Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme*

Joppe W. Bos[1], Kristin Lauter[1], Jake Loftus[2], and Michael Naehrig[1]

[1] Microsoft Research
{jbos,klauter,mnaehrig}@microsoft.com
[2] University of Bristol
loftus@cs.bris.ac.uk

Abstract. In 1996, Hoffstein, Pipher and Silverman introduced an efficient lattice based encryption scheme dubbed `NTRUEncrypt`. Unfortunately, this scheme lacks a proof of security. However, in 2011, Stehlé and Steinfeld showed how to modify `NTRUEncrypt` to reduce security to standard problems in ideal lattices. In 2012, López-Alt, Tromer and Vaikuntanathan proposed a fully homomorphic scheme based on this modified system. However, to allow homomorphic operations and prove security, a non-standard assumption is required. In this paper, we show how to remove this non-standard assumption via techniques introduced by Brakerski and construct a new fully homomorphic encryption scheme from the Stehlé and Steinfeld version based on standard lattice assumptions and a circular security assumption. The scheme is scale-invariant and therefore avoids modulus switching and the size of ciphertexts is one ring element. Moreover, we present a practical variant of our scheme, which is secure under stronger assumptions, along with parameter recommendations and promising implementation results. Finally, we present an approach for encrypting larger input sizes by extending ciphertexts to several ring elements via the CRT on the message space.

1 Introduction

Fully homomorphic encryption (FHE) is a powerful form of encryption which allows an untrusted server to carry out arbitrary computation on encrypted data on behalf of a client. Introduced in [21] by Adleman, Dertouzos and Rivest, the problem of constructing a scheme which can evaluate any function on encrypted data remained open until 2009, when Gentry constructed an FHE scheme based on ideal lattices [10]. Gentry's scheme effectively laid down a blueprint for constructing FHE schemes and paved the way for many further constructions [26,3,4,6,5,24,20,11,9]. The main focus of the cryptologic research community has been on improving the efficiency of FHE and basing its security on standard assumptions.

* Most of this work was done while the third author was an intern in the Cryptography Research group at Microsoft Research.

M. Stam (Ed.): IMACC 2013, LNCS 8308, pp. 45–64, 2013.

Recently, López-Alt et al. [16] proposed a (multi-key) FHE scheme based on the work by Stehlé and Steinfeld [25] in which a provably secure version of `NTRUEncrypt` [13] is presented with security based on standard problems in ideal lattices. Unfortunately, the FHE scheme from [16] needs to make an additional assumption relating to the uniformity of the public key, the so-called decisional small polynomial ratio (DSPR) assumption, to allow homomorphic operations and remain semantically secure. We show how to avoid this additional assumption and transform the results from [25] into a fully homomorphic encryption scheme based on standard lattice assumptions only. This is achieved by limiting noise growth during homomorphic operations via a tensoring technique recently introduced by Brakerski [3]. Besides this theoretical advantage, our scheme has other attractive properties. Firstly, this new scheme is *scale-invariant* in the sense of [3], i.e. it avoids the modulus-switching technique of Brakerski, Gentry and Vaikuntanathan [4]. Secondly, we keep the property of the scheme in [16] that a ciphertext consists of only a single ring element as opposed to the two or more ring elements for schemes based purely on the (ring) learning with errors (RLWE) assumption [17]. This decreases the ciphertext size since parameters are comparable in both settings. Finally, we present a technique to increase the size of the input space by working with separate, small plaintext moduli in ciphertexts of multiple ring elements, which are later combined via the Chinese remainder theorem into a larger plaintext modulus. For some applications, this additional flexibility to increase the message space without changing parameters at the cost of increasing ciphertext size can prove especially useful.

Our main contribution is an FHE scheme based on the schemes by Stehlé and Steinfeld [25] and López-Alt et al. [16] that does not need the DSPR assumption and thus is secure under the RLWE and circular security assumptions only. The public key in both schemes is the fraction $h = gf^{-1} \mod q$ of two polynomials f and g in a cyclotomic polynomial ring modulo an integer modulus q that are sampled from a discrete Gaussian distribution. The DSPR assumption is the assumption that such a fraction is indistinguishable from uniform random in the ring modulo q. Stehlé and Steinfeld show that this assumption holds if the Gaussian is wide enough. Unfortunately, the scheme by López-Alt et al. cannot use such a wide Gaussian for key generation. Since the norms of f and g contribute to the noise growth during homomorphic multiplication, using a wide enough Gaussian means that the scheme is not guaranteed to be capable of doing even a single multiplication. We solve this problem by using decompositions and Brakerski's [3] tensoring technique. During the homomorphic multiplication procedure which includes a key switching step, we decompose the polynomial f into its bit decomposition, i.e. into a vector of polynomials with binary coefficients. This technique replaces the ring product of polynomials by a scalar product of binary decomposition vectors with vectors of polynomials multiplied by powers of 2 modulo q. The noise growth introduced in such a scalar product is bounded by a polynomial in $\log(q)$ and the degree of f, replacing the square of the norm of f in the bounds of the original scheme. Noise growth is much smaller now and it is possible to sample from a wide Gaussian to ensure the Stehlé-Steinfeld

conditions. As noted in Appendix A.1 of [16], any FHE scheme is inherently a multi-key scheme for a constant number of parties, but this construction is rather inefficient. The original scheme in [16], however, directly yields the multi-key property for a non-constant number of parties, which is much more efficient. Our scheme is not a multi-key scheme in that sense because decryption of a multi-key ciphertext would require a multiplication by the product of all keys that were involved in the generation of the ciphertext. With keys generated in the setting of Stehlé and Steinfeld, multiplying by a product of only two keys would already lead to a noise overflow, making it impossible to decrypt correctly.

The second part of the paper describes a more practical variant of the above scheme, along with details on parameter selection and implementation results. The price for obtaining security without the DSPR assumption in the above scheme lies in a large evaluation key and a complicated key switching procedure, both of which are a consequence of using the tensoring approach. Any possibility, which we are aware of, to avoid the tensor products, leads to an increase in the noise bounds that makes it necessary to reintroduce the DSPR assumption. However, if one is willing to make this assumption, there are several efficiency advantages and possible trade-offs as shown in our more practical variant. This variant keeps the general characteristics of the scheme, but simplifies key switching and avoids tensor products. A much shorter evaluation key can be achieved by using base-w instead of base-2 decompositions for a $w > 2$, e.g. $w = 2^{32}$. This increases noise growth, but ensures that the evaluation key contains only a few ring elements. Since the key switching is the main cost in homomorphic multiplication, the choice of w provides an important trade-off between homomorphic capability and multiplication efficiency. We also point out that it is possible to weaken the DSPR assumption by allowing the polynomial g to be sampled from a wider Gaussian than f. The proofs of most lemmas and theorems are given in the full version of this paper [2].

2 Preliminaries

In this section, we define all basic notation that is needed in the paper. The most important structure is the ring R. Let d be a positive integer and define $R = \mathbb{Z}[X]/(\Phi_d(X))$ as the ring of polynomials with integer coefficients modulo the d-th cyclotomic polynomial $\Phi_d(X) \in \mathbb{Z}[X]$. The degree of Φ_d is $n = \varphi(d)$, where φ is Euler's totient function. The elements of R can be uniquely represented by all polynomials in $\mathbb{Z}[X]$ of degree less than n. Arithmetic in R is arithmetic modulo $\Phi_d(X)$, which is implicit whenever we write down terms or equalities involving elements in R. An arbitrary element $a \in R$ can be written as $a = \sum_{i=0}^{n-1} a_i X^i$ with $a_i \in \mathbb{Z}$ and we identify a with its vector of coefficients $(a_0, a_1, \ldots, a_{n-1})$. In particular, a can be viewed as an element of the $\mathbb{R}$-vector space $\mathbb{R}^n$. We choose the maximum norm on $\mathbb{R}^n$ to measure the size of elements in R. The maximum norm of a is defined as $\|a\|_\infty = \max_i\{|a_i|\}$.

When multiplying two elements $g, h \in R$, the norm of their product gh expands with respect to the individual norms of g and h. The maximal norm

expansion that can occur is $\delta = \sup\{\|g \cdot h\|_\infty / (\|g\|_\infty \|h\|_\infty) : g, h \in R\}$, which is a ring constant. When d is a power of 2 and thus $\Phi_d(X) = X^n + 1$, we have $\delta = n$ [10, Section 3.4]. To keep the exposition more general, we do not restrict to this special case and work with general δ in most of what follows.

Let χ be a probability distribution on R. We assume that we can efficiently sample elements from R according to χ, and we use the standard notation $a \leftarrow \chi$ to denote that $a \in R$ is sampled from χ. The distribution χ on R is called B-bounded for some $B > 0$ if for all $a \leftarrow \chi$ we have $\|a\|_\infty < B$, i.e. a is B-bounded (see [4, Def. 3] and [16, Def. 3.1 and 3.2]). Let us introduce a specific example of a distribution on R. First, define the discrete Gaussian distribution $D_{\mathbb{Z},\sigma}$ with mean 0 and standard deviation σ over the integers, which assigns a probability proportional to $\exp(-\pi|x|^2/\sigma^2)$ to each $x \in \mathbb{Z}$. When d is a power of 2 and $\Phi_d(X) = X^n + 1$, we can take χ to be the spherical discrete Gaussian $\chi = \mathcal{D}_{\mathbb{Z}^n,\sigma}$, where each coefficient of the polynomial is sampled according to the one-dimensional distribution $D_{\mathbb{Z},\sigma}$ (see [17] for more details and why $\chi = \mathcal{D}_{\mathbb{Z}^n,\sigma}$ is the right choice in that case). The distribution χ is used in many fully homomorphic encryption schemes based on RLWE to sample random error polynomials that have small coefficients with high probability. Such polynomials are a significant part of the noise terms used in the encryption process. To deduce meaningful bounds on noise size and noise growth during homomorphic operations, we assume that the distribution we are working with is B-bounded for some B. For the discrete Gaussian, this is a reasonable assumption since sampled elements tend to be small with high probability. By rejecting samples with norm larger than B, we can sample from a truncated Gaussian distribution that is statistically close to the true discrete Gaussian if B is chosen large enough. For example, if we take $B = 6\sigma$, all samples are B-bounded with very high probability [18, Lemma 4.4].

Although the principal object of interest for our scheme is the ring R, and all polynomials that we deal with are considered to be elements of R, we often reduce polynomial coefficients modulo an integer modulus q. We denote the map that reduces an integer x modulo q and uniquely represents the result by an element in the interval $(-q/2, q/2]$ by $[\cdot]_q$. We extend this map to polynomials in $\mathbb{Z}[X]$ and thus also to elements of R by applying it to their coefficients separately, i.e. $[\cdot]_q : R \to R,\ a = \sum_{i=0}^{n-1} a_i X^i \mapsto \sum_{i=0}^{n-1} [a_i]_q X^i$. Furthermore, we extend this notation to vectors of polynomials by applying it to the entries of the vectors separately. Sometimes we reduce an integer modulo q and uniquely represent the result by an element in $[0, q)$. In this case, we write $r_q(x)$ to mean the reduction of x into $[0, q)$. A polynomial $f \in R$ is invertible modulo q if there exists a polynomial $f^{-1} \in R$ such that $ff^{-1} = \tilde{f}$, where $\tilde{f}(X) = \sum_i a_i X^i$ with $a_0 = 1 \mod q$ and $a_j = 0 \mod q$ for all $j \neq 0$. Our homomorphic encryption scheme uses two different moduli. In addition to a modulus q that is used to reduce the coefficients of the elements that represent ciphertexts, there is a second modulus $t < q$ that determines the message space R/tR, i.e. messages are polynomials in R modulo t. We make frequent use of the quantity $\Delta = \lfloor q/t \rfloor$ and it is readily verified that $q - r_t(q) = \Delta \cdot t$.

In [3], functions called BitDecomp and PowersOfTwo are used. We slightly generalize these to an arbitrary base and describe our notation next. Fix a positive integer $w > 1$ that is used to represent integers in a radix-w system. Let $\ell_{w,q} = \lfloor \log_w(q) \rfloor + 2$, then a non-negative integer $z < q$ can be written as $\sum_{i=0}^{\ell_{w,q}-2} z_i w^i$ where the z_i are integers such that $0 \leq z_i < w$. If z is an integer in the interval $(-q/2, q/2]$, it can be written uniquely as $\sum_{i=0}^{\ell_{w,q}-1} z_i w^i$ with $z_i \in (-w/2, w/2]$. With this, an element $x \in R$ with coefficients in $(-q/2, q/2]$ can be written as $\sum_{i=0}^{\ell_{w,q}-1} x_i w^i$, where $x_i \in R$ with coefficients in $(-w/2, w/2]$. Since then $x_i = [x_i]_w$, we write $x = \sum_{i=0}^{\ell_{w,q}-1} [x_i]_w w^i$ to make clear that the norm of the coefficient polynomials x_i is at most $w/2$. With this notation, define

$$D_{w,q} : R \to R^{\ell_{w,q}}, \quad x \mapsto ([x_0]_w, [x_1]_w, \ldots, [x_{\ell_{w,q}-1}]_w) = ([x_i]_w)_{i=0}^{\ell_{w,q}-1},$$

this function for $w = 2$ is called BitDecomp in [3]. We define a second function

$$P_{w,q} : R \to R^{\ell_{w,q}}, \quad x \mapsto ([x]_q, [xw]_q, \ldots, [xw^{\ell_{w,q}-1}]_q) = ([xw^i]_q)_{i=0}^{\ell_{w,q}-1},$$

which is called PowersOfTwo in [3] for $w = 2$. For any two $x, y \in R$, we see that the scalar product of the vectors $D_{w,q}(x)$ and $P_{w,q}(y)$ is the same as the product xy modulo q, because

$$\langle D_{w,q}(x), P_{w,q}(y) \rangle = \sum_{i=0}^{\ell_{w,q}-1} [x_i]_w [yw^i]_q \equiv y \sum_{i=0}^{\ell_{w,q}-1} [x_i]_w w^i \equiv xy \pmod{q}.$$

Note that when $\|f\|_\infty < B$ for some $B < q$, then only the $\ell_{w,B} := \lfloor \log_w(B) \rfloor + 2$ least significant polynomials in $D_{w,q}(f)$ can be non-zero. We use the tensor product of two vectors in the usual way, i.e. for a positive integer ℓ and two vectors $a, b \in R^\ell$, the tensor $a \otimes b \in R^{\ell^2}$ is the concatenation of the $a_i b$ for $i \in \{1, 2, \ldots, \ell\}$. We extend the functions $D_{w,q}$ and $P_{w,q}$ to vectors. For $v = (v_1, v_2, \ldots, v_\ell) \in R^\ell$ denote the vector $(D_{w,q}(v_1), \ldots, D_{w,q}(v_\ell)) \in R^{\ell \cdot \ell_{w,q}}$ by $D_{w,q}(v)$, likewise we extend $P_{w,q}$.

Several operations in the scheme require scaling by rational numbers such that the resulting polynomials do not necessarily belong to R but instead have rational coefficients. In that case, a rounding procedure is applied to get back to integer coefficients. The usual rounding of a rational number a to the nearest integer is denoted by $\lfloor a \rceil$.

The Ring Learning With Errors (RLWE) Problem. Our scheme relies on the hardness of the (decisional) ring learning with errors problem, which was first introduced by Lyubashevsky, Peikert and Regev [17].

Definition 1 (Decision-RLWE). *Given a security parameter λ, let d and q be integers depending on λ, let $R = \mathbb{Z}[X]/(\Phi_d(X))$ and let $R_q = R/qR$. Given a distribution χ over R_q that depends on λ, the Decision-$\mathsf{RLWE}_{d,q,\chi}$ problem is to distinguish the following two distributions. The first distribution consists of pairs (a, u), where $a, u \leftarrow R_q$ are drawn uniformly at random from R_q. The second*

distribution consists of pairs of the form $(a, a \cdot s + e)$. *The element* $s \leftarrow R_q$ *is drawn uniformly at random and is fixed for all samples. For each sample,* $a \leftarrow R_q$ *is drawn uniformly at random, and* $e \leftarrow \chi$. *The Decision-*$\mathsf{RLWE}_{d,q,\chi}$ *assumption is that the Decision-*$\mathsf{RLWE}_{d,q,\chi}$ *problem is hard.*

In [17], it was shown that the hardness of RLWE can be established by a quantum reduction to worst-case shortest vector problems in ideal lattices over the ring R, see also [4, Thm. 2]. It is known that the *search* variant of $\mathsf{RLWE}_{d,q,\chi}$, in which we are required to explicitly find the secret s given an $\mathsf{RLWE}_{d,q,\chi}$ instance, is equivalent to the decision problem [17]. There are a number of variants of RLWE which are as hard as RLWE, for example we can restrict the sampling of a and e to invertible elements only [25]. And we can also choose s from χ without incurring any loss of security [1].

The Decisional Small Polynomial Ratio (DSPR) Problem. In [16], López-Alt, Tromer and Vaikuntanathan introduced the decisional small polynomial ratio problem. They describe a multi-key fully homomorphic encryption scheme with security based on the assumption that the DSPR problem is hard in the ring R_q where $R = \mathbb{Z}[x]/(x^n + 1)$ for n a power of 2 and $t = 2$. We state a more general form of the problem for any cyclotomic ring $R = \mathbb{Z}[x]/(\Phi_d(x))$ and general $1 < t < q$. Let $h = tg/f \pmod q$ where $f = 1 + tf'$ and $f', g \leftarrow \chi$ where χ is a truncated Gaussian distribution. In [16], the problem of distinguishing such an element h from a uniformly random element of $R_q = R/qR$ was formalized as the DSPR problem. Assuming the hardness of DSPR and RLWE, the scheme in [16] is secure. To state the problem, define the following: for a distribution χ on R_q and $z \in R_q$ we define $\chi_z = \chi + z$ to be the distribution shifted by z. Also, let $R_q^\times$ be the set of all invertible elements in R_q.

Definition 2 (DSPR). *For security parameter* λ, *let* d *and* q *be integers, let* $R = \mathbb{Z}[X]/(\Phi_d(X))$ *and* $R_q = R/qR$ *and let* χ *be a distribution over* R_q, *all depending on* λ. *Let* $t \in R_q^\times$ *be invertible in* R_q, $y_i \in R_q$ *and* $z_i = -y_i t^{-1} \pmod q$ *for* $i \in \{1, 2\}$. *The* $\mathsf{DSPR}_{d,q,\chi}$ *problem is to distinguish elements of the form* $h = a/b$ *where* $a \leftarrow y_1 + t \cdot \chi_{z_1}, b \leftarrow y_2 + t \cdot \chi_{z_2}$ *from uniformly random elements of* R_q. *The* $\mathsf{DSPR}_{d,q,\chi}$ *assumption is that the* $\mathsf{DSPR}_{d,q,\chi}$ *problem is hard.*

Theorem 4.1 in the full version of [25] shows that $\mathsf{DSPR}_{d,q,\chi}$ is hard when the χ_{z_i} are shifted versions of a discrete Gaussian distributions χ which is $\mathcal{D}_{\mathbb{Z}^n,\sigma}$ restricted to $R_q^\times$ for a large enough deviation σ. For convenience, we state the theorem in the full version of this paper [2, Appendix A]. A discrete Gaussian on $R_q^\times$ can be obtained from a discrete Gaussian on R_q by rejecting non-invertible elements.

3 Basic Scheme

In this section, we describe the basic public key encryption scheme that is the foundation for the leveled schemes of the next sections. The scheme is parameterized by a modulus q and a plaintext modulus $1 < t < q$. Ciphertexts are elements

of $R = \mathbb{Z}[X]/(\Phi_d(X))$ and plaintexts are elements of R/tR (see Section 2). Secret keys and errors are generated from different distributions, for example Gaussian distributions of different width. The secret key is derived from the distribution χ_{key}, and errors are sampled from the distribution χ_{err}. We use "Regev-style" encryption as in [3] and [9]. The scheme consists of the following algorithms.

- $\mathsf{Basic.ParamsGen}(\lambda)$: Given the security parameter λ, fix a positive integer d that determines R, moduli q and t with $1 < t < q$, and distributions $\chi_{\text{key}}, \chi_{\text{err}}$ on R. Output $(d, q, t, \chi_{key}, \chi_{err})$.
- $\mathsf{Basic.KeyGen}(d, q, t, \chi_{\text{key}}, \chi_{\text{err}})$: Sample $f', g \leftarrow \chi_{\text{key}}$ and let $f = [tf' + 1]_q$. If f is not invertible modulo q, choose a new f'. Compute the inverse $f^{-1} \in R$ of f modulo q and set $h = [tgf^{-1}]_q$. Output the public and private key pair $(\mathsf{pk}, \mathsf{sk}) = (h, f) \in R^2$.
- $\mathsf{Basic.Encrypt}(h, m)$: The message space is R/tR. For a message $m + tR$, choose $[m]_t$ as its representative. Sample $s, e \leftarrow \chi_{\text{err}}$, and output the ciphertext $c = [\lfloor q/t \rfloor [m]_t + e + hs]_q \in R$.
- $\mathsf{Basic.Decrypt}(f, c)$: To decrypt a ciphertext c, compute

$$m = \left[\left\lfloor \frac{t}{q} \cdot [fc]_q \right\rceil \right]_t \in R.$$

In the following, we often refer to a message as an element m in the ring R although the message space is R/tR, keeping in mind that encryption always takes place on the representative $[m]_t$ and that by decrypting, all that can be recovered is m modulo t.

Correctness. The following lemma states conditions for a ciphertext c such that the decryption algorithm outputs the message m that was originally encrypted.

Lemma 1. *Let q, t, and $\Delta = \lfloor q/t \rfloor$ be as above and let $c, f, m \in R$. If there exists $v \in R$ such that*

$$fc = \Delta[m]_t + v \pmod{q} \text{ and } \|v\|_\infty < (\Delta - r_t(q))/2,$$

then $\mathsf{Basic.Decrypt}(f, c) = [m]_t$*, i.e. c decrypts correctly under the secret key f.*

Of course, for any given c, f and m, there always exists a $v \in R$ such that $fc = \Delta[m]_t + v \pmod{q}$. But only a v of small norm allows one to recover $[m]_t$ from c. Since we are always free to vary v modulo q, i.e. to add any multiple of q to it, we choose v to be the canonical element $[v]_q$. This means that we choose v with the smallest possible norm among all polynomials that satisfy the equation. We call this specific v the *inherent noise in c with respect to m and f*. The previous lemma says that if the inherent noise in a ciphertext is small enough, then decryption works correctly.

Inherent Noise in Initial Ciphertexts. The following lemma derives a bound on the inherent noise in a freshly encrypted ciphertext output by $\mathsf{Basic.Encrypt}$, assuming bounds B_{key} on the key and B_{err} on the error distributions. Note that since $f', g \leftarrow \chi_{\text{key}}$we have $\|f'\|_\infty, \|g\|_\infty < B_{\text{key}}$ and it follows that $\|tg\|_\infty < tB_{\text{key}}$ and $\|f\|_\infty = \|1 + tf'\|_\infty < tB_{\text{key}}$ since $t \geq 2$.

Lemma 2. *Let the key and error distributions be B_{key}-bounded and B_{err}-bounded, respectively. Given $m \in R$, a public key $h = [tgf^{-1}]_q \in R$ with secret key $f = [1 + tf']_q$, $f', g \leftarrow \chi_{\mathrm{key}}$, and let $c = \mathsf{Basic.Encrypt}(h, m)$. There exists $v \in R$ such that $fc = \Delta[m]_t + v \pmod q$ and*

$$\|v\|_\infty < \delta t B_{\mathrm{key}} \left(2B_{\mathrm{err}} + \frac{1}{2} r_t(q)\right).$$

In particular, by Lemma 1, decryption works correctly if $2\delta t B_{\mathrm{key}}(2B_{\mathrm{err}} + \frac{1}{2} r_t(q)) + r_t(q) < \Delta$.

4 Leveled Homomorphic Scheme

In this section, we state our leveled homomorphic encryption scheme YASHE[1] based on the Basic scheme from the previous section. We then analyze the homomorphic operations and deduce bounds on the noise growth that occurs during these operations.

- $\mathsf{YASHE.ParamsGen}(\lambda)$: Given the security parameter λ, output the parameters $(d, q, t, \chi_{\mathrm{key}}, \chi_{\mathrm{err}}, w)$, where $(d, q, t, \chi_{\mathrm{key}}, \chi_{\mathrm{err}}) \leftarrow \mathsf{BasicParamsGen}(\lambda)$ and $w > 1$ is an integer.
- $\mathsf{YASHE.KeyGen}(d, q, t, \chi_{\mathrm{key}}, \chi_{\mathrm{err}}, w)$: Compute

  $$h, f \leftarrow \mathsf{Basic.KeyGen}(d, q, t, \chi_{\mathrm{key}}, \chi_{\mathrm{err}}).$$

  Sample $\boldsymbol{e}, \boldsymbol{s} \leftarrow \chi_{\mathrm{err}}^{\ell^3_{w,q}}$, compute

  $$\boldsymbol{\gamma} = [f^{-1} P_{w,q}(D_{w,q}(f) \otimes D_{w,q}(f)) + \boldsymbol{e} + h \cdot \boldsymbol{s}]_q \in R^{\ell^3_{w,q}},$$

  and output $(\mathsf{pk}, \mathsf{sk}, \mathsf{evk}) = (h, f, \boldsymbol{\gamma})$.
- $\mathsf{YASHE.Encrypt}(\mathsf{pk}, m)$: Encrypt $m \in R$ by $c \leftarrow \mathsf{Basic.Encrypt}(\mathsf{pk}, m) \in R$.
- $\mathsf{YASHE.Decrypt}(\mathsf{sk}, c)$: Output the message $m \leftarrow \mathsf{Basic.Decrypt}(\mathsf{sk}, c) \in R$.
- $\mathsf{YASHE.KeySwitch}(\tilde{c}_{\mathrm{mult}}, \mathsf{evk})$: Output $\left[\langle D_{w,q}(\tilde{c}_{\mathrm{mult}}), \mathsf{evk}\rangle\right]_q \in R$.
- $\mathsf{YASHE.Add}(c_1, c_2)$: Compute the addition of c_1 and c_2 as $c_{\mathrm{add}} = [c_1 + c_2]_q$.
- $\mathsf{YASHE.Mult}(c_1, c_2, \mathsf{evk})$: Compute

  $$\tilde{c}_{\mathrm{mult}} = \left[\left\lfloor \frac{t}{q} P_{w,q}(c_1) \otimes P_{w,q}(c_2) \right\rceil\right]_q \in R^{\ell^2_{w,q}},$$

  and output $c_{\mathrm{mult}} = \mathsf{YASHE.KeySwitch}(\tilde{c}_{\mathrm{mult}}, \mathsf{evk})$.

Since encryption and decryption are the same as in the Basic scheme from Section 3, the correctness bound does not change and Lemmas 1 and 2 hold for YASHE as well. Next, we analyze the homomorphic operations YASHE.Add and YASHE.Mult.

[1] Yet Another Somewhat Homomorphic Encryption scheme.

Homomorphic Addition. Given two ciphertexts $c_1, c_2 \in R$, which encrypt two messages m_1, m_2 with inherent noise terms v_1, v_2, their sum *modulo* q, $c_{\text{add}} = [c_1 + c_2]_q$, encrypts the sum of the messages *modulo* t, $[m_1 + m_2]_t$. Indeed, we can write $[m_1]_t + [m_2]_t = [m_1 + m_2]_t + tr_{\text{add}}$ for some $r_{\text{add}} \in R$ with $\|r_{\text{add}}\|_\infty \leq 1$. Since

$$\begin{aligned} f[c_1 + c_2]_q = fc_1 + fc_2 &= \Delta([m_1]_t + [m_2]_t) + (v_1 + v_2) \\ &= \Delta([m_1 + m_2]_t + tr_{\text{add}}) + (v_1 + v_2) \pmod q, \end{aligned}$$

we obtain $f[c_1 + c_2]_q = \Delta[m_1 + m_2]_t + (v_1 + v_2 - r_t(q)r_{\text{add}}) \pmod q$ because $\Delta t \equiv -r_t(q) \pmod q$. This means that the size of the inherent noise v_{add} of c_{add} is bounded by

$$\|v_{\text{add}}\|_\infty \leq \|v_1\|_\infty + \|v_2\|_\infty + r_t(q). \tag{1}$$

Up to the term $r_t(q) < t$, the inherent noise terms are added during homomorphic addition.

Homomorphic Multiplication. The homomorphic multiplication operation is divided into two parts. The first part describes a basic procedure to obtain an intermediate ciphertext that encrypts the product $[m_1 m_2]_t$ modulo t of two messages m_1 and m_2. However, the intermediate ciphertext can not be decrypted with Basic.Decrypt using the secret key f. The second part performs a procedure which allows a public transformation of this intermediate ciphertext to a ciphertext that can be decrypted with f. This latter procedure was introduced in [6] in the form of relinearization and was later expanded in [4] into a method called key switching, which transforms a ciphertext decryptable under one secret key to one decryptable under any other secret key. For our analysis, we assume that χ_{key} and χ_{err} are B_{key}- and B_{err}-bounded, respectively. Even if we work with unbounded Gaussian distributions, this is a valid assumption since elements drawn from either distribution have bounded norm for suitable bounds with high probability. The deduction of noise bounds mostly follows the basic multiplication section of [9], since ciphertexts and the decryption algorithm in YASHE have a very similar structure to those in [9].

First Step. Let $c_1, c_2 \in R$ be ciphertexts that encrypt messages $m_1, m_2 \in R$. In the first step of the homomorphic multiplication operation, we compute

$$\tilde{c}_{\text{mult}} = \left[\left\lfloor \frac{t}{q} P_{w,q}(c_1) \otimes P_{w,q}(c_2) \right\rceil \right]_q .$$

The following theorem shows that $\langle \tilde{c}_{\text{mult}}, D_{w,q}(f) \otimes D_{w,q}(f) \rangle = \Delta[m_1 m_2]_t + \tilde{v}_{\text{mult}} \pmod q$, and it provides a bound for the size of $\tilde{v}_{\text{mult}}$. Thus, $\tilde{c}_{\text{mult}}$ can be viewed as an encryption of $[m_1 m_2]_t$ under $D_{w,q}(f) \otimes D_{w,q}(f)$ if the inherent noise term $\tilde{v}_{\text{mult}}$ is small enough.

Theorem 1 (Multiplication Noise). *Let $c_1, c_2 \in R$ be ciphertexts encrypting $m_1, m_2 \in R$, decryptable with the secret key f. Let $v_1, v_2 \in R$ be the inherent noise terms in c_1, c_2 and let $V > 0$ such that $\|v_i\|_\infty \leq V < \Delta/2$,*

$i \in \{1,2\}$. Let $\tilde{c}_{\text{mult}}$ be the intermediate ciphertext in YASHE.Mult*, and let $\ell_{w,tB_{\text{key}}} = \lfloor \log_w(tB_{\text{key}}) \rfloor + 2$. Then $\langle \tilde{c}_{\text{mult}}, D_{w,q}(f) \otimes D_{w,q}(f) \rangle = \Delta[m_1 m_2]_t + \tilde{v}_{\text{mult}} \pmod q$ where*

$$\|\tilde{v}_{\text{mult}}\|_\infty < \delta t(2 + \delta \ell_{w,tB_{\text{key}}} w)V + \frac{\delta t^2}{2}(3 + \delta \ell_{w,tB_{\text{key}}} w) + \frac{1}{8}(\delta \ell_{w,tB_{\text{key}}} w)^2 + \frac{1}{2}.$$

Starting with two ciphertexts at a given inherent noise level, the first step of the multiplication increases the inherent noise level by a multiplicative factor of roughly $\delta^2 t \ell_{w,tB_{\text{key}}} w$ and an additive term of $\frac{\delta^2}{2} \ell_{w,tB_{\text{key}}} w(t^2 + \frac{1}{4} \ell_{w,tB_{\text{key}}} w)$.

Key Switching. The second part in the homomorphic multiplication procedure is a key switching step, which transforms the ciphertext $\tilde{c}_{\text{mult}}$ into a ciphertext c_{mult} that is decryptable under the original secret key f. We use the evaluation key

$$\mathsf{evk} = [f^{-1} P_{w,q}(D_{w,q}(f) \otimes D_{w,q}(f)) + \boldsymbol{e} + h \cdot \boldsymbol{s}]_q,$$

output by YASHE.KeyGen where $\boldsymbol{e}, \boldsymbol{s} \leftarrow \chi_{\text{err}}^{\ell_{w,q}^3}$ are vectors of polynomials sampled from the error distribution χ_{err} and $[\cdot]_q$ is applied to each coefficient of the vector. Note that this key is a vector of quasi-encryptions of $f^{-1} P_{w,q}(D_{w,q}(f) \otimes D_{w,q}(f))$ that depend on the secret key f, under its corresponding public key and that it is made public because it is needed for the homomorphic multiplication operation. Therefore, we need to make a circular security assumption, namely that the scheme is still secure even given that evk is publicly known (see Section 4.2). The following lemma deduces a bound on the noise caused by the key switching procedure and states an overall bound on the noise growth during a single homomorphic multiplication operation.

Lemma 3. *Let notation be as in Theorem 1 and as above. In particular, let $\tilde{c}_{\text{mult}}$ be the intermediate ciphertext in* YASHE.Mult *with inherent noise term $\tilde{v}_{\text{mult}}$. Let* evk *be the evaluation key and $c_{\text{mult}} =$* YASHE.KeySwitch*$(\tilde{c}_{\text{mult}}, \mathsf{evk})$. Then $f c_{\text{mult}} = \Delta[m_1 m_2]_t + v_{\text{mult}} \pmod q$, where*

$$\|v_{\text{mult}}\|_\infty < \|\tilde{v}_{\text{mult}}\|_\infty + \delta^2 t \ell_{w,q}^3 w B_{\text{err}} B_{\text{key}}.$$

Theorem 1 and Lemma 3 give an overall upper bound on the noise growth during a homomorphic multiplication. This clearly dominates the noise growth for homomorphic addition.[2]

4.1 Correctness

This section discusses the correctness of YASHE and shows that it is a leveled homomorphic encryption scheme. We state correctness by giving an asymptotic bound on the number of multiplicative levels in an arithmetic circuit that can be correctly evaluated. For this, we concretely focus on a parameter setting such

[2] As noted in [3] the number of elements in $D_{w,q}(f) \otimes D_{w,q}(f)$ can be reduced from $\ell_{w,q}^2$ to $\binom{\ell_{w,q}}{2}$ which correspondingly reduces the number of ring elements in evk.

that the assumptions of the theorem by Stehlé and Steinfeld (see [2, Appendix A]) hold. This means that the DSPR problem is hard in R_q. We therefore fix the following parameters: let d be a power of 2, $n = \varphi(d)$, $\epsilon \in (0, 1)$, $k \in (1/2, 1)$ and let $q = 2^{d^\epsilon}$ be a prime such that $\Phi_d(X) = X^n + 1$ splits into n irreducible factors modulo q. Let χ_{key} be a discrete Gaussian distribution on R_q with deviation $\sigma_{\text{key}} \geq d\sqrt{\log(8dq)} \cdot q^k$, and let χ_{err} be an asymptotically $\omega(\sqrt{d \log(d)})$-bounded Gaussian distribution on R where d tends to infinity. Finally, we fix $w = 2$ and $t = 2$, but note that similar results hold for general w, t – this restriction is merely for the purpose of exposition.

Theorem 2 (Correctness of YASHE). *For the parameter choices above,* YASHE *can evaluate any circuit of depth*

$$L = \mathcal{O}\left(\frac{(1-k)\log(q)}{\log(\log(q)) + \log(d)}\right).$$

4.2 Security

To prove security of YASHE, we need to assume that IND-CPA security can be maintained even when an adversary has access to elements of the evaluation key evk. Due to the way we construct evk it is not sufficient to simply replace f by L distinct secret keys f_i, as has been done in previous works – a specific assumption is still required. This is a form of key dependent message security, for the family of functions defining the evaluation key. Under this "circular security" assumption, the IND-CPA security of YASHE follows from the IND-CPA security of the scheme Basic described in Section 3 and the RLWE assumption.

Theorem 3 (Security of YASHE). *The scheme* YASHE *is IND-CPA secure under the* $\mathsf{RLWE}_{d,q,\chi_{\text{err}}}$ *assumption and the assumption that the scheme remains IND-CPA secure, even when an adversary has access to* evk *output by* $\mathsf{YASHE.KeyGen}(d, q, 2, \chi_{\text{key}}, \chi_{\text{err}}, 2)$.

Proof. Since $\sigma_{\text{key}} \geq d\sqrt{\log(8dq)} \cdot q^k$ for some $k > 1/2 + \nu$ with $\nu > 0$, the conditions of [2, Theorem 7] (see also [25]) are satisfied. Hence the public key is indistinguishable from a uniform element of $R_q^\times$. It follows from [25] that the scheme Basic is IND-CPA secure under the $\mathsf{RLWE}_{d,q,\chi_{\text{err}}}$ assumption in R_q. Under the circular security assumption outlined above, the IND-CPA security of YASHE follows. □

For the proof of Theorem 3, we only need parameters that satisfy the assumptions in [2, Theorem 7]. For the parameters outlined at the beginning of this subsection, the RLWE assumption is believed to be hard based on standard worst-case lattice problems.

4.3 From Leveled to Fully Homomorphic Encryption

In [10], Gentry showed how a fully homomorphic scheme can be obtained from a leveled homomorphic scheme supporting computation of circuits of sufficient

depth. If a scheme can evaluate its own decryption circuit and one additional multiplication, then that scheme can be converted to a fully homomorphic scheme. The only caveat is that we have to make an additional assumption: to execute the bootstrapping procedure, it is necessary to augment the public key with encryptions $\mathsf{YASHE.Encrypt}(\mathsf{pk}, \mathsf{sk}[j])$ of the bits of the secret key, under its corresponding public key. Similarly to the assumption on the evaluation key, we need to make an additional assumption that including encryptions of bits of the secret key does not affect security.

To achieve a fully homomorphic scheme, we simply view the decryption circuit as a circuit computed on the bits of the secret key at a ciphertext c we wish to refresh. The noise in the resulting *fresh* ciphertext will be of fixed size depending on the noise in the encryptions of the bits of the secret key. In the full version [2, (Lemma 6, Theorem 8)] we show that YASHE can be bootstrapped to a fully homomorphic scheme.

5 A More Practical Variant of the Scheme

In this section, we propose a more practical variant YASHE′ of YASHE. The difference to YASHE lies in the homomorphic multiplication procedure. In YASHE′, an intermediate ciphertext is simply a single polynomial while it is a vector of polynomials in YASHE. This results in an evaluation key that consists of only $\ell_{w,q}$ polynomials instead of $\ell_{w,q}^3$ for YASHE and thus in a simpler key switching procedure. We now state the scheme and discuss the noise growth during the simplified homomorphic multiplication operation YASHE′.Mult.

- $\mathsf{YASHE'.ParamsGen}(\lambda)$: Output $(d, q, t, \chi_{key}, \chi_{err}) \leftarrow \mathsf{BasicParamsGen}(\lambda)$.
- $\mathsf{YASHE'.KeyGen}(d, q, t, \chi_{key}, \chi_{err}, w)$: Compute

$$h, f \leftarrow \mathsf{Basic.KeyGen}(d, q, t, \chi_{\mathrm{key}}, \chi_{\mathrm{err}}).$$

  Sample $\boldsymbol{e}, \boldsymbol{s} \leftarrow \chi_{err}^{\ell_{w,q}}$, compute $\boldsymbol{\gamma} = [P_{w,q}(f) + \boldsymbol{e} + h \cdot \boldsymbol{s}]_q \in R^{\ell_{w,q}}$. and output $(\mathsf{pk}, \mathsf{sk}, \mathsf{evk}) = (h, f, \boldsymbol{\gamma})$.
- $\mathsf{YASHE'.Encrypt}(\mathsf{pk}, m)$: Encrypt $m \in R$ as $c \leftarrow \mathsf{Basic.Encrypt}(\mathsf{pk}, m) \in R$.
- $\mathsf{YASHE'.Decrypt}(\mathsf{sk}, c)$: Output the message $m \leftarrow \mathsf{Basic.Decrypt}(\mathsf{sk}, c) \in R$.
- $\mathsf{YASHE'.KeySwitch}(\tilde{c}_{\mathrm{mult}}, \mathsf{evk})$: Output the ciphertext $[\langle D_{w,q}(\tilde{c}_{\mathrm{mult}}), \mathsf{evk}\rangle]_q$.
- $\mathsf{YASHE'.Add}(c_1, c_2)$: Output $c_{add} \leftarrow \mathsf{YASHE.Add}(c_1, c_2) = [c_1 + c_2]_q$.
- $\mathsf{YASHE'.Mult}(c_1, c_2, \mathsf{evk})$: Output the ciphertext

$$c_{\mathrm{mult}} = \mathsf{YASHE'.KeySwitch}(\tilde{c}_{\mathrm{mult}}, \mathsf{evk}), \text{ where } \tilde{c}_{\mathrm{mult}} = \left[\left\lfloor \frac{t}{q} c_1 c_2 \right\rceil\right]_q.$$

For two ciphertexts $c_1, c_2 \in R$ that encrypt $m_1, m_2 \in R$, the intermediate ciphertext $\tilde{c}_{\mathrm{mult}}$ during homomorphic multiplication $\mathsf{YASHE'.Mult}$ satisfies $f^2 \tilde{c}_{\mathrm{mult}} = \Delta[m_1 m_2]_t + \tilde{v}_{\mathrm{mult}} \pmod q$ as shown in the following theorem. This means that $\tilde{c}_{\mathrm{mult}}$ is an encryption of $[m_1 m_2]_t$ under f^2. The theorem also provides an upper bound on the inherent noise term in the intermediate ciphertext. We assume that the error distribution χ_{err} is B_{err}-bounded and that the key distribution χ_{key} is B_{key}-bounded.

Theorem 4 (Multiplication Noise). *Let $c_1, c_2 \in R$ be ciphertexts encrypting $m_1, m_2 \in R$, which are decryptable with the secret key f. Let $v_1, v_2 \in R$ be the inherent noise terms in c_1, c_2 and let $V > 0$ such that $\|v_i\|_\infty \leq V < \Delta/2$, $i \in \{1, 2\}$. Let $\tilde{c}_{\text{mult}}$ be the intermediate ciphertext in* YASHE′.Mult. *Then $f^2\tilde{c}_{\text{mult}} = \Delta[m_1 m_2]_t + \tilde{v}_{\text{mult}} \pmod q$ where*

$$\|\tilde{v}_{\text{mult}}\|_\infty < \delta t(4 + \delta t B_{\text{key}})V + \delta^2 t^2 B_{\text{key}}(B_{\text{key}} + t).$$

Key Switching. The key switching algorithm now transforms such an intermediate encryption into a ciphertext that can be decrypted with f itself. The evaluation key is $\mathsf{evk} = [P_{w,q}(f) + \boldsymbol{e} + h \cdot \boldsymbol{s}]_q$, where $\boldsymbol{e}, \boldsymbol{s} \leftarrow \chi_{\text{err}}^{\ell_{w,q}}$ are vectors of polynomials sampled from the error distribution χ_{err}. Again, this key is a vector of quasi-encryptions of the secret key f under its corresponding public key. It is required for the homomorphic multiplication operation and is therefore made public. This means, we need to make a circular security assumption as for YASHE, namely that the scheme is still secure even given that evk is publicly known. The following lemma gives a bound on the key switching noise.

Lemma 4. *Let $\tilde{c}_{\text{mult}}$ be the intermediate ciphertext in* YASHE′.Mult. *Its inherent noise term is denoted by $\tilde{v}_{\text{mult}}$. Let $\boldsymbol{\gamma}$ be the evaluation key from above and $c_{\text{mult}} =$* YASHE′.KeySwitch$(\tilde{c}_{\text{mult}}, \boldsymbol{\gamma})$. *Then $fc_{\text{mult}} = \Delta[m_1 m_2]_t + v_{\text{mult}} \pmod q$, where*

$$\|v_{\text{mult}}\|_\infty < \|\tilde{v}_{\text{mult}}\|_\infty + \delta^2 t \ell_{w,q} w B_{\text{err}} B_{\text{key}}.$$

5.1 Correctness and Security of YASHE′

In the following theorem, we give an explicit bound for correctness of a homomorphic evaluation of an arithmetic circuit in R/tR of multiplicative depth L that is organized in a leveled tree structure of multiplications without any additions. At each level all ciphertexts are assumed to have inherent noise terms of roughly the same size. The bounds that we obtain might be too large and could be significantly reduced for computations that involve more additions and less multiplications as well as multiplications of ciphertexts with imbalanced inherent noise terms. In favor of simplicity, we restrict to the above setting.

Theorem 5 (Correctness of YASHE′**).** *Let $\epsilon_1 = 4(\delta t B_{\text{key}})^{-1}$. The scheme* YASHE′ *can correctly evaluate an arithmetic circuit consisting of L-levels of multiplications in R/tR on ciphertexts with inherent noise of size at most V that are arranged in a binary tree of L levels of multiplications if*

$$2(1+\epsilon_1)^{L-1}\delta^{2L} t^{2L-1} B_{\text{key}}^L((1+\epsilon_1)tV + L(tB_{\text{key}} + t^2 + \ell_{w,q} w B_{\text{err}})) < \Delta - r_t(q).$$

Appendix K in [2] gives detailed bounds on the increase of the inherent noise terms in ciphertexts during homomorphic addition and multiplication. One can take these bounds to deduce overall bounds for the exact computation that is supposed to be carried out on encrypted data. The obtained bounds can then be used to deduce tailored parameters for the scheme to ensure correctness and

security for that particular setting, possibly resulting in more efficient parameters for the specific computation.

The security of YASHE′ is based on the RLWE assumption and a circular security assumption similar to the one for YASHE. The price we pay for a simpler homomorphic multiplication operation lies in an additional security assumption. Since YASHE′ only works for a much narrower key distribution that does not satisfy the requirements for applying the Stehlé and Steinfeld result ([25, Thm. 4.1]), security also relies on the Decisional Small Polynomial Ratio (DSPR) assumption, as stated in Section 2. In YASHE, this assumption could be avoided by making the scheme work with a key distribution as demanded by [25]. Following the same hybrid argument as in [16], one can prove that the scheme described in this section is secure under the DSPR assumption and the RLWE assumption (see [16, Section 3.3]). If a, b are two elements sampled from a Gaussian with very small standard deviation or from a different distribution that yields polynomials with very small coefficients only, the ratio $h = a/b$ can clearly not be uniform because the number of elements for a and b is too small and produces only a small number of values for h when compared to all elements in R_q. Still, a computationally bounded adversary might not be able to distinguish such a case from uniform randomly chosen h.

Theorem 6 (Security of YASHE′). *Let d be a positive integer, q and $t < q$ be two moduli, w be a fixed positive integer, and let χ_{key} and χ_{err} be distributions on R. The scheme* YASHE′ *is IND-CPA secure under the* $\mathsf{RLWE}_{d,q,\chi_{\text{err}}}$ *assumption, the* $\mathsf{DSPR}_{d,q,\chi_{\text{key}}}$ *assumption, and the assumption that the scheme remains IND-CPA secure even when the evaluation key* evk *which is output by* $\mathsf{YASHE'.KeyGen}(d, q, t, \chi_{\text{key}}, \chi_{\text{err}})$ *is known to the adversary.*

Remark 1. The $\mathsf{DSPR}_{d,q,\chi_{\text{key}}}$ assumption can be replaced by a weaker assumption $\mathsf{DSPR}_{d,q,\chi_f,\chi_g}$, where the elements f and g that are used for the public key $h = [tgf^{-1}]_q$ are sampled from distributions of different width with bounds B_f and B_g, respectively. This new assumption can be made weaker than the original assumption since the element g can be sampled from a much wider distribution than f. Introducing these two distributions means that the noise bound for the inherent noise in a fresh ciphertext is changed to $\delta t(B_{\text{err}}(B_f + B_g) + r_t(q)B_f/2)$. The proofs of the noise bounds for YASHE′.Mult show that the bound B_g only influences the constant C_2 in Lemma 9 in the full version [2]. The contributions of B_g in the noise bounds for L levels of multiplications are merely a constant factor independent of L. Therefore, the scheme is still leveled homomorphic with the weaker assumption.

Remark 2. For YASHE′, since private keys are sampled with very small norm, the circular security assumption can be avoided in the usual way by providing a different public/private key pair (h_i, f_i) for each level i of multiplications for $0 \leq i \leq L$. The evaluation key has to be extended to L vectors $\gamma_i = [P_{w,q}(f_{i-1}^2) + \boldsymbol{e} + h_i \cdot \boldsymbol{s}]_q$, $1 \leq i \leq L$, such that the key switching step $\mathsf{YASHE'.KeySwitch}(\tilde{c}_{\text{mult}}, \mathsf{evk}_i)$ transforms the intermediate ciphertext $\tilde{c}_{\text{mult}}$ decryptable under f_{i-1}^2 (obtained from two ciphertexts at level $i-1$) into one decryptable under f_i at level i.

Table 1. Parameters that guarantee security of $\lambda = 80$ bits against the distinguishing attack with advantage $\epsilon = 2^{-80}$. We fix $w = 2^{32}$, the key distribution is assumed to be bounded by $B_{\text{key}} = 1$, and we use $\sigma_{\text{err}} = 8$ and $B_{\text{err}} = 6\sigma_{\text{err}}$. Either for fixed sizes of q, we give the minimal degree $n_{\min}$ (left part), or for fixed dimension n, we give the maximal size $\log(q_{\max})$ (right part). For each pair (q, n) according to the given sizes, and different values of t, correctness is guaranteed for at most $L_{\max}$ multiplicative levels.

$\lceil \log(q) \rceil$	$n_{\min}$	t	$L_{\max}$
128	3329	2	3
		256	2
		1024	1
192	5018	2	5
		256	3
		1024	3
256	6707	2	7
		256	5
		1024	4
512	13463	2	15
		256	10
		1024	9
1024	26974	2	31
		256	21
		1024	19

n	$\log(q_{\max})$	t	$L_{\max}$
2^{12}	157	2	4
		256	2
		1024	2
2^{13}	312	2	9
		256	6
		1024	5
2^{14}	622	2	19
		256	13
		1024	11
2^{15}	1243	2	37
		256	25
		1024	23
2^{16}	2485	2	71
		256	50
		1024	46

5.2 Parameters

In this section, we give suggestions for choosing concrete parameters which can be used as a guideline to instantiate practical schemes with varying complexity. There are multiple parameters one can adjust, so we restrict ourselves to a subset of choices which we think are most relevant. We consider two settings. In the first, we fix a specific size for the modulus q. This is interesting for instance when a fast modular multiplication implementation (in either hard- or software) is already available, and one prefers to use this to boost the scheme's performance. We fix different sizes for the modulus q starting from 128 bits up to 1024 bits. The other setting focuses on special-purpose polynomial arithmetic. Here, we fix the degree $n = \varphi(d)$ to be a power of 2 between 2^{12} and 2^{16}.

The parameters presented in Table 1 are obtained by following the security analysis of Lindner and Peikert [15] under the assumption that the results from [15] in the LWE setting carry over to the RLWE setting, and assuming that the assumptions in Section 5.1 hold. This analysis is similar to the ones from [12,9,14] and we refer to [12] for a more complete discussion of assumptions made in deriving parameters. Note that recent results by Chen and Nguyen [7] are considered to be more accurate for estimating the security of specific parameters using the simulation of the BKZ 2.0 algorithm for assessing the runtime of lattice basis reduction. Selecting parameters for YASHE′ with this method is ongoing work at the time of writing this paper. However, it is expected that the parameters presented in this paper which are obtained by using the Lindner-Peikert method are more conservative than those obtained with the BKZ 2.0 simulation.

Next, we discuss in more detail the parameter selection recommendations made in Table 1. We use $B_{\text{key}} = 1$, in other words we are assuming that even when the polynomials f', g have coefficients in $\{-1, 0, 1\}$, the public key $h = [tgf^{-1}]_q$ is indistinguishable from uniform. The standard deviation of the error distribution is fixed at $\sigma_{\text{err}} = 8$; this is consistent with [19]. The high probability bound on the size of the coefficients of errors drawn from Gaussian distributions is chosen as $6\sigma_{\text{err}}$.

To distinguish with an advantage of ϵ in the RLWE problem, an adversary is required to find vectors of length at most $\alpha \cdot (q/\sigma)$ where $\alpha = \sqrt{\ln(1/\epsilon)/\pi}$. In our specific parameter examples, we use $\epsilon = 2^{-80}$, which results in $\alpha \approx 4.201$. We refer to [15] for a more complete description of a distinguishing attack and the precise lattices we are required to find short vectors in. Running Schnorr-Euchner's BKZ [22], the best known lattice reduction algorithm in practice, and its successor BKZ 2.0 [7] for security parameter λ (following [12] we use $\lambda = 80$) one expects to find vectors of length $2^{2\sqrt{n \log_2(q) \log_2(\delta_{\text{RHF}})}}$ in time $T_{\text{BKZ}} = 2^{\lambda}$ where δ_{RHF} is the so-called root Hermite factor. This latter quantity is the overwhelming factor determining the quality of the basis which can be achieved in a given time and is computed as in [15] from

$$\log_2(T_{\text{BKZ}}) = 1.8/\log_2(\delta_{\text{RHF}}) - 110.$$

It is currently infeasible to achieve a target root Hermite factor $\delta_{\text{RHF}} < 1.005$ [7]. To guarantee security, we require that the shortest vector obtained through lattice reduction is longer than a vector which could give an adversary a non-negligible advantage ϵ in the Ring-LWE distinguishing problem. This means that for security we thus require

$$\alpha \cdot q/\sigma < 2^{2\sqrt{n \log_2(q) \log_2(\delta_{\text{RHF}})}}.$$

For fixed parameters α and δ_{RHF}, this inequality provides bounds on the remaining parameters q, σ_{err} and n. Fixing σ_{err} too ($\sigma_{\text{err}} = 8$ here), we get a dependency between q and n that is expressed in the two settings discussed above as follows. When we fix q, we obtain a lower bound $n_{\min}$ for the dimension n to guarantee security against the distinguishing attack. For the example values for the sizes of q given in the first column of the left part of Table 1, we list this minimal degree in the second column. We used the worst case bound for a modulus q of that size. Vice versa, first fixing the degree n means that we get an upper bound $q_{\max}$ for q. We display the relation between n and the size $\log(q_{\max})$ in the first two columns of the right part of Table 1.

For guaranteeing correctness, we use the noise bounds derived in the previous section. As mentioned in Section 2, when d is a power of 2 and thus $\Phi_d(X) = X^n + 1$, the expansion factor is $\delta = n$. Then, by Lemma 1 and Lemma 9 in the full version [2] we know that our scheme can correctly evaluate a depth L circuit as long as

$$(1+\epsilon_1)^{L-1} n^{2L} t^{2L-1} B_{\text{key}}^{L} \left((1+\epsilon_1) tV + L\left(t(B_{\text{key}} + t) + \ell_{w,q} w B_{\text{err}}\right)\right)$$

is less than $(\Delta - r_t(q))/2$, where $\epsilon_1 = 4(ntB_{\text{key}})^{-1}$ and $V = ntB_{\text{key}}(2B_{\text{err}} + r_t(q)/2)$ is the inherent noise of fresh ciphertexts by Lemma 2. For each row in

either the left or the right part of Table 1, we take the given values for q and n together with different values for t and check what is the maximum number of levels $L_{\max}$ for which the correctness bound still holds. Note that in the left part, we take the minimal degree $n_{\min}$. This means that when choosing a power of 2 for the degree, the values for $L_{\max}$ might change. In the right part, we take the largest possible value for q with the given maximal bit size.

It is important to ensure that the security bounds as well as the correctness bounds are both satisfied. Note that the authors of [9] failed to check their parameters presumably obtained from the correctness bound in the security bound, too, resulting in insecure parameters of $q = 2^{1358}$ and $n = 2^{10}$.

5.3 Implementation

Currently there are not many known implementation results for FHE schemes. Some of those which have been published demonstrate that the current state-of-the-art's performance is still rather unsatisfactory, see for example the implementations which are capable of computing AES homomorphically [12,8]. Other people have focused on implementing relatively simple schemes that require only a few levels of multiplications [14]. When using the ring $R = \mathbb{Z}[X]/(X^{4096} + 1)$, $t = 2^{10}$ and a 130-bit prime q, the authors of [14] present implementation results on an Intel Core 2 Duo running at 2.1 GHz. Encryption takes 756 ms, addition of ciphertexts 4 ms, multiplication of ciphertexts 1590 ms (this includes the degree reduction) and decryption 57 ms.

We have implemented the YASHE′ variant proposed in Section 5 in a C-library. All the arithmetic has been built from scratch and we do not depend on any external number theory library. Using almost the same parameters (we use a 127-bit prime q) with $w = 2^{32}$ we obtained the following results on an Intel Core i7-3520M at 2893.484 MHz with hyperthreading turned off and over-clocking ("turbo boost") disabled. Encryption runs in 79.2 million cycles (27 ms), addition of ciphertexts in 70 thousand cycles (0.024 ms), multiplication of ciphertexts (including the key-switching) in 90.7 million cycles (31 ms) and decryption in 14.1 million cycles (5 ms).

This performance increase by at least one order of magnitude (for the decryption) to two orders of magnitude (for the addition of ciphertexts) can be partially explained by the fact that we are running on a more recent processor and that we implemented the scheme directly in C (avoiding the overhead incurred by using a computer algebra system as in [14]). The remainder of the speed-up is due to our newly proposed scheme, in particular due to a simpler multiplication operation on ciphertexts that uses a more compact evaluation key consisting of only 4 elements. These performance numbers highlight the fact that HE is much more practical for schemes which do not require very deep circuits (like AES) but instead only need a few (around 2^2 to 2^5) multiplications.

5.4 Truncating Ciphertext Words

Brakerski [3, Section 4.2] first suggested for his scale-invariant LWE scheme to discard some least significant bits of the ciphertext. Based on this idea, we

describe an optimization to our scheme which significantly reduces both the ciphertext length and the number of elements in the evaluation key. By aligning the number of bits we discard with a multiple of w used in YASHE.KeySwitch, the number of elements required to switch keys is reduced per multiplication.

Define $\mathsf{YASHE.Discard_w}(c, i)$ as the function which takes as input a ciphertext and the number $0 \leq i < \ell_{w,q}$ of w-words to be truncated and outputs $c' = \mathsf{YASHE.Discard_w}(c, i) = \lfloor w^{-i} c \rfloor$. Then, $w^i c'$ is equal to c with the i least significant w-words of c being set to 0. If $cf = \Delta m + v \pmod q$, then $w^i c' f = \Delta m + v' \pmod q$ with $\|v'\|_\infty \leq \|v\|_\infty + \frac{1}{2}\delta w^i \|f\|_\infty$. For a constant $B > 0$ such that $2B > \delta\|f\|_\infty/2$, if we discard $\log_w(2B) - \log_w(\delta\|f\|_\infty)$ words, we incur an additional noise term of size B, but the ciphertext can now be represented by $\log_w(q/B) + \log_w(\delta\|f\|_\infty/2)$ words. This means that, with discarding, the length of ciphertexts does not depend on the absolute value of q but only on the ratio of q to the noise in the ciphertext. Perhaps more importantly, this means that when we consider $D_{w,q}(c)$ for a ciphertext c with coefficients represented by roughly $\log_w(q/B)$ words, all the lowest $\log_w(B)$ words are now zero. If c is a ciphertext decryptable under f^2, in the key switching step, we only need the top $\log_w(q/B)$ elements from the evaluation key to carry out the switch.

5.5 Encoding Input Data via the CRT

For our leveled homomorphic encryption scheme, we have given bounds on parameters and input data to ensure correctness and security. For applications such as outsourcing of storage and computation on private data to the cloud, it could be the case that the user requires a flexible system which allows for additional computation, more computation than was planned for when setting system parameters. We propose a way to extend the system to allow additional computation without resetting the parameters. For computations on integer values, the encoding of larger integers using the Chinese Remainder theorem allows for either greater precision of computation or larger integer inputs, using the same underlying field size and lattice dimension but at the cost of increasing the number of ciphertexts to be operated on.

Integer computations with results up to a bound B are done by encoding each input as a collection of integers modulo coprime t_i via the CRT. Computations are then carried out on the collection and correctly reflect the integer operations not involving any modular reductions, as long as the product of the t_i is greater than B. Each integer in the collection is encrypted as a separate ciphertext with respect to its corresponding plain text modulus t_i and those ciphertexts can be processed in parallel to return encrypted collections. After they are decrypted, the CRT is used to recover the output as an actual integer.

This approach is different than the ones introduced in [23] and [8], since in contrast to these schemes, we do not use the CRT to pack information into different plain text slots of a single ciphertext. Instead, we simply encrypt each part of the CRT encoding in a separate ciphertext with respect to its plain text modulus t_i. This introduces a different way of flexibility. Ciphertexts now consist of several ring elements, but can be processed in parallel. For example, this allows

to work on integers of double bit length by keeping the same parameters, only extending to two ciphertexts with different values for t_0 and t_1.

6 Conclusions

We have proposed a new fully homomorphic encryption scheme based on the scheme by Stehlé and Steinfeld which removes the non-standard decisional small polynomial ratio assumption needed in the homomorphic encryption scheme by López-Alt, Tromer and Vaikuntanathan. Hence, the security is solely based on standard lattice assumptions and a circular security assumption. Our new scheme avoids modulus switching and keeps the size of ciphertexts to a single ring element. Furthermore, we have presented a more practical variant of our scheme which does need the decisional small polynomial ratio assumption. For this latter scheme we presented parameters and implementation results.

Acknowledgments. The authors thank Adriana López-Alt for many useful suggestions and discussions, in particular for pointing out the possibility of a weaker assumption in Remark 1, Tancrède Lepoint for his comments and for noticing an error in an earlier version of Table 1, Nigel P. Smart for helpful advice and the anonymous reviewers for their constructive feedback.

References

1. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In: Halevi, S. (ed.) CRYPTO 2009. LNCS, vol. 5677, pp. 595–618. Springer, Heidelberg (2009)
2. Bos, J.W., Lauter, K., Loftus, J., Naehrig, M.: Improved security for a ring-based fully homomorphic encryption scheme (full version). Cryptology ePrint Archive, Report 2013/075 (2013), `http://eprint.iacr.org/`
3. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical gapSVP. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 868–886. Springer, Heidelberg (2012)
4. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: Fully homomorphic encryption without bootstrapping. In: ITCS, pp. 309–325 (2012)
5. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS, pp. 97–106 (2011)
6. Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505–524. Springer, Heidelberg (2011)
7. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 1–20. Springer, Heidelberg (2011)
8. Cheon, J.H., Coron, J.-S., Kim, J., Lee, M.S., Lepoint, T., Tibouchi, M., Yun, A.: Batch fully homomorphic encryption over the integers. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 315–335. Springer, Heidelberg (2013)
9. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012), `http://eprint.iacr.org/`

10. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, pp. 169–178 (2009)
11. Gentry, C., Halevi, S.: Implementing Gentry's fully-homomorphic encryption scheme. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 129–148. Springer, Heidelberg (2011)
12. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the AES circuit. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 850–867. Springer, Heidelberg (2012)
13. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: A ring-based public key cryptosystem. In: Buhler, J.P. (ed.) ANTS 1998. LNCS, vol. 1423, pp. 267–288. Springer, Heidelberg (1998)
14. Lauter, K., Naehrig, M., Vaikuntanathan, V.: Can homomorphic encryption be practical?. In: Cachin, C., Ristenpart, T. (eds.) CCSW, pp. 113–124. ACM (2011)
15. Lindner, R., Peikert, C.: Better key sizes (and attacks) for LWE-based encryption. In: Kiayias, A. (ed.) CT-RSA 2011. LNCS, vol. 6558, pp. 319–339. Springer, Heidelberg (2011)
16. López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: STOC, pp. 1219–1234 (2012)
17. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 1–23. Springer, Heidelberg (2010)
18. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. In: SIAM J. on Comp., pp. 372–381. IEEE Computer Society (2004)
19. Micciancio, D., Regev, O.: Lattice-based cryptography. In: Bernstein, D.J., Buchmann, J., Dahmen, E. (eds.) Post-Quantum Cryptography, pp. 147–191. Springer, Heidelberg (2009)
20. Coron, J.-S., Mandal, A., Naccache, D., Tibouchi, M.: Fully homomorphic encryption over the integers with shorter public keys. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 487–504. Springer, Heidelberg (2011)
21. Rivest, R.L., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. In: Foundations of Secure Computation, vol. 4, pp. 169–180. Academic Press, New-York (1978)
22. Schnorr, C.-P.: A hierarchy of polynomial time lattice basis reduction algorithms. Theor. Comput. Sci. 53, 201–224 (1987)
23. Smart, N., Vercauteren, F.: Fully homomorphic SIMD operations. Cryptology ePrint Archive, Report 2011/133 (2011), `http://eprint.iacr.org/`
24. Smart, N.P., Vercauteren, F.: Fully homomorphic encryption with relatively small key and ciphertext sizes. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 420–443. Springer, Heidelberg (2010)
25. Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 27–47. Springer, Heidelberg (2011)
26. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 24–43. Springer, Heidelberg (2010)