind rotation. In particular, we give a concise representation of the final noise term. Furthermore, we show that each noise coefficient and decomposition term of the randomized decomposition appears only once in the final noise term. The second part of the proof consists of a hybrid argument, where we argue step-by-step that the distribution of the extracted and “masked” LWE sample is statistically close to a “freshly” sampled LWE sample of the same message.

Below we give the first part of the proof. But to further tame complexity we split this part in three more sub-parts. First we analyze a single external product, then a single MUX gate and we finalize this part with blind rotation and extraction.

Single External Product. First let us remind that for $j \in [\ell_{\text{br}}]$ we have

$$\begin{aligned} \mathbf{C}_G[* , j] &= \text{RLWE}_{\sigma}(\mathbf{s}, \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{j-1}) \quad \text{and} \\ \mathbf{C}_G[* , j + \ell_{\text{br}}] &= \text{RLWE}_{\sigma}(\mathbf{s}, -\mathbf{s} \cdot \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{j-1}). \end{aligned}$$

Denote $\mathbf{C}_G = [\mathbf{a}_j, \mathbf{b}_j]_{j=1}^{2\ell_{\text{br}}}$, and $\mathbf{d} = [\mathbf{a}_d, \mathbf{b}_d]$ where $\mathbf{b}_d = \mathbf{a}_d \cdot \mathbf{s} + \mathbf{e}_d + \mathbf{m}_d$. We analyze the sample $\mathbf{c}_{\text{prod}} \leftarrow \text{extProd}_{\text{simul}}(\mathbf{C}_G, \mathbf{d})$. Then in Steps 1 and 2 of the $\text{extProd}_{\text{simul}}$ algorithm we compute $\mathbf{c}_{\text{prod}} = \mathbf{C}_G \cdot \mathbf{G}_{\text{simul}}^{-1}(\mathbf{d}, \mathbf{L}_{\text{br}}) = [\mathbf{a}_{\text{prod}}, \mathbf{b}_{\text{prod}}]$. Let us denote the vector $[\mathbf{r}_j]_{j=1}^{2\ell_{\text{br}}} = \mathbf{G}_{\text{simul}}^{-1}(\mathbf{d}, \mathbf{L}_{\text{br}})$. We can write

$$\mathbf{a}_{\text{prod}} = \sum_{j=1}^{2\ell_{\text{br}}} \mathbf{a}_j \cdot \mathbf{r}_j.$$

Furthermore, we can write

$$\mathbf{b}_{\text{prod}} = \mathbf{a}_{\text{prod}} \cdot \mathbf{s} + \mathbf{m}_G \cdot \mathbf{e}_d + \hat{\mathbf{e}} + \mathbf{m}_G \cdot \mathbf{m}_g, \quad (1)$$

where $\hat{\mathbf{e}} = \sum_{j=1}^{2\ell_{\text{br}}} \mathbf{e}_j \cdot \mathbf{r}_j$. Equation 1 holds because we have

$$\begin{aligned} \sum_{j=1}^{2\ell_{\text{br}}} \mathbf{b}_j \cdot \mathbf{r}_j &= \mathbf{a}_{\text{prod}} \cdot \mathbf{s} + \sum_{j=1}^{\ell_{\text{br}}} (\mathbf{e}_j + \mathbf{m}_G \cdot \mathbf{L}^{j-1}) \cdot \mathbf{r}_j + \sum_{j=1}^{\ell_{\text{br}}} (\mathbf{e}_{j+\ell_{\text{br}}} - \mathbf{s} \cdot \mathbf{m}_G \cdot \mathbf{L}^{j-1}) \cdot \mathbf{r}_{j+\ell_{\text{br}}} \\ &= \mathbf{a}_{\text{prod}} \cdot \mathbf{s} + \sum_{j=1}^{2\ell_{\text{br}}} \mathbf{r}_j \cdot \mathbf{e}_j + \mathbf{m}_G \cdot \left(\sum_{j=1}^{\ell_{\text{br}}} \mathbf{r}_j \cdot \mathbf{L}^{j-1} - \mathbf{s} \cdot \sum_{j=\ell_{\text{br}}+1}^{2\ell_{\text{br}}} \mathbf{r}_j \cdot \mathbf{L}^{j-1-\ell_{\text{br}}} \right) \end{aligned}$$

and in particular from the properties of the Gaussian sampling algorithm (Lemma 8) we have that

$$\mathbf{m}_G \cdot \left(\sum_{j=1}^{\ell_{\text{br}}} \mathbf{r}_j \cdot \mathbf{L}^{j-1} - \mathbf{s} \cdot \sum_{j=\ell_{\text{br}}+1}^{2\ell_{\text{br}}} \mathbf{r}_j \cdot \mathbf{L}^{j-1-\ell_{\text{br}}} \right) = \mathbf{m}_G(\mathbf{b}_d - \mathbf{a}_d \cdot \mathbf{s}) = \mathbf{m}_G(\mathbf{e}_d + \mathbf{m}_d).$$

Single MUX Gate. Let us analyze a single execution of the MUX gate $\mathbf{c}_{\text{out}} \leftarrow \text{Mux}(\mathbf{C}_G, \mathbf{d}, \mathbf{h})$, where \mathbf{C}_G and \mathbf{d} are RGSW and RLWE samples as above, and $\mathbf{h} \in \text{RLWE}_\sigma(\mathbf{s}, \mathbf{m}_h) = [\mathbf{a}_h, \mathbf{b}_h]$. Since $\mathbf{c}_{\text{out}} \leftarrow \text{extProd}_{\text{simul}}(\mathbf{C}_G, \mathbf{d} - \mathbf{h}) + \mathbf{h}$, we can write $\mathbf{c}_{\text{out}} = [\mathbf{a}_{\text{prod}}, \mathbf{b}_{\text{prod}}] + [\mathbf{a}_h, \mathbf{b}_h]$. Remind that the message encoded in \mathbf{C}_G is a single bit $m_G \in \{0, 1\}$. That is, we are only interested in the special case where the RGSW message is a single bit. We have two cases:

- The case with $m_G = 0$. In this case we can write $\mathbf{b}_{\text{prod}} = \mathbf{a}_{\text{prod}} \cdot \mathbf{s} + \boldsymbol{\epsilon}$. In other words, the error from $\mathbf{d} - \mathbf{h}$ cancels out. And we have that $\mathbf{b}_{\text{out}} = \mathbf{b}_{\text{prod}} + \mathbf{b}_h = \mathbf{a}_{\text{out}} \cdot \mathbf{s} + \boldsymbol{\epsilon} + \mathbf{m}_h + \boldsymbol{\epsilon}_h$.
- The case with $m_G = 1$. In this case we can write $\mathbf{b}_{\text{prod}} = \mathbf{a}_{\text{prod}} \cdot \mathbf{s} + \boldsymbol{\epsilon} + \mathbf{m}_d - \mathbf{m}_h + \boldsymbol{\epsilon}_d - \boldsymbol{\epsilon}_h$. And we have that $\mathbf{b}_{\text{out}} = \mathbf{b}_{\text{prod}} + \mathbf{b}_h = \mathbf{a}_{\text{out}} \cdot \mathbf{s} + \mathbf{m}_d + \boldsymbol{\epsilon} + \boldsymbol{\epsilon}_d$.

Finally, note that in blind rotation we have $\mathbf{d} = \mathbf{h} \cdot X^l \in \mathcal{R}_Q^2$ for some $l \in \mathbb{Z}_{2N}$. That is, the two ring LWE samples are negacyclic rotations of one another. This means that $\boldsymbol{\epsilon}_d = \boldsymbol{\epsilon}_h \cdot X^l \in \mathcal{R}_Q$.

Blind Rotation and Extraction. First we set the accumulator to $\mathbf{c}_{\text{acc},0} = [0, \widehat{\mathbf{a}}_{\text{rot}} \cdot X^b]$. Note that the first accumulator is special because its noise term is zero. Let us denote the error term that is added in the i th iteration of the blind rotation loop by $\widehat{\boldsymbol{\epsilon}}_i$. Particularly, this noise term is $\widehat{\boldsymbol{\epsilon}}_i = \sum_{j=1}^{2\ell_{\text{br}}} \boldsymbol{\epsilon}_{i,j} \cdot \mathbf{r}_{i,j}$, where $\boldsymbol{\epsilon}_{i,j}$ is the noise term of the blind rotation keys, $\mathbf{r}_{i,j}$ is the component from the randomized decomposition algorithm. At the i th iteration we run a homomorphic MUX gate that multiplies the input RLWE sample's noise and message by $X^{-\mathbf{a}^{[i] \cdot \mathbf{s}^{[i]}}$, and adds a new noise term $\widehat{\boldsymbol{\epsilon}}_i$. Remind that the noise term of $\mathbf{c}_{\text{acc},0}$ is zero, thus at iteration $i = 1$, the resulting RLWE sample $\mathbf{c}_{\text{acc},1}$ has noise term $\widehat{\boldsymbol{\epsilon}}_1$. Then $\mathbf{c}_{\text{acc},1}$ has noise term $\widehat{\boldsymbol{\epsilon}}_1 \cdot X^{-\mathbf{a}^{[2] \cdot \mathbf{s}^{[2]}} + \widehat{\boldsymbol{\epsilon}}_2$. Finally after n iterations we have

$$\text{Error}(\mathbf{c}_{\text{acc}}) = \text{Error}(\mathbf{c}_{\text{acc},n}) = \sum_{i=1}^n \widehat{\boldsymbol{\epsilon}}_i \cdot X^{\sum_{j=i+1}^n -\mathbf{a}^{[j] \cdot \mathbf{s}^{[j]}}} \quad (2)$$

and the message is $\mathbf{a}_{\text{rot}} \cdot X^{\text{Phase}(\mathbf{c})}$.

Let us denote $\mathbf{c}_{\text{acc}} = [\mathbf{a}_{\text{acc}}, \mathbf{b}_{\text{acc}}]$ and the extracted LWE sample $\mathbf{c}_{\text{ext}} = [\mathbf{a}_{\text{ext}}, \mathbf{b}_{\text{ext}}]$. Note that $\mathbf{a}_{\text{ext}} \in \mathbb{Z}_Q^N$ and $\mathbf{b}_{\text{ext}} = \mathbf{b}_{\text{acc}}[1] = \langle \mathbf{a}_{\text{acc}}, \mathbf{s}' \rangle + \mathbf{a}_{\text{rot}} \cdot X^{\text{Phase}(\mathbf{c})}[1] + (\sum_{i=1}^n \widehat{\boldsymbol{\epsilon}}_i \cdot X^{\sum_{j=i+1}^n -\mathbf{a}^{[j] \cdot \mathbf{s}^{[j]}}})[1]$. In particular, note that

$$\begin{aligned} \text{Error}(\mathbf{c}_{\text{ext}}) &= \left(\sum_{i=1}^n \widehat{\boldsymbol{\epsilon}}_i \cdot X^{\sum_{j=i+1}^n -\mathbf{a}^{[j] \cdot \mathbf{s}^{[j]}} \right)[1] \\ &= \sum_{i=1}^n \sum_{j=1}^{2\ell_{\text{br}}} (\boldsymbol{\epsilon}_{i,j} \cdot X^{\sum_{j=i+1}^n -\mathbf{a}^{[j] \cdot \mathbf{s}^{[j]}} \cdot \mathbf{r}_{i,j})[1] \\ &= \sum_{i=1}^n \sum_{j=1}^{2\ell_{\text{br}}} \sum_{k=1}^N \mathbf{e}_{i,j}[k] \cdot \mathbf{r}_{i,j}[k] \end{aligned}$$

where $\mathbf{e}_{i,j}$ is a vector of discrete Gaussian random variables centered at zero of parameter σ_{br} . Note that in the ring \mathcal{R}_Q and assuming the coefficients of $\boldsymbol{\epsilon}_{i,j}$ are centered at zero, we have that $\boldsymbol{\epsilon}_{i,j} \cdot X^{\sum_{j=i+1}^n -\mathbf{a}^{[j] \cdot \mathbf{s}^{[j]}} = \boldsymbol{\epsilon}'_{i,j}$ rotates the coefficients of $\boldsymbol{\epsilon}_{i,j}$ negacyclicly, meaning that $\boldsymbol{\epsilon}'_{i,j}$ has the same distribution as $\boldsymbol{\epsilon}_{i,j}$. Similarly, we have that the constant coefficient of $\boldsymbol{\epsilon}'_{i,j} \cdot \mathbf{r}_{i,j}$ in \mathcal{R}_Q is $(\boldsymbol{\epsilon}'_{i,j} \cdot \mathbf{r}_{i,j})[1] = \mathbf{r}_{i,j}[1] \cdot \boldsymbol{\epsilon}'_{i,j}[1] + \sum_{k=2}^N -\boldsymbol{\epsilon}'_{i,j}[N-k+1] \cdot \mathbf{r}_{i,j}[k]$. Hence we can write $\mathbf{e}_{i,j}[k] = -\boldsymbol{\epsilon}'_{i,j}[N-k+1]$ for $k \in [2, N]$ and $\mathbf{e}_{i,j}[1] = \boldsymbol{\epsilon}'_{i,j}[1]$. Therefore $\mathbf{e}_{i,j}[k]$ is from the discrete Gaussian distribution of the same parameter as $\boldsymbol{\epsilon}_{i,j}$ given that the distribution of the coefficients are centered at zero.

Distribution of the Extracted LWE Sample. Now we are ready to argue that the $\mathbf{c}_{\text{out}} = \mathbf{c}_{\text{ext}} + \mathbf{r}^\top \cdot \mathbf{v} + y$ sample is statistically close to a LWE sample of the same message that is independent of the input ciphertext. The proof is due to the following hybrid argument.

Hybrid 0. In this hybrid the sample is as in the original scheme. Specifically, we have $\mathbf{c}_{\text{out}} \leftarrow \mathbf{c}_{\text{ext}} + \mathbf{c}_{\text{rand}} + y$, where $\mathbf{c}_{\text{rand}} \leftarrow \mathbf{v}^\top \cdot \mathbf{r}$ with $\mathbf{r} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}^h, \sigma_{\text{rand}}}$, $y \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}}}$ and \mathbf{v} is a vector of size h of LWE samples of zero with noise parameter σ_R .

Hybrid 1. As Hybrid 0, but we set the message in \mathbf{c}_{out} to $f(m)$ instead of $\mathbf{a}_{\text{rot}} \cdot X^{\text{Phase}(\mathbf{c})}[1]$. Assuming that $\text{Error}(\mathbf{c}) \leq \frac{N}{t}$ and \mathbf{a}_{rot} is such that $f(m) = \mathbf{a}_{\text{rot}} \cdot X^{\text{Phase}(\mathbf{c})}[1]$ this change is only syntactical by correctness of the bootstrapping algorithm.

Hybrid 2. This hybrid is as Hybrid 1, but instead of computing $\mathbf{c}_{\text{rand}} \leftarrow \mathbf{r}^\top \cdot \mathbf{v} = [\mathbf{a}_{\text{rand}}, b_{\text{rand}}]$, we take $\mathbf{c}_{\text{rand}} = [\mathbf{a}_{\text{rand}}, b_{\text{rand}}]$ to be a fresh LWE encryption of zero. In particular, we take $\mathbf{a}_{\text{rand}} \leftarrow_{\S} \mathbb{Z}_Q^N$ from the uniform distribution and take $b_{\text{rand}} = \langle \mathbf{a}_{\text{rand}}, \mathbf{s}' \rangle + e_{\text{rand}}$ where $e_{\text{rand}} \leftarrow_{\S} \tilde{\mathbf{e}}^\top \cdot \mathbf{r}$ with $\mathbf{r} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}^h, \sigma_{\text{rand}}}$ and where $\tilde{\mathbf{e}}$ denotes the errors of the LWE samples in \mathbf{v} .

Claim. *Given that $\sigma_{\text{rand}} \geq C_{\epsilon, h}$ the statistical distance between Hybrid 2 and Hybrid 1 is at most*

$$\frac{1}{2} \sqrt{\frac{2^{(N+1) \log(Q)}}{2^{\log(1-\epsilon)+h \log(\sigma_{\text{rand}})}}}$$

for some $\epsilon > 0$.

Proof. Denote $\mathbf{v} = [a_{v,i}, b_{v,i}]_{i=1}^h$ where $b_{v,i} = \langle \mathbf{a}_{v,i}, \mathbf{s}' \rangle + e_{v,i}$. Let $\tilde{\mathbf{b}} = [b_{v,i}]_{i=1}^h$ and $\tilde{\mathbf{e}} = [e_{v,i}]_{i=1}^h$. We can write $\mathbf{A} = [\mathbf{a}_{v,1}, \dots, \mathbf{a}_{v,h}] \in \mathbb{Z}_Q^{N \times h}$, $\mathbf{a}_{\text{rand}} \leftarrow \tilde{\mathbf{A}} \cdot \mathbf{r}$ and $b_{\text{rand}} \leftarrow \tilde{\mathbf{b}}^\top \cdot \mathbf{r} = \langle \mathbf{a}_{\text{rand}}, \mathbf{s}' \rangle + \tilde{\mathbf{e}}^\top \cdot \mathbf{r}$.

Now it is easy to see that the rest of the proof follows directly by applying Lemma 17. Note that the notation in Lemma 17 is mostly already in place, except that we set the m from the lemma to h and the n from the lemma to N . Furthermore, the matrix \mathbf{A} from the lemma is the matrix $\tilde{\mathbf{A}}$ in this hybrid and the error \mathbf{e} from the lemma is the error $\tilde{\mathbf{e}}$ in this hybrid. Note that Lemma 17 requires Q to be an odd prime. In Remark 1 we discuss how we handle a non-prime Q . Finally, the uniform vector \mathbf{u} from the lemma corresponds to \mathbf{a}_{rand} . \square

Hybrid 3. This hybrid is as Hybrid 2, except that we choose \mathbf{a}_{out} from the uniform distribution. Note that both hybrids are in fact identical as from Hybrid 1, we have that \mathbf{a}_{rand} is sampled from the uniform distribution over \mathbb{Z}_Q^N , and we have $\mathbf{a}_{\text{out}} = \mathbf{a}_{\text{ext}} + \mathbf{a}_{\text{rand}} \in \mathbb{Z}_Q^N$. This hybrid is only a syntactic change.

Hybrid 4. This hybrid is as Hybrid 3 except that we compute $\mathbf{b}_{\text{out}} = \langle \mathbf{a}_{\text{out}}, \mathbf{s} \rangle + f(m) + e_{\text{rand}} + e_{\text{out}}$, where $e_{\text{out}} \leftarrow_{\S} \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{x}} \cdot \sqrt{1+\|\mathbf{e}\|}}$. In particular, the noise term is independent from the ciphertext \mathbf{c} and the blind rotation key \mathbf{br} .

Claim. *Given that*

$$\sigma_{\mathbf{x}} \geq \sqrt{1 + B_{\text{br}}} \cdot \max(\|\mathbf{q}_{\text{L-br}, Q}\|, \sqrt{L_{\text{br}}^2 + 1}) \cdot C_{\delta, 2 \cdot n \cdot N \cdot \ell_{\text{br}}}$$

the statistical distance between Hybrid 3 and Hybrid 4 is at most 2δ for some $\delta > 0$. The notation in Lemma 15 is mostly already in place, except we set $m = 2 \cdot n \cdot N \cdot \ell_{\text{br}}$ and $L = L_{\text{br}}$.

Proof. Note that we have

$$\text{Error}(\mathbf{c}_{\text{ext}}) = \sum_{i=1}^n \sum_{j=1}^{2\ell_{\text{br}}} \sum_{k=1}^N \mathbf{e}_{i,j}[k] \cdot \mathbf{r}_{i,j}[k] = \sum_{i=1}^n \sum_{k=1}^N \sum_{j=1}^{2\ell_{\text{br}}} \mathbf{e}_{i,j}[k] \cdot \mathbf{r}_{i,j}[k].$$

We will group the terms $\mathbf{e}_{i,j}[k] \cdot \mathbf{r}_{i,j}[k]$ into vectors by the j iterator. We write $\widehat{\mathbf{e}}_{i,k} = [\mathbf{e}_{i,j}[k]]_{j=1}^{2\ell_{\text{br}}}$ and $\widehat{\mathbf{x}}_{i,k} = [\mathbf{r}_{i,j}[k]]_{j=1}^{2\ell_{\text{br}}}$. Note that $\widehat{\mathbf{x}}_{i,k} \in \mathcal{D}_{\Lambda_Q^\perp(\mathbf{G}_{\text{Lbr},2}) + \mathbf{G}_{\text{det}}^{-1}(\mathbf{c}_{\text{acc},i}[k]), \sigma_{\mathbf{x}}}$. In other words $\widehat{\mathbf{x}}_{i,k}$ is the Gaussian sampling of the k th coefficient (where $k \in [2\ell_{\text{br}}]$, meaning that we take the concatenation of the two polynomials in the accumulator $\mathbf{c}_{\text{acc},i}$) in the i th iteration of the blind rotation algorithm. Now we can write

$$\text{Error}(\mathbf{c}_{\text{out}}) = e_{\text{rand}} + \text{Error}(\mathbf{c}_{\text{ext}}) + y = e_{\text{rand}} + y + \sum_{i=1}^n \sum_{k=1}^N \widehat{\mathbf{e}}_{i,k}^\top \cdot \widehat{\mathbf{x}}_{i,k}.$$

We can further represent $\text{Error}(\mathbf{c}_{\text{ext}})$ as a product $\mathbf{x}^\top \cdot \mathbf{e}$ of two vectors $\mathbf{e} = [\widehat{\mathbf{e}}_{i,k}^\top]_{i=1,k=1}^{n,N}$ and $\mathbf{x} = [\widehat{\mathbf{x}}_{i,k}]_{i=1,k=1}^{n,N}$. Note that

$$\mathbf{x} \in \mathcal{D}_{\Lambda_Q^\perp(\mathbf{G}_{\text{Lbr},Q,2 \cdot n \cdot N}) + \mathbf{G}_{\text{det}}^{-1}([\mathbf{c}_{\text{acc},i}[k]]_{i=1,k=1}^{n,N}), \sigma_{\mathbf{x}}}$$

since it is just a concatenation of $n \cdot N$ vectors from $\mathcal{D}_{\Lambda_Q^\perp(\mathbf{G}_{\text{Lbr},2}) + \mathbf{G}_{\text{det}}^{-1}(\mathbf{c}_{\text{acc},i}[k]), \sigma_{\mathbf{x}}}$. Finally, note that \mathbf{a}_{out} and the message are independent of $\text{Error}(\mathbf{c}_{\text{ext}})$. Therefore we can apply Lemma 15 to $\mathbf{x}^\top \cdot \mathbf{e} + y$ with $m = 2 \cdot n \cdot N \cdot \ell_{\text{br}}$ and $\mathbf{L} = \mathbf{L}_{\text{br}}$. \square

Finally, we have that \mathbf{c}_{out} is distributed as in the theorem statement, and is in particular independent of the input ciphertext \mathbf{c} and bootstrapping key br . \square

Proof of Circuit Privacy. Recall that according to Definition 4, to prove circuit privacy we have to show a simulator that, on input ek and $m_{\text{out}} = \mathcal{C}(m_1, \dots, m_n)$ outputs a ciphertexts of m_{out} that is distributed as an output of the `Eval` algorithm when evaluating \mathcal{C} on encryptions of m_1, \dots, m_n . We build such simulator by sampling an encryption of zero, bootstrapping in `simul` mode and adding m_{out} to the resulting ciphertext. Circuit privacy then follows from Theorem 1.

Theorem 2. *Let $(\text{ek}, \text{sk}) \leftarrow \text{Setup}(\lambda)$. Let \mathcal{C} be a polynomial size circuit and $\text{ct}_{\text{out}} \leftarrow \text{Eval}(\text{ek}, [\text{ct}_i]_{i=1}^n, \mathcal{C})$, where $\text{ct}_i \leftarrow \text{Dec}(\text{sk}, m_i)$ and `Eval` is as described in Section 3, where the bootstrapping algorithm for the output gate is set to `ver = simul`. If the parameters of the FHE scheme are chosen such that $\Delta(\mathbf{c}_{\text{out}}, \mathbf{c}_{\text{fresh}}) \leq \text{negl}(\lambda)$, where $\mathbf{c}_{\text{fresh}}$ is as in Theorem 1, then the evaluation process is circuit private.*

Proof. To show circuit privacy, we need to show a simulator `Sim` that gets as input ek and m_{out} . The proof follows nearly immediately from Theorem 1. Denote as \mathbf{c}_{out} the ciphertext returned by `Eval`. Recall that \mathbf{c}_{out} is distributed as given by Lemma 1, because `Eval` ends with an invocation of `Bootstrap` in simulation mode. Set $\mathbf{c} = \mathbf{0}$ and $\mathbf{a}_{\text{rot}} = \mathbf{o}$. The simulator runs and outputs

$$\mathbf{c}_{\text{fresh}} \leftarrow \text{Bootstrap}_{\text{simul}}(\text{br}, \text{ksK}, \mathbf{c}, \mathbf{a}_{\text{rot}}, \mathbf{v}, \sigma_{\text{rand}}, \sigma_{\mathbf{x}}) + [\mathbf{0}, m_{\text{out}}].$$

Denote $\mathbf{c}_{\text{fresh}} = [\mathbf{a}_{\text{fresh}}, b_{\text{fresh}}]$. From Lemma 1 we have that \mathbf{c}_{out} is distributed as in `Eval`. Namely, $\mathbf{a}_{\text{fresh}}$ is statistically close to uniform, $b_{\text{fresh}} = \langle \mathbf{a}_{\text{fresh}}, \mathbf{s}' \rangle + m_{\text{out}} + e_{\text{rand}} + e_{\text{out}}$ with $e_{\text{rand}} + e_{\text{out}}$ distributed as in Lemma 1. Hence the $\mathbf{c}_{\text{fresh}}$ is statistically close to \mathbf{c}_{out} , and in particular independent from the circuit \mathcal{C} . \square

5 Parameters, Implementation, and Experiments

In this section, we discuss our parameter choices, implementation, and experiments for our method as well as for the washing machine method by Ducas and Stehlé [DS16]. Remind

Table 1: Parameter Sets. We list base-two logarithms of the parameters. Recall that Q denotes the coefficient modulus, \mathbf{L}_{Bk} is the decomposition basis of the RLWE samples in the bootstrapping key, σ_{Ksk} is the standard deviation for the key switching key, σ_{rand} is the standard deviation and h is the size of the LWE vector used for generating fresh LWE samples, and $\sigma_{\mathbf{x}}$ is the standard deviation for the randomized blind rotation. For all sets we set $n = 912$, $N = 2^{11}$, $\mathbf{L}_{\text{Ksk}} = 2^7$ and $\sigma_{\text{br}} = 3.2$. Note that the column $\mathbf{L}_{\text{Bk-simul}}$ gives the decomposition base for $\text{ver} = \text{simul}$ and $\mathbf{L}_{\text{Bk-det}}$ for $\text{ver} = \text{det}$. The column U_x refers to the uniform distribution interval for the noise flooding in DS-WM. The column Δ refers to the statistical distance from a random ciphertext after a single bootstrapping operation. Consequently, for DS-WM-Int and DS-WM-Double, we have to run the bootstrapping 5 and 12 times, respectively.

Alg.	Param. Q	simul		det		σ_{Ksk}	$\sigma_{\mathbf{x}}$	U_x	σ_{rand}	h	Δ
		$\mathbf{L}_{\text{Bk}}, \ell_{\text{Bk}}$	$\mathbf{L}_{\text{Bk}}, \ell_{\text{Bk}}$	$\mathbf{L}_{\text{Bk}}, \ell_{\text{Bk}}$	$\mathbf{L}_{\text{Bk}}, \ell_{\text{Bk}}$						
Ours-Int	48	8, 6	24, 2	26	17.7	12.8	-	12.3	80		
DS-WM-Int	48	8, 6	24, 2	26	-	41.2	12.3	12.8	16		
Ours-Double	36	4, 9	12, 3	14	8.9	-	-	11.7	80		
DS-WM-Double	36	4, 9	12, 3	14	-	30.8	21.9	11.7	6.6		

that Ducas and Stehlé [DS16] left finding correct parameters as an open problem. In this section, we address this problem and rule out the possibility of instantiating FHEW/TFHE negligible statistical security over low-degree rings that are often used for efficiency. We give a comparison of both methods when parameters are targeted towards an implementation over integers modulo an NTT-friendly prime number and over integers modulo a power of two represented as double-precision floating point numbers.

Remark 1. In the case of a power of two modulus Q , we slightly modify the scheme and choose the LWE samples with respect to a prime modulus that is slightly larger than Q . Then, after Step 8 of Algorithm 3 we switch the modulus to Q . This change is made because the leftover hash lemma (Lemma 17) requires a universal hash function, which is satisfied if the modulus is prime but not when the modulus is a power of two.

5.1 Parameters.

We choose our parameter sets to target 128-bit security for the (R)LWE samples and 80-bits statistical security when running the bootstrapping in simul mode. Remind that, in contrast to computational security, the advantage in breaking a property for statistical security does not increase with advances in high-performance computing. Since we measure the distance between samples, the ability to distinguish two distributions depends on the number of samples given. The parameters are listed in Table 1. We estimate the (R)LWE security using the latest commit of the LWE estimator [APS15]. We wrote a python script that we published alongside our source code to estimate the statistical security. We choose similar parameters to make a good comparison between our method and the Ducas-Stehlé washing machine method [DS16] that we refer to as DS-WM. For completeness we recall the relevant lemmas from [DS16] that we used for our estimations in Appendix A. For all parameters we chose the same ring dimension $N = 2^{11}$. The strategy is to choose the highest modulus such that: (1) the RLWE problem remains 128-bit secure according to the LWE estimator [APS15], (2) the modulus is below 50-bits for Ours-Int and DS-WM-Int to allow for faster multiplication of ring elements, and (3) computing convolutions does not introduce significant numerical errors (see Appendix B for estimations). Furthermore, we choose the LWE parameters for the key switching key and the masking key \mathbf{v} , the

same for all solutions. Then we choose the decomposition bases. We chose the highest decomposition base that: (1) gives us required correctness, (2) maximizes $L_{Bk}\text{-simul}$, (3) $L_{Bk}\text{-det} = L_{Bk}\text{-simul}^k$ for some $k \geq 1$ for deterministic gadget decomposition gives us the required correctness as well. The last condition allows us to significantly optimize computation since deterministic mode requires roughly $k \times$ fewer ring multiplications. Finally, we compute the noise parameters for the Gaussian sampling and noise flooding given the desired security level and all other parameters.

As we can read from Table 1, we managed to find parameters where Ours-Int method needs $7 = \ell_{br} + 1$ ring multiplications per gadget multiplication in `simul` mode, whereas DS-WM-Int requires $11 = \ell_{br} + 1$. We stress that the noise flooding method is already incorrect for a decomposition base $L_{br} = 2^6$ meaning that this is the best base choice for performance. Nevertheless, for the flooding set, we can use a similar decomposition basis in `det` mode as in our set resulting in the same efficiency for deterministic bootstrapping. This parameter choice allows us to compare both methods better as we can now focus solely on discussing the differences between the sets in `simul` mode. For the `det` mode efficiency of the sets is the same and correctness is very similar. In the case, where the modulus is a power of L_{br} and implementation uses double precision floating point arithmetic to compute polynomial multiplication, we need a much lower modulus, and decomposition bases. In this case we found sets where all decomposition parameters are the same for our method on for DS-WM. In Table 2, we list the probabilities of having an error while bootstrapping. The errors are given as base-two logarithms for readability. As we may see, our method has a different characteristic when it comes to correctness than the DS-WM. First observe that correctness of \mathbf{c}_{in} is lower than correctness of \mathbf{c}_{out} . This is due to the modulus switching to a much smaller modulus $2N$ and the rounding error. Then note that when increasing the message space, our method is still correct for \mathbf{c}_{out} , while DS-WM’s correctness collapses already at $t = 5$. This is the consequence of needing to run the bootstrapping step numerous times to sanitize a ciphertext. However, since our method requires only a single bootstrapping invocation, we can output ciphertexts of larger precision. Finally, we describe correctness issues when instantiating the system over over degree $N = 2^{10}$ rings in Appendix C.

5.2 Implementation and Performance.

We implemented the schemes in C++11 and tested it on a machine with 11th Gen Intel(R) Core(TM) i7-11850H 2.50GHz processor that supports AVX2 and AVX-512 instructions. The timing results and the size of the evaluation keys for the bootstrapping algorithms are given in Figure 3. To implement negacyclic convolutions for Ours-Int and DS-WM-Int, we used the Intel Hexl library [BKS⁺21]. The library gives a high-performance implementation of Number Theoretic Transforms optimized for the ring $\mathbb{Z}_Q[X]/(X^N + 1)$ and takes full advantage of Intel AVX instructions. To compute the negacyclic convolutions for Ours-Double and DS-WM-Double we use the FFTW library [FJ21] that implements fast Fourier transforms on IEEE-754 double precision floating point arithmetic. For the Gaussian sampling algorithm, we implemented two methods. For the case where the modulus is a power of L_{Bk} we implement the simple and very efficient Gaussian sampler from [MP12]. For general moduli, which is the case for Our-Int and DM-WM-Int, we implement a version of the method by Genise and Micciancio [GM18, CDCG⁺18]. Both samplers are instantiated with either the Karney method [Kar16] to sample for the exact discrete Gaussian distribution, or (somewhat heuristically) we use the Box-Muller transform [BM58] with rounding to the closest integer. The implementation using the Box-Muller transform serves mostly for comparison and to showcase potential speedups. Our proofs require a discrete Gaussian sampler for simplicity. Nevertheless, we note that in some cases, proofs can be generalized [HLS18] to support rounded continuous Gaussian distributions while preserving security. We leave such generalizations as future work.

Table 2: Correctness Estimates. We give the probability of failure to correctly decrypt c_{out} and c_{in} for a given message space t . We give the correctness estimates as base-two logarithm. We mark failure probabilities below 2^{-80} with green, and above that threshold with red.

Ours-Int					Ducas-Stehlé: DS-WM-Int				
t	det		simul		t	det		simul	
	c_{out}	c_{in}	c_{out}	c_{in}		c_{out}	c_{in}	c_{out}	c_{in}
4	-771	-239	-494	-154	4	-771	-239	-198	-198
5	-495	-99	-317	-64	5	-495	-99	-69	-69
6	-345	-40	-221	-26	6	-345	-40	-20	-20
10	-126	0.90	-82	0.82	10	-126	0	0	0
11	-105	0.99	-68	0.97	11	-105	0	0	0

Ours-Double					Ducas-Stehlé: DS-WM-Double				
t	det		simul		t	det		simul	
	c_{out}	c_{in}	c_{out}	c_{in}		c_{out}	c_{in}	c_{out}	c_{in}
4	-841	-260	-307	-96	4	-841	-260	-96	-96
5	-540	-108	-198	-41	5	-613	-108	-38	-38
6	-376	-43	-138	-17	6	-376	-43	-14	-14
10	-137	0	-51	0	10	-137	0	0	0
11	-114	0	-41	0	11	-114	0	0	0

Table 3: Performance. The BL and KS columns give the blind rotation and key switching timings. The suffix “-C” in the parameter sets stands for using rounded continuous Gaussian sampling. In the “Total” column we give the over time to run a bootstrapping operation. The G_{simul}^{-1} column represents the proportion of the Gaussian sampling in the total computation. Remind that in the DS-WM, we must run bootstrapping several times. The Ksk, Bk, and v columns give sizes of the respective public keys.

	Total [s]		BL [s]		G_{simul}^{-1}	Ksk [MB]	Bk [MB]	v [MB]
	det	simul	det	simul				
Ours-Int		1.36		1.36	79%			
Ours-Int-C	0.14	1.01	0.14	1.01	72%	79	134	186
DS-WM-Int		2.03		0.39	–			
Ours-Double		1.33		1.33	59%			
Ours-Double-C	0.27	0.91	0.27	0.91	40%	56	168	69
DS-WM-Double		7.10		0.59	–			

As we can see from Table 3, our method is faster when compared to both DS-WM methods. In particular, Ours-Int is roughly $1.56\times$ faster than DS-WM-Int ($2.03\times$ for rounded Box-Muller). Then, Our-Double is about $5.46\times$ faster than its DS-WM-Double analog ($7.1\times$ for rounded Box-Muller). The “Double” parameters are slower than the “Int” parameters in `det` mode. The reason for this is that the “Double” parameters require computing more negacyclic convolutions than the “Int” parameters due to their different modulus and decomposition factors. It’s important to note that all parameter sets should achieve similar correctness levels. Therefore, for the “Double” parameter sets, which are intended to be implemented using IEEE-754 double precision floating point format, we are constrained by the precision of the arithmetic. When estimating the key size, we list two methods. One assumes storing integers or (discretized) floating point numbers in an array of 8-bit bytes. This could potentially give an edge to the “Double” parameter sets since these sets use a smaller modulus. However, these sets still require larger key material because we need to store more ring elements. Although the byte array representation can be used when transmitting a key, when computing the bootstrapping operation, the integers or floats are stored in random access memory in 64-bit registers, regardless of whether the modulus is a 48-bit or 36-bit number. In any case the Our-Int parameter set outperforms all other sets. The size of a ciphertext is the same for every parameter set. Note that in all sets we can send the first ciphertexts with a modulus equal to $2 \cdot N = 2^{12}$, and all sets have the same LWE dimension $n = 912$. Consequently, the ciphertexts will take approximately 0.46 [MB]. Finally, note that in the case of Our-Int, sanitization is roughly $9\times$ slower ($7\times$ for rounded Box-Muller) than deterministic computation. For Our-Double, sanitization is only $4.8\times$ slower ($3.3\times$ for rounded Box-Muller) than deterministic computation. Partially, the reason for this is that the sanitization algorithms have smaller decomposition bases, but a notable portion of the computation is spent on computing Gaussian sampling. For Gaussian sampling, in the special case where the modulus is a power of the decomposition basis, as is the case in Our-Double, approximately half of the computation is spent on sampling Gaussians. In the general case, such as Our-Int, where the modulus is an NTT-friendly prime number, Gaussian sampling constitutes approximately 78% of the entire computation. We stress that the Gaussian sampling step is a straightforward implementation of the method from [GM18, CDCG⁺18]. There is still much room for improvement in the implementation. In particular, an optimized implementation could take advantage of AVX vector processor extensions to parallelize parts or even the entire sampling algorithm. We do not see much room for improvement in the washing machine method because its only difference from a deterministic bootstrap is the choice of uniform flooding noise.

6 Conclusions and Open Problems

We showed that it is practically feasible to build an efficient FHE scheme with circuit privacy, that outperforms the Ducas-Stehlé washing machine method [DS16]. We believe that an optimized implementation and using faster discrete Gaussian sampling algorithms [MW17, DFW22] instead of Karney [Kar16] may further improve the performance. An interesting problem is to analyze whether we can use randomized gadget decomposition that has output from other distributions and discrete Gaussian. In particular, it may be worth exploring the use of more efficient subgaussian samplers [GMP19, ZY22, JLP21] in place of the Gaussian sampling algorithms [GM18, CDCG⁺18].

References

- [ABSdV19] Mark Abspoel, Niek J. Bouman, Berry Schoenmakers, and Niels de Vreede. Fast secure comparison for medium-sized integers and its application in binarized neural networks. In Mitsuru Matsui, editor, *CT-RSA 2019*, volume 11405 of *LNCS*, pages 453–472. Springer, Cham, March 2019. doi:[10.1007/978-3-030-12612-4_23](https://doi.org/10.1007/978-3-030-12612-4_23).
- [ACC⁺18] Martin Albrecht, Melissa Chase, Hao Chen, Jintai Ding, Shafi Goldwasser, Sergey Gorbunov, Shai Halevi, Jeffrey Hoffstein, Kim Laine, Kristin Lauter, Satya Lokam, Daniele Micciancio, Dustin Moody, Travis Morrison, Amit Sahai, and Vinod Vaikuntanathan. Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, Toronto, Canada, November 2018.
- [ADDG24] Martin R. Albrecht, Alex Davidson, Amit Deo, and Daniel Gardham. Crypto dark matter on the torus - oblivious PRFs from shallow PRFs and TFHE. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part VI*, volume 14656 of *LNCS*, pages 447–476. Springer, Cham, May 2024. doi:[10.1007/978-3-031-58751-1_16](https://doi.org/10.1007/978-3-031-58751-1_16).
- [AGHS13] Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian leftover hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 97–116. Springer, Berlin, Heidelberg, December 2013. doi:[10.1007/978-3-642-42033-7_6](https://doi.org/10.1007/978-3-642-42033-7_6).
- [AGHV22] Adi Akavia, Craig Gentry, Shai Halevi, and Margarita Vald. Achievable CCA2 relaxation for homomorphic encryption. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part II*, volume 13748 of *LNCS*, pages 70–99. Springer, Cham, November 2022. doi:[10.1007/978-3-031-22365-5_3](https://doi.org/10.1007/978-3-031-22365-5_3).
- [AJL⁺12] Gilad Asharov, Abhishek Jain, Adriana López-Alt, Eran Tromer, Vinod Vaikuntanathan, and Daniel Wichs. Multiparty computation with low communication, computation and interaction via threshold FHE. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 483–501. Springer, Berlin, Heidelberg, April 2012. doi:[10.1007/978-3-642-29011-4_29](https://doi.org/10.1007/978-3-642-29011-4_29).
- [AP13] Jacob Alperin-Sheriff and Chris Peikert. Practical bootstrapping in quasilinear time. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 1–20. Springer, Berlin, Heidelberg, August 2013. doi:[10.1007/978-3-642-40041-4_1](https://doi.org/10.1007/978-3-642-40041-4_1).
- [AP14] Jacob Alperin-Sheriff and Chris Peikert. Faster bootstrapping with polynomial error. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 297–314. Springer, Berlin, Heidelberg, August 2014. doi:[10.1007/978-3-662-44371-2_17](https://doi.org/10.1007/978-3-662-44371-2_17).
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. URL: <https://doi.org/10.1515/jmc-2015-0016>, doi:[10.1515/jmc-2015-0016](https://doi.org/10.1515/jmc-2015-0016).
- [AR13] Divesh Aggarwal and Oded Regev. A note on discrete gaussian combinations of lattice vectors, 2013. URL: <https://arxiv.org/abs/1308.2405>, doi:[10.48550/ARXIV.1308.2405](https://doi.org/10.48550/ARXIV.1308.2405).

- [BCGZ24] Weihao Bai, Long Chen, Qianwen Gao, and Zhenfeng Zhang. MPC-in-the-head framework without repetition and its applications to the lattice-based cryptography. In *2024 IEEE Symposium on Security and Privacy*, pages 578–596. IEEE Computer Society Press, May 2024. doi:10.1109/SP54263.2024.00157.
- [BdPMW16] Florian Bourse, Rafaël del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 62–89. Springer, Berlin, Heidelberg, August 2016. doi:10.1007/978-3-662-53008-5_3.
- [BGGJ20] Christina Boura, Nicolas Gama, Mariya Georgieva, and Dimitar Jetchev. Chimera: Combining ring-lwe-based fully homomorphic encryption schemes. *Journal of Mathematical Cryptology*, 14(1):316–338, 2020. URL: <https://doi.org/10.1515/jmc-2019-0026> [cited 2024-12-24], doi:doi:10.1515/jmc-2019-0026.
- [BGGJ20] Marcelo Blatt, Alexander Gusev, Yuriy Polyakov, and Shafi Goldwasser. Secure large-scale genome-wide association studies using homomorphic encryption. *Proceedings of the National Academy of Sciences*, 117(21):11608–11613, 2020. doi:10.1073/pnas.1918257117.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In Shafi Goldwasser, editor, *ITCS 2012*, pages 309–325. ACM, January 2012. doi:10.1145/2090236.2090262.
- [BI22] Florian Bourse and Malika Izabachène. Plug-and-play sanitization for TFHE. Cryptology ePrint Archive, Report 2022/1438, 2022. URL: <https://eprint.iacr.org/2022/1438>.
- [BIP⁺18] Dan Boneh, Yuval Ishai, Alain Passelègue, Amit Sahai, and David J. Wu. Exploring crypto dark matter: New simple PRF candidates and their applications. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 699–729. Springer, Cham, November 2018. doi:10.1007/978-3-030-03810-6_25.
- [BKS⁺21] Fabian Boemer, Sejun Kim, Gelila Seifu, Fillipe DM de Souza, Vinodh Gopal, et al. Intel HEXL (release 1.2). <https://github.com/intel/hexl>, 2021.
- [BM58] George EP Box and Mervin E Muller. A note on the generation of random normal deviates. *The annals of mathematical statistics*, 29(2):610–611, 1958.
- [BN20] Carsten Baum and Ariel Nof. Concretely-efficient zero-knowledge arguments for arithmetic circuits and their application to lattice-based cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 495–526. Springer, Cham, May 2020. doi:10.1007/978-3-030-45374-9_17.
- [Bot24] Enrico Bottazzi. Greco: Fast zero-knowledge proofs for valid FHE RLWE ciphertexts formation. Cryptology ePrint Archive, Report 2024/594, 2024. URL: <https://eprint.iacr.org/2024/594>.
- [BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011. doi:10.1109/FOCS.2011.12.

- [BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014*, pages 1–12. ACM, January 2014. doi:10.1145/2554797.2554799.
- [CDCG⁺18] David Bruce Cousins, Giovanni Di Crescenzo, Kamil Doruk Gür, Kevin King, Yuriy Polyakov, Kurt Rohloff, Gerard W. Ryan, and Erkay Savas. Implementing conjunction obfuscation under entropic ring lwe. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 354–371, 2018. doi:10.1109/SP.2018.00007.
- [CDKS19] Hao Chen, Wei Dai, Miran Kim, and Yongsoo Song. Efficient multi-key homomorphic encryption with packed ciphertexts with application to oblivious neural network inference. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 395–412. ACM Press, November 2019. doi:10.1145/3319535.3363207.
- [CdWM⁺17] Hervé Chabanne, Amaury de Wargny, Jonathan Milgram, Constance Morel, and Emmanuel Prouff. Privacy-preserving classification on deep neural network. Cryptology ePrint Archive, Report 2017/035, 2017. URL: <https://eprint.iacr.org/2017/035>.
- [CGGI16a] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part I*, volume 10031 of *LNCS*, pages 3–33. Springer, Berlin, Heidelberg, December 2016. doi:10.1007/978-3-662-53887-6_1.
- [CGGI16b] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption library, August 2016. <https://tfhe.github.io/tfhe/>.
- [CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. TFHE: Fast fully homomorphic encryption over the torus. *Journal of Cryptology*, 33(1):34–91, January 2020. doi:10.1007/s00145-019-09319-x.
- [CH18] Hao Chen and Kyoohyung Han. Homomorphic lower digits removal and improved FHE bootstrapping. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 315–337. Springer, Cham, April / May 2018. doi:10.1007/978-3-319-78381-9_12.
- [CJL⁺20] Ilaria Chillotti, Marc Joye, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Concrete: Concrete operates on ciphertexts rapidly by extending tfhe. In *WAHC 2020–8th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, volume 15, 2020.
- [CLOT21] Ilaria Chillotti, Damien Ligier, Jean-Baptiste Orfila, and Samuel Tap. Improved programmable bootstrapping with larger precision and efficient arithmetic circuits for TFHE. In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part III*, volume 13092 of *LNCS*, pages 670–699. Springer, Cham, December 2021. doi:10.1007/978-3-030-92078-4_23.
- [CLR17] Hao Chen, Kim Laine, and Peter Rindal. Fast private set intersection from homomorphic encryption. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1243–1255. ACM Press, October / November 2017. doi:10.1145/3133956.3134061.

- [CO17] Wutichai Chongchitmate and Rafail Ostrovsky. Circuit-private multi-key FHE. In Serge Fehr, editor, *PKC 2017, Part II*, volume 10175 of *LNCS*, pages 241–270. Springer, Berlin, Heidelberg, March 2017. doi:10.1007/978-3-662-54388-7_9.
- [DFW22] Yusong Du, Baoying Fan, and Baodian Wei. An improved exact sampling algorithm for the standard normal distribution. *Computational Statistics*, 37(2):721–737, 2022.
- [DGBL⁺16] Nathan Dowlin, Ran Gilad-Bachrach, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. *ICML'16*, page 201–210. JMLR.org, 2016.
- [DM15] Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 617–640. Springer, Berlin, Heidelberg, April 2015. doi:10.1007/978-3-662-46800-5_24.
- [DORS08] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM journal on computing*, 38(1):97–139, mar 2008. doi:10.1137/060651380.
- [DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 523–540. Springer, Berlin, Heidelberg, May 2004. doi:10.1007/978-3-540-24676-3_31.
- [DS16] Léo Ducas and Damien Stehlé. Sanitization of FHE ciphertexts. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part I*, volume 9665 of *LNCS*, pages 294–310. Springer, Berlin, Heidelberg, May 2016. doi:10.1007/978-3-662-49890-3_12.
- [DSGKS21] Dana Dachman-Soled, Huijing Gong, Mukul Kulkarni, and Aria Shahverdi. Towards a ring analogue of the leftover hash lemma. *Journal of Mathematical Cryptology*, 15(1):87–110, 2021. URL: <https://doi.org/10.1515/jmc-2020-0076> [cited 2022-10-06], doi:doi:10.1515/jmc-2020-0076.
- [FJ21] Matteo Frigo and Steven G. Johnson. Fftw. <https://www.fftw.org>, 2021.
- [Gen09a] Craig Gentry. *A Fully Homomorphic Encryption Scheme*. PhD thesis, Stanford University, Stanford, CA, USA, 2009.
- [Gen09b] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009. doi:10.1145/1536414.1536440.
- [GH11] Craig Gentry and Shai Halevi. Fully homomorphic encryption without squashing using depth-3 arithmetic circuits. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 107–109. IEEE Computer Society Press, October 2011. doi:10.1109/FOCS.2011.94.

- [GHS12] Craig Gentry, Shai Halevi, and Nigel P. Smart. Better bootstrapping in fully homomorphic encryption. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 1–16. Springer, Berlin, Heidelberg, May 2012. doi:10.1007/978-3-642-30057-8_1.
- [GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. i-Hop homomorphic encryption and rerandomizable Yao circuits. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 155–172. Springer, Berlin, Heidelberg, August 2010. doi:10.1007/978-3-642-14623-7_9.
- [GM18] Nicholas Genise and Daniele Micciancio. Faster Gaussian sampling for trapdoor lattices with arbitrary modulus. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 174–203. Springer, Cham, April / May 2018. doi:10.1007/978-3-319-78381-9_7.
- [GMP19] Nicholas Genise, Daniele Micciancio, and Yuriy Polyakov. Building an efficient lattice gadget toolkit: Subgaussian sampling and more. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part II*, volume 11477 of *LNCS*, pages 655–684. Springer, Cham, May 2019. doi:10.1007/978-3-030-17656-3_23.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. doi:10.1145/1374376.1374407.
- [GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 75–92. Springer, Berlin, Heidelberg, August 2013. doi:10.1007/978-3-642-40041-4_5.
- [HFH99] Bernardo A. Huberman, Matt Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In *Proceedings of the 1st ACM Conference on Electronic Commerce, EC '99*, page 78–86, New York, NY, USA, 1999. Association for Computing Machinery. doi:10.1145/336992.337012.
- [HLS18] Andreas Hülsing, Tanja Lange, and Kit Smeets. Rounded Gaussians - fast and secure constant-time sampling for lattice-based crypto. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018, Part II*, volume 10770 of *LNCS*, pages 728–757. Springer, Cham, March 2018. doi:10.1007/978-3-319-76581-5_25.
- [HS15] Shai Halevi and Victor Shoup. Bootstrapping for HELib. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 641–670. Springer, Berlin, Heidelberg, April 2015. doi:10.1007/978-3-662-46800-5_25.
- [HS21] Shai Halevi and Victor Shoup. Bootstrapping for HELib. *Journal of Cryptology*, 34(1):7, January 2021. doi:10.1007/s00145-020-09368-7.
- [IP07] Yuval Ishai and Anat Paskin. Evaluating branching programs on encrypted data. In Salil P. Vadhan, editor, *TCC 2007*, volume 4392 of *LNCS*, pages 575–594. Springer, Berlin, Heidelberg, February 2007. doi:10.1007/978-3-540-70936-7_31.

- [JKLS18] Xiaoqian Jiang, Miran Kim, Kristin E. Lauter, and Yongsoo Song. Secure outsourced matrix computation and application to neural networks. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 1209–1222. ACM Press, October 2018. doi: [10.1145/3243734.3243837](https://doi.org/10.1145/3243734.3243837).
- [JLP21] Sohyun Jeon, Hyang-Sook Lee, and Jeongeun Park. Efficient lattice gadget decomposition algorithm with bounded uniform distribution. *IEEE Access*, 9:17429–17437, 2021. doi: [10.1109/ACCESS.2021.3053288](https://doi.org/10.1109/ACCESS.2021.3053288).
- [JVC18] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. GAZELLE: A low latency framework for secure neural network inference. In William Enck and Adrienne Porter Felt, editors, *USENIX Security 2018*, pages 1651–1669. USENIX Association, August 2018.
- [Kar16] Charles F. F. Karney. Sampling exactly from the normal distribution. *ACM Trans. Math. Softw.*, 42(1), jan 2016. doi: [10.1145/2710016](https://doi.org/10.1145/2710016).
- [Klu22] Kamil Kluczniak. NTRU-v-um: Secure fully homomorphic encryption from NTRU with small modulus. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022*, pages 1783–1797. ACM Press, November 2022. doi: [10.1145/3548606.3560700](https://doi.org/10.1145/3548606.3560700).
- [KS22] Kamil Kluczniak and Leonard Schild. Fdfb: Full domain functional bootstrapping towards practical fully homomorphic encryption. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(1):501–537, Nov. 2022. URL: <https://tches.iacr.org/index.php/TCHES/article/view/9960>, doi: [10.46586/tches.v2023.i1.501-537](https://doi.org/10.46586/tches.v2023.i1.501-537).
- [KSK⁺18] Andrey Kim, Yongsoo Song, Miran Kim, Keewoo Lee, and Jung Hee Cheon. Logistic regression model training based on the approximate homomorphic encryption. *BMC medical genomics*, 11(4):23–31, 2018. doi: [10.1186/s12920-018-0401-7](https://doi.org/10.1186/s12920-018-0401-7).
- [KSK⁺20] Duhyeong Kim, Yongha Son, Dongwoo Kim, Andrey Kim, Seungwan Hong, and Jung Hee Cheon. Privacy-preserving approximate gwas computation based on homomorphic encryption. *BMC Medical Genomics*, 13(7):1–12, 2020. doi: [10.1186/s12920-020-0722-1](https://doi.org/10.1186/s12920-020-0722-1).
- [Lat22] Lattigo v4. Online: <https://github.com/tuneinsight/lattigo>, August 2022. EPFL-LDS, Tune Insight SA.
- [LJLA17] Jian Liu, Mika Juuti, Yao Lu, and N. Asokan. Oblivious neural network predictions via MiniONN transformations. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 619–631. ACM Press, October / November 2017. doi: [10.1145/3133956.3134056](https://doi.org/10.1145/3133956.3134056).
- [LMP22] Zeyu Liu, Daniele Micciancio, and Yuriy Polyakov. Large-precision homomorphic sign evaluation using FHEW/TFHE bootstrapping. In Shweta Agrawal and Dongdai Lin, editors, *ASIACRYPT 2022, Part II*, volume 13792 of *LNCS*, pages 130–160. Springer, Cham, December 2022. doi: [10.1007/978-3-031-22966-4_5](https://doi.org/10.1007/978-3-031-22966-4_5).
- [LNP22] Vadim Lyubashevsky, Ngoc Khanh Nguyen, and Maxime Plançon. Lattice-based zero-knowledge proofs and applications: Shorter, simpler, and more general. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022*,

- Part II*, volume 13508 of *LNCS*, pages 71–101. Springer, Cham, August 2022. doi:10.1007/978-3-031-15979-4_3.
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. doi:10.1007/978-3-642-38348-9_3.
- [Mea86] Catherine Meadows. A more efficient cryptographic matchmaking protocol for use in the absence of a continuously available third party. In *1986 IEEE Symposium on Security and Privacy*, pages 134–134, 1986. doi:10.1109/SP.1986.10022.
- [Mic02] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *43rd FOCS*, pages 356–365. IEEE Computer Society Press, November 2002. doi:10.1109/SFCS.2002.1181960.
- [MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Berlin, Heidelberg, April 2012. doi:10.1007/978-3-642-29011-4_41.
- [MP21] Daniele Micciancio and Yuriy Polyakov. *Bootstrapping in FHEW-like Cryptosystems*, page 17–28. Association for Computing Machinery, New York, NY, USA, 2021. URL: <https://doi.org/10.1145/3474366.3486924>.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 372–381, 2004. doi:10.1109/FOCS.2004.72.
- [MW17] Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 455–485. Springer, Cham, August 2017. doi:10.1007/978-3-319-63715-0_16.
- [OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Berlin, Heidelberg, August 2014. doi:10.1007/978-3-662-44371-2_30.
- [PAL21] PALISADE Lattice Cryptography Library (release 1.11.5). <https://palisade-crypto.org/>, 2021.
- [RAD78] RL Rivest, L Adleman, and ML Dertouzos. On data banks and privacy homomorphisms. *foundations of secure computation (1978)*, 169–180. *Search in*, 1978.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. doi:10.1145/1060590.1060603.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), September 2009. doi:10.1145/1568318.1568324.

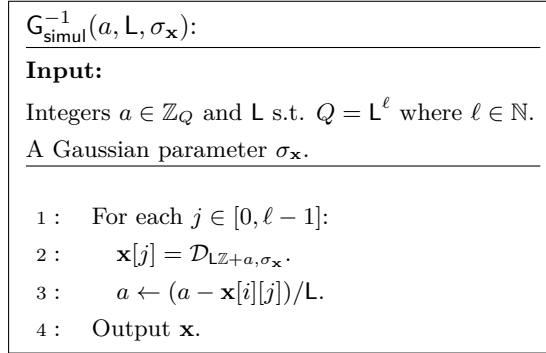


Figure 4: Gaussian Sampling Algorithm [MP12].

- [RSC⁺19] M. Sadegh Riazi, Mohammad Samragh, Hao Chen, Kim Laine, Kristin E. Lauter, and Farinaz Koushanfar. XONN: XNOR-based oblivious deep neural network inference. In Nadia Heninger and Patrick Traynor, editors, *USENIX Security 2019*, pages 1501–1518. USENIX Association, August 2019.
- [SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 27–47. Springer, Berlin, Heidelberg, May 2011. doi:10.1007/978-3-642-20465-4_4.
- [YXS⁺21] Zhaomin Yang, Xiang Xie, Huajie Shen, Shiyong Chen, and Jun Zhou. TOTA: Fully homomorphic encryption with smaller parameters and stronger security. Cryptology ePrint Archive, Report 2021/1347, 2021. URL: <https://eprint.iacr.org/2021/1347>.
- [ZY22] Shiduo Zhang and Yang Yu. Towards a simpler lattice gadget toolkit. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022, Part I*, volume 13177 of *LNCS*, pages 498–520. Springer, Cham, March 2022. doi:10.1007/978-3-030-97121-2_18.

A Additional Preliminaries

In this section we recall some useful lemmas and definitions.

A.1 Gaussian Sampling Algorithms

We recall the Gaussian sampling algorithm from [GM18] on Figure 5. For completeness, we also recall the Gaussian sampling algorithm from [MP12] on Figure 4 that is specialized for modulus in the form $Q = \mathbf{L}^\ell$. We recall only the specification for sampling given a one-dimensional integer target. Remind that we use the sampling algorithm separately on every coefficient when given as input a polynomial.

Note that for the special case modulus, the standard deviation can be bounded by $\mathbf{L} \cdot C_{\epsilon, \ell}$ which is much smaller than for the general case of $Q < \mathbf{L}^\ell$. Our correctness estimation scripts take all the decomposition algorithms into account. In practice, however, a modulus of the form $Q = \mathbf{L}^\ell$ forces us to implement negacyclic convolution of polynomials with fast Fourier transforms on floating point arithmetic. As discussed in Section 5, we found it infeasible to instantiate the scheme such that no numerical errors are induced by ring multiplications.

The algorithm on Figure 5 takes additionally the following precomputed vectors as input. The vector \mathbf{l} is such that $\mathbf{l}[1]^2 = L(1 + 1/\ell)$ and $\mathbf{l}[i]^2 = L(1 + 1/(\ell - 1))$. The vector \mathbf{h} is such that $\mathbf{h}[1] = 0$ and $\mathbf{h}[i + 1]^2 = L(1 + 1/(\ell - 1))$ for $i \in [2, \ell]$. Finally, we assume that any vector at index 0 and $\ell + 1$ is set to zero. For more details on the correctness of the Gaussian sampling algorithm for any $Q < L^\ell$ we refer to [GM18]. Furthermore, we refer to [MP12] for the analysis of the Gaussian sampling algorithm on Figure 4 for the special case $Q = L^\ell$.

Instantiation over degree $N = 2^{10}$ Rings. Numerous works like [DM15, CGGI16a, CGGI20] choose parameters for the ring $\mathbb{Z}_Q[X]/(X^N + 1)$ setting the degree to $N = 2^{10}$. The obvious benefit is that the timings are fast. What is important is that when assessing correctness, these works often report only on the correctness \mathbf{c}_{out} and ignore the correctness of \mathbf{c}_{in} . Notably, Ducas and Stehlé [DS16] propose to instantiate their washing machine method on a parameter set from the Ducas, and Micciancio’s FHEW bootstrapping [DM15], albeit they note that their instantiation is heuristic and leave a serious analysis as an open problem. We investigated the possibility of choosing a parameter set for the rings of dimension $N = 2^{10}$, and, unfortunately, we found it to be infeasible. We chose 32 bit modulus and set all decomposition bases to 2 (smallest and least efficient possible) to maximize correctness. We set $n = 810$, and the LWE standard deviations to 3.2, which gives us a low-security level of only 84 bits according to the LWE estimator [APS15]. Our estimates show that the correctness of \mathbf{c}_{in} was around the 2^{-33} level already for deterministic computation. We noticed that increasing the modulus Q , and hence dropping LWE security below 80 bits, does not change the correctness level. The reason for this is the rounding error when modulus switching from Q to $2N$. In other words, the ring degree is already so small that ciphertexts modulo $2N$ cannot accommodate the rounding error within the interval N/t . Note that we use binary keys here, which is even more beneficial for correctness than if we would use ternary or Gaussian distributed keys as in FHEW-style schemes [DM15]. To conclude, our analysis shows that the parameter choice in [DS16] for DS-WM cannot give circuit privacy is better than 30-bits, and we need to instantiate the method with a larger ring. Furthermore, we also rule out the possibility of running parameters from [CGGI16a, CGGI20] in simulation mode due to the small ring degree with statistical security larger than 30-bits.

B Error Analysis and Missing Algorithms

In this section we give the noise analysis and correctness proofs.

Proof. (External Product, Lemma 9). Denote $\mathbf{x} = \mathbf{G}_{\text{ver}}^{-1}(\mathbf{c}, \mathbf{L}_{\text{br}})$ and $\mathbf{c} = [\mathbf{a}, \mathbf{b}]$. Then we compute

$$\begin{aligned} \mathbf{C}_G \cdot \mathbf{x} &= \sum_{i=1}^{\ell_{\text{br}}} \text{RLWE}_{\sigma_G}(\mathbf{s}, \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{i-1}) \cdot \mathbf{x}[i] \\ &\quad + \sum_{i=\ell_{\text{br}}+1}^{2\ell_{\text{br}}} \text{RLWE}_{\sigma}(\mathbf{s}, -\mathbf{s} \cdot \mathbf{m}_G \cdot \mathbf{L}_{\text{br}}^{i-1}) \cdot \mathbf{x}[i] \\ &= \text{RLWE}_{\sigma}(\mathbf{s}, \mathbf{m}_G \cdot \mathbf{b}) + \text{RLWE}_{\sigma}(\mathbf{s}, -\mathbf{s} \cdot \mathbf{m}_G \cdot \mathbf{a}) \\ &= \text{RLWE}_{\sigma_1}(\mathbf{s}, \mathbf{m}_G \cdot (\mathbf{m} + \mathbf{e})) = \text{RLWE}_{\sigma_{\text{out}}}(\mathbf{s}, \mathbf{m}_G \cdot \mathbf{m}). \end{aligned}$$

The following $\sigma_1^2 \leq 2\ell_{\text{br}} \cdot N \cdot \sigma_G^2 \cdot \mathbf{B}(\mathbf{G}_{\text{ver}}^{-1}(\cdot, \mathbf{L}_{\text{br}}))$ holds because we compute the multisum $\sum_{i=1}^{2\ell_{\text{br}}} \mathbf{e}_i \cdot \mathbf{x}[i]$ where \mathbf{e}_i is the error of the i th RLWE sample in \mathbf{C}_G . Note that each coefficient of $\mathbf{e}_i \cdot \mathbf{x}[i]$ in the ring \mathcal{R}_Q is a negacyclic convolution of the coefficients in \mathbf{e}_i

$G_{\text{simul}}^{-1}(a, L, \sigma_{\mathbf{x}})$:
Input:
Integers $a \in R_Q$ and L s.t. $Q < L^\ell$, and a Gaussian parameter $\sigma_{\mathbf{x}}$.
1: Set $\mathbf{q} \leftarrow G_{\text{det}}^{-1}(Q, L)$ and $\mathbf{u} \leftarrow G_{\text{det}}^{-1}(a, L)$.
2: $\sigma = \sigma_{\mathbf{x}}/(L + 1)$.
3: $\mathbf{p} \leftarrow \text{Perturb}(\sigma, \sigma_{\mathbf{x}})$.
4: For each $i \in [\ell]$
5: $\mathbf{c}[i] \leftarrow (\mathbf{c}[i - 1] - \mathbf{u}[i] - \mathbf{p}[i])/L$.
6: $\mathbf{z} \leftarrow \text{SampleD}(\sigma, \mathbf{c}, \sigma_{\mathbf{x}})$.
7: For each $i \in [\ell - 1]$
8: $\mathbf{t}[i] \leftarrow L \cdot \mathbf{z}[i] - \mathbf{z}[i - 1] + \mathbf{q}[i] \cdot \mathbf{z}[\ell - i] + \mathbf{u}[i]$.
9: $\mathbf{t}[\ell] \leftarrow \mathbf{q}[\ell] \cdot \mathbf{z}[\ell] - \mathbf{z}[\ell - 1] + \mathbf{u}[\ell]$.
10: Return \mathbf{t} .
$\text{SampleD}(\sigma, \mathbf{c}, \sigma_{\mathbf{x}})$:
Input:
Gaussian parameters σ and $\sigma_{\mathbf{x}}$, and a vector $\mathbf{c} \in \mathbb{R}^\ell$.
1: $\mathbf{z}[\ell] \leftarrow \lfloor -\mathbf{c}[\ell]/\mathbf{d}[\ell] \rfloor$.
2: $\mathbf{z}[\ell] \leftarrow \mathbf{z}[\ell] + \text{SampleZ}_t(\sigma/\mathbf{d}[\ell], \lfloor -\mathbf{c}[\ell]/\mathbf{d}[\ell] \rfloor_{[0,1)}, \sigma_{\mathbf{x}})$.
3: $\mathbf{c} \leftarrow \mathbf{c} - \mathbf{z}[\ell] \cdot \mathbf{d}$.
4: For all $i \in [\ell - 1]$.
5: $\mathbf{z}[i] \leftarrow \lfloor -\mathbf{c}[i] \rfloor + \text{SampleZ}_t(\sigma, \lfloor -\mathbf{c}[i] \rfloor_{[0,1)}, \sigma_{\mathbf{x}})$.
6: Return \mathbf{z} .
$\text{Perturb}(\sigma, \sigma_{\mathbf{x}})$:
Input:
A Gaussian parameters σ and $\sigma_{\mathbf{x}}$.
1: $\beta \leftarrow 0$.
2: For i in $[\ell]$:
3: $\mathbf{c} \leftarrow \beta/\mathbf{1}[i]$ and $\sigma[i] \leftarrow \sigma/\mathbf{1}[i]$
4: $\mathbf{z} \leftarrow \lfloor \mathbf{c}[i] \rfloor + \text{SampleZ}_t(\sigma[i], \lfloor \mathbf{c}[i] \rfloor_{[0,1)}, \sigma_{\mathbf{x}})$.
5: $\beta \leftarrow -\mathbf{z}[i]\mathbf{h}[i]$.
6: $\mathbf{p}[1] \leftarrow (2L + 1)\mathbf{z}[1] + L\mathbf{z}[1]$.
7: For i in $[2, \ell]$:
8: $\mathbf{p} \leftarrow L(\mathbf{z}[i - 1] + 2\mathbf{z}[i] + [i + 1])$.
9: Return \mathbf{p} .

Figure 5: Gaussian Sampling Algorithm [GM18] for $Q < L^\ell$. We denote $\lfloor c \rfloor_{[0,1)} = c - \lfloor c \rfloor$. The algorithm $\text{SampleZ}_t(\sigma, c, \sigma_{\max})$ is any Gaussian sampling algorithm that samples over $\mathbb{Z} \cup [c - t \cdot \sigma_{\max}, c + t \cdot \sigma_{\max}]$ with mean c . We assume

and $\mathbf{x}[i]$. Finally, $\sigma_{\text{out}}^2 \leq \sigma^2 + \mathbf{m}_G \cdot \sigma_1^2$ holds because, $\mathbf{m}_G \cdot \boldsymbol{\epsilon}$ is 0 or $\boldsymbol{\epsilon}$ depending on the bit \mathbf{m}_G . \square

Proof. (Mux Gate, Lemma 10). Note that the RGSW sample \mathbf{C} encodes a bit $\mathbf{m}_c \in \{0, 1\}$. As in the proof of Lemma 9 we have $\mathbf{c}_{\text{out}} = \text{RLWE}_{\sigma_1}(\mathbf{s}, \mathbf{m}_c \cdot (\mathbf{m}_d - \mathbf{m}_h + \boldsymbol{\epsilon}_d - \boldsymbol{\epsilon}_h) + \mathbf{h})$. Hence, for $\mathbf{m}_C = 0$, we get

$$\mathbf{c}_{\text{out}} = \text{RLWE}_{\sigma_1}(\mathbf{s}, \mathbf{m}_h + \boldsymbol{\epsilon}_h) = \text{RLWE}_{\sigma_{\text{out}}}(\mathbf{s}, \mathbf{m}_h),$$

and for $\mathbf{m}_C = 1$, we get

$$\mathbf{c}_{\text{out}} = \text{RLWE}_{\sigma_1}(\mathbf{s}, \mathbf{m}_d + \boldsymbol{\epsilon}_g) = \text{RLWE}_{\sigma_{\text{out}}}(\mathbf{s}, \mathbf{m}_d).$$

In either case, we have that $\sigma_{\text{out}}^2 \leq 2\ell_{\text{br}} \cdot N \cdot \sigma_G^2 \cdot \text{Var}(\mathbf{G}_{\text{ver}}^{-1}(\cdot, \mathbf{L}_{\text{br}})) + \sigma^2$, because as we assumed the noise parameter for \mathbf{d} and \mathbf{h} is the same. \square

Proof. (Modulus Switching, Lemma 11). Denote $\mathbf{c} = (b, \mathbf{a})$, where $b = \mathbf{a}^\top \cdot \mathbf{s} + \mathbf{m} + e \in \mathbb{Z}_Q$ where e has variance σ^2 . Then we have the following:

$$\begin{aligned} \text{Phase}(\lfloor \frac{q}{Q} \cdot \mathbf{c} \rfloor) &= \lfloor \frac{q}{Q} \cdot b \rfloor - \lfloor \frac{q}{Q} \cdot \mathbf{a}^\top \rfloor \cdot \mathbf{s} = \frac{q}{Q} \cdot b + r - \frac{q}{Q} \cdot \mathbf{a}^\top \cdot \mathbf{s} + \mathbf{r}^\top \cdot \mathbf{s} \\ &= \frac{q}{Q} \cdot \mathbf{m} + \frac{q}{Q} \cdot e + r + \mathbf{r}^\top \cdot \mathbf{s} \end{aligned}$$

where $r \in \mathbb{R}$ and $\mathbf{r} \in \mathbb{R}^n$ are in $[-\frac{1}{2}, \frac{1}{2}]$. Then we have

$$\begin{aligned} \sigma_{\text{out}}^2 &= \text{Var}(\frac{q}{Q} \cdot e + r + \mathbf{r}^\top \cdot \mathbf{s}) = \text{Var}(\frac{q}{Q} \cdot e) + \text{Var}(\mathbf{r}^\top \cdot \mathbf{s}) \\ &= \frac{q^2}{Q^2} \cdot \sigma^2 + \sum_{i=1}^n \text{Var}(\mathbf{r}[i] \cdot \mathbf{s}[i]) \leq \frac{q^2}{Q^2} \cdot \sigma^2 + \frac{1}{4} \cdot \text{Hw}(\mathbf{s}) \cdot \text{Var}(\mathbf{s}). \end{aligned}$$

The expectation of the output noise satisfies

$$\begin{aligned} &|\frac{q}{Q} \cdot \mathbb{E}(\text{Error}(\mathbf{c})) + \mathbb{E}(\frac{q}{Q} \cdot e + r + \mathbf{r}^\top \cdot \mathbf{s})| \\ &= |\frac{q}{Q} \cdot \mathbb{E}(\text{Error}(\mathbf{c}))| + |\mathbb{E}(r + \mathbf{r}^\top \cdot \mathbf{s})| \\ &= |\frac{q}{Q} \cdot \mathbb{E}(\text{Error}(\mathbf{c}))| + |\mathbb{E}(r) + \sum_{i=1}^n \mathbb{E}(\mathbf{r} \cdot \mathbf{s})| \\ &\leq |\frac{q}{Q} \cdot \mathbb{E}(\text{Error}(\mathbf{c}))| + 1/2 + 1/2 \cdot \text{Hw}(\mathbf{s}) \cdot |\mathbb{E}(\mathbf{s})| \end{aligned}$$

given that the expectation of e is 0. \square

Proof. (Sample Extraction, Lemma 12). Denote $\mathbf{s} = \boldsymbol{\varsigma} \in \mathbb{Z}_Q^N$ and $b = \mathbf{b}[1] \in \mathbb{Z}_Q$. Denote $\mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \mathbf{m} + \boldsymbol{\epsilon} \in \mathcal{R}_Q$, $\mathbf{m} = \sum_{i=1}^N \mathbf{m}[i] \cdot X^{i-1}$ and $\boldsymbol{\epsilon} = \sum_{i=1}^N \boldsymbol{\epsilon}[i] \cdot X^{i-1}$ then it is easy to see, that $b = (\mathbf{a} \cdot \mathbf{s})[1] + \mathbf{m}[1] + \boldsymbol{\epsilon}[1]$. Furthermore, denote $\mathbf{a} = \sum_{i=1}^N \mathbf{a}[i] \cdot X^{i-1}$ and $\boldsymbol{\varsigma} = \sum_{i=1}^N \boldsymbol{\varsigma}[i] \cdot X^{i-1}$. Denote $\boldsymbol{\varsigma} \cdot \mathbf{a} = (\sum_{i=1}^N \mathbf{a}[i] \cdot X^{i-1}) \cdot (\sum_{i=1}^N \boldsymbol{\varsigma}[i] \cdot X^{i-1})$. By expanding the product we have that the constant coefficient is given by $(\boldsymbol{\varsigma} \cdot \mathbf{a})[1] = \boldsymbol{\varsigma}[1] \cdot \mathbf{a}[1] - \sum_{i=2}^N \boldsymbol{\varsigma}[i] \cdot \mathbf{a}[N-i+2]$.

If we set $\mathbf{s} = \boldsymbol{\varsigma}$ and \mathbf{a} such that $\mathbf{a}[1] = \mathbf{a}[1]$ and $\mathbf{a}[i] = -\mathbf{a}[i]$ for $i = 2 \dots N$, then (b, \mathbf{a}) is a valid LWE sample with respect to \mathbf{s} . \square

KeySwitchSetup ($\sigma_{\text{ksK}}, \mathbf{s}, \mathbf{s}'$):
Input: A bound $\sigma_{\text{ksK}} \in \mathbb{N}$. Secret keys $\mathbf{s} \in \mathbb{Z}_Q^n$, and $\mathbf{s}' \in \mathbb{Z}_Q^N$.
<ol style="list-style-type: none"> 1 : For $i \in [N], j \in [\ell_{\text{ksK}}]$ 2 : Set $\text{ksK}[\ell_{\text{ksK}}(i-1) + j] \leftarrow \text{LWE}_{\sigma_{\text{ksK}}, n, Q}(\mathbf{s}, \mathbf{s}'[i] \cdot \mathbf{L}_{\text{ksK}}^{j-1})$. 3 : Output $\text{ksK} \in \text{LWE}_{\sigma_{\text{ksK}}, n, Q}(\mathbf{s}', \cdot)^{N \ell_{\text{ksK}}}$.
KeySwitch (\mathbf{c}, ksK):
Input: A LWE ciphertext $\mathbf{c} = [b, \mathbf{a}] \in \text{LWE}_{\sigma, N, Q}(\mathbf{s}', m)$ A key switching key $\text{ksK} \in \text{LWE}_{\sigma_{\text{ksK}}, n, Q}(\mathbf{s}, \cdot)^{N \ell_{\text{ksK}}}$.
<ol style="list-style-type: none"> 1 : Compute $\mathbf{x} \leftarrow \mathbf{G}_{\det}^{-1}(\mathbf{a}, \mathbf{L}_{\text{ksK}}) \in \mathbb{Z}_{\mathbf{L}_{\text{ksK}}}^N \ell_{\text{ksK}}$. 2 : Output $c_{\text{out}} \leftarrow [b, \mathbf{0}] - \mathbf{x}^\top \cdot \text{ksK} \in \mathbb{Z}_Q^{n+1}$.

Figure 6: Key switching algorithm and its setup.

Proof. (Key Switching, Lemma 13). Let us first note that for all $i \in [n]$ we have

$$\begin{aligned} \mathbf{x}^\top \cdot \text{ksK} &= \sum_{i=1}^N \sum_{j=1}^{\ell_{\text{ksK}}} \mathbf{x}[\ell_{\text{ksK}}(i-1) + j] \cdot \text{ksK}[\ell_{\text{ksK}}(i-1) + j] \\ &= \text{LWE}_{\sigma_1, n, Q}(\mathbf{s}, \sum_{i=1}^N \mathbf{a}[i] \cdot \mathbf{s}'[i]) \end{aligned}$$

where

$$\sigma_1^2 \leq \sum_{i=1}^{\ell_{\text{ksK}}} N \cdot \mathbf{B}(\mathbf{G}_{\det}^{-1}(\cdot, \mathbf{L}_{\text{ksK}})) \cdot \sigma_{\text{ksK}}^2 \leq N \cdot \ell_{\text{ksK}} \cdot \mathbf{B}(\mathbf{G}_{\det}^{-1}(\cdot, \mathbf{L}_{\text{ksK}})) \cdot \sigma_{\text{ksK}}^2.$$

The bound follows from the fact that we have a multisum of scalars in \mathbf{x} and LWE samples from the key switching key.

Let us denote $b = \mathbf{a}^\top \cdot \mathbf{s}' + m + e$ and $\mathbf{x}^\top \cdot \text{ksK} = [\hat{b}, \hat{\mathbf{a}}]$ where $\hat{b} = \hat{\mathbf{a}}^\top \mathbf{s} + \mathbf{a}^\top \cdot \mathbf{s}' + \hat{e}$ then

$$\mathbf{c}_{\text{out}} = [b, \mathbf{0}] - \mathbf{x}^\top \cdot \text{ksK} = [b - \hat{b}, -\hat{\mathbf{a}}] = [-\hat{\mathbf{a}}^\top \mathbf{s} + m + e - \hat{e}, -\hat{\mathbf{a}}]$$

Hence, \mathbf{c}_{out} is a valid LWE sample of m with respect to key \mathbf{s} and

$$\sigma_{\text{out}}^2 \leq N \cdot \ell_{\text{ksK}} \cdot \mathbf{B}(\mathbf{G}_{\det}^{-1}(\cdot, \mathbf{L}_{\text{ksK}})) \cdot \sigma_{\text{ksK}}^2 + \sigma^2.$$

□

Proof. (Bootstrapping, Lemma 14). The correctness of blind rotation follows from two observations. First, is that multiplying a RLWE sample with X^k for some $k \in \mathbb{Z}_N$ does not change the parameter of its noise, because the error polynomial is only rotated, and we change the sign of some of the coefficients. Second, we run n times the homomorphic CMux gate, thus the variance of the output noise follows from Lemma 10. Finally, note that each iteration rotates the message by $X^{-\mathbf{a}[i] \cdot \mathbf{s}[i]}$. Denote $\mathbf{c} = [\mathbf{a}, b]$, where

Table 4: The standard deviation of the numerical errors when multiplying degree $N = 2^{11}$ polynomials. One polynomial has coefficient modulus Q , and the other one has coefficient modulus B . We choose both polynomials uniformly at random. We denote by “-” when the $\log_2(Q) + \log_2(B) + \log_2(N) \geq 64$, in which case we exceed 64 bits and end up with random-looking polynomials.

$\log_2(B) \backslash \log_2(Q)$	32	34	36	38	40	42	44	46
4	0	0	0	0.5	1.14	5.5	17.3	77.7
6	0	0	0.23	1.25	4.1	24.1	72.4	325.9
8	0	0.5	1.9	4.2	17.5	108.1	263.3	-
10	0.6	1.4	4.1	19.1	74.7	290.2	-	-
12	1.3	3.9	19.0	78.1	326	-	-	-
14	4.6	17.8	68.2	505.7	-	-	-	-
16	15.8	87.8	282.2	-	-	-	-	-
18	78.2	490.9	-	-	-	-	-	-

$b = \mathbf{a}^\top \mathbf{s} + m + e \in \mathbb{Z}_{2N}$. After n iterations we obtain $\mathbf{a}_{\text{rot}} \cdot X^{b-\mathbf{a}^\top \mathbf{s}} = \mathbf{a}_{\text{rot}} \cdot X^{m+e}$. What follows is $(\mathbf{a}_{\text{rot}} \cdot X^{m+e})[1] = f(m+e) \in \mathbb{Z}_Q$ from the assumption on \mathbf{a}_{rot} .

Correctness of bootstrapping trivially follows from the correctness of the underlying algorithms. In particular, the noise parameter of the \mathbf{c}_{in} ciphertext follows from the fact that we run the key switching procedure on \mathbf{c} and then switch the modulus to $2N$. Finally, the noise parameter of \mathbf{c}_{out} follows from the correctness of blind rotation, Lemma 13, Lemma 12 and Lemma 11. Finally, if $\text{ver} = \text{simul}$, then we additionally compute a linear combination of of LWE samples of zero from the vector \mathbf{v} . Hence the additional part $h \cdot \sigma_R^2 \cdot \sigma_{\text{rand}}^2$ of the noise follows from linear homomorphism of LWE samples and the fact that all error terms are uncorrelated. \square

Numerical Error. Finally, let us address the issue of numerical errors when performing ring operations. In particular, we focus on computing products of ring elements (or negacyclic convolutions of polynomials for our ring choice). We measured the numerical error when computing products of ring elements with the `fftw` library [FJ21]. The result is depicted in Table 4. Based on this table, we ruled out choosing certain moduli and decomposition bases while preserving correctness.

Note that Table 4 gives only the error of polynomial multiplication. Note that the impact on the ciphertext error of an external product is much higher and dependent on the secret norm of the secret key. Let (\mathbf{a}, \mathbf{b}) by a RLWE ciphertext such that $\mathbf{b} - \mathbf{a} \cdot \mathbf{s} = \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is small. Denote

$$(\mathbf{a}', \mathbf{b}') = (\mathbf{Mul}(\mathbf{a}, \mathbf{c}), \mathbf{Mul}(\mathbf{b}, \mathbf{c})) = (\mathbf{a} \cdot \mathbf{c} + \boldsymbol{\tau}_2, \mathbf{b} \cdot \mathbf{c} + \boldsymbol{\tau}_1),$$

where $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ is the numerical error introduced by the multiplication algorithm \mathbf{Mul} . Then we can see that the phase $\mathbf{b}' - \mathbf{a}' \cdot \mathbf{s} = \boldsymbol{\epsilon} - \boldsymbol{\tau}_2 \cdot \mathbf{s} + \boldsymbol{\tau}_1$. We obtain an additional error which infinity norm is

$$\|\boldsymbol{\tau}_2 \cdot \mathbf{s} + \boldsymbol{\tau}_1\| \leq NB(Q, \mathbf{c}) \cdot (\|\mathbf{s}\|_\infty + 1),$$

where $\|\boldsymbol{\tau}_1\|_\infty, \|\boldsymbol{\tau}_2\|_\infty < B(Q, \mathbf{c})$ and $B(Q, \mathbf{c})$ being an error function determined by the modulus Q and the polynomial \mathbf{c} . If the external product is implemented using such erroneous multiplication algorithm then we need to add $2\ell_{\text{br}} \cdot N \cdot B(Q, \mathbf{c}) \cdot (\|\mathbf{s}\|_\infty + 1)$ to the variance σ_1 assuming that the error function $B(Q, \mathbf{c})$ is modeled by a discrete Gaussian.

Consequently we need to update the bound on σ_{out} on the \mathbf{c}_{out} error of the bootstrapping algorithm as follows. For $\text{ver} = \text{det}$ we have

$$\sqrt{2n \cdot \ell_{\text{br}} \cdot N \cdot (\sigma_{\text{br}}^2 \cdot \mathbf{B}(\mathbf{G}_{\text{det}}^{-1}(\cdot, \mathbf{L}_{\text{br}})) + B(Q, \mathbf{c}) \cdot (\|s\|_{\infty} + 1))},$$

and for $\text{ver} = \text{simul}$ we have

$$\sqrt{2n \cdot \ell_{\text{br}} \cdot N \cdot (\sigma_{\text{br}}^2 \cdot \mathbf{B}(\mathbf{G}_{\text{simul}}^{-1}(\cdot, \mathbf{L}_{\text{br}}; \sigma_{\mathbf{x}})) + \cdot B(Q, \mathbf{c}) \cdot (\|s\|_{\infty} + 1)) + h \cdot \sigma_R^2 \cdot \sigma_{\text{rand}}^2}.$$

C Discussion on the Noise Flooding Technique

Here we give a rough parameter estimate for TFHE using noise flooding. We use our estimator for the DS-WM method setting the number of washing cycles to 1. In fact the noise flooding method is a special case of the DS-WM method. In fact we will start by modifying the DS-WM-Int parameter set. Let's set the modulus to 2^{110} and ring to 2^{12} . Note that the ring is bigger than the rings that we used, but security is below 128-bit, according to the FHE-Standard [ACC⁺18]! So the parameter set doesn't satisfy out conditions, but let's make it easier. The blind rotation error is already above $B = 26$ bits, so we need at least $B * 2^{80} = 2^{106}$ bits of flooding noise according to the smudging lemma (Lemma 1). A 110-bit modulus should be large enough to accommodate the message. All other parameters (the decomposition base etc.) stay the same so that we don't need to increase the modulus anymore. Then we need $(14 + 1) \cdot 2$ convolutions per external product ($912 \cdot 30 = 27360$ total in a larger ring), a masking key of size 3977 MB (in comparison to 186 MB), a 1464 MB bootstrapping key (we have 134 MB), and 837 MB Key switching key (we have 79 MB). Quadruple precision FFT or Intel Hexl don't handle such big numbers. But we can use the RNS representation at the cost of roughly 2times more convolutions. In total, we get $27360 \cdot 2 = 54720$ convolutions, in comparison to 12768 convolutions in Our-Int parameter set. To summarize, naive flooding has 4.28 times more convolutions with a ring that is twice as large and requires 6278 MB of key material in comparison to our 399 MB. Remind that this parameter set doesn't give 128-bit security, but is close to. To satisfy our security constraint we would actually need to take a 2^{13} dimension ring! Note that the number of convolutions is just an indication of how much slower an implementation can be in best case. To implement RNS there is much more effort necessary, that is going to slow down computation due to composition, memory access, cache related issue etc.