

Flattening NTRU for Evaluation Key Free Homomorphic Encryption

Yarkin Doröz¹ and Berk Sunar¹

Worcester Polytechnic Institute
{ydoroz,sunar}@wpi.edu

Abstract. In less than a decade fully homomorphic encryption has made significant advances. Despite all these improvements it is quite a challenge to reduce the parameter sizes and specifically evaluation keys. Eliminating the need for such prohibitively large evaluation keys and expensive noise management techniques has become a significant thrust among homomorphic encryption researchers. In a notable attempt, Gentry, Sahai, and Waters (GSW) introduced a scheme based on the approximate eigenvector problem that eliminates evaluation keys and costly key switching operations. In another very recent development, the Subfield Lattice Attack was introduced by Albrecht, Bai, and Ducas showing that the asymptotic security level with narrow key distributions may be far less than assumed in NTRU based FHE proposals.

In this paper, we propose a new FHE scheme **F-NTRU** that adopts the flattening technique proposed in GSW to derive an NTRU based scheme that (similar to GSW) does not require evaluation keys or key switching. Our scheme eliminates the decision small polynomial ratio (DSPR) assumption but relies only on the standard R-LWE assumption. Our scheme uses wide key distributions, and hence is immune to the Subfield Lattice Attack. We provide implementation results which show reasonable evaluation times compared to existing schemes while eliminating the need for storing and managing costly evaluation keys.

Keywords: Homomorphic Encryption, Flattening, NTRU, DSPR, Subfield Lattice Attack

1 Introduction

The notion of fully homomorphic encryption (FHE) scheme stayed as an open question for a few decades since its introduction by Rivest et al. [36]. In the meantime, numerous schemes featured limited homomorphic functionality, e.g. restricted to evaluate a non-universal set of gates, or prohibitive growth in noise or ciphertext sizes restricting the evaluation depth. These schemes are referred commonly referred to as *partially* homomorphic cryptosystems. Some of these schemes and their homomorphic properties are Elgamal encryption system [18] additive homomorphism in \mathbb{Z}_q , Goldwasser-Micali cryptosystem [26] additive homomorphism in \mathbb{Z}_2 , RSA cryptosystem [37] multiplicative homomorphism on \mathbb{Z}_q

and Paillier cryptosystem [35] additive homomorphism in \mathbb{Z}_q . Later, in 2009 the first working FHE scheme was constructed by Gentry [20, 21]. However, the proposed scheme was lacking in performance and it was not possible to implement anything practical. For instance, its most crucial operation called *bootstrapping*, which is used to restore the noise in ciphertext to an acceptable level, was taking 30 seconds in this very first implementation. After the introduction of the first FHE scheme, we witnessed the introduction of many new FHE constructions aimed at bringing FHE closer to real-world applications and with fewer assumptions. Several integer-based constructions and learning with error (LWE)-based constructions may be found in [15, 10, 11] and in [5, 22, 23], respectively.

These recent constructions brought along an impressive array of optimization techniques that can compete with the expensive bootstrapping operation. Such one scheme, which is based on LWE, was constructed by Brakerski, Gentry and Vaikuntanathan (BGV) [4]. The scheme uses a method called *modulus switching* to mitigate noise growth in ciphertexts. By applying modulus switching at each multiplicative level, exponential noise growth is prevented thereby maintaining noise at a fixed level throughout the homomorphic evaluation levels. The scheme was implemented as a software library HELib [27] using C++. The library was used by Gentry, Halevi and Smart [24] to implement the AES circuit homomorphically. The authors also introduced SIMD techniques [39] to batch multiple messages and process homomorphic AES circuits in parallel. They achieved amortized runtime of 2 seconds for 120 parallel blocks. Later, a new technique is introduced by Brakerski [3] that is applicable to LWE schemes. It uses tensor products to decrease noise growth from quadratic to linear.

Recently, another leveled FHE scheme was presented by López-Alt, Tromer, Vaikuntanathan (LTV) in [31]. Their scheme is based on a variant of NTRU [28] constructed earlier by Stehlé and Steinfeld [40]. The scheme is capable of processing homomorphic functions for various users with each having different public-secret key pairs. The authors introduced a new key switching technique called *relinearization*. By using it alongside with modulus switching, they are capable to mitigate the noise growth and keep the growth linear in size over the levels. However, the relinearization method requires evaluation keys to operate for each level of computation. This brings a prohibitive memory requirement especially for *deep* circuit evaluations.

The NTRU variant in [40] was later modified and implemented by Bos et al. in [1]. They adopt the tensor product technique in [3] and achieved a scale-invariant scheme with limited noise growth on homomorphic operations. Also, with the use of the tensor product technique the authors managed to improve the security of the LTV scheme [31] by using much higher levels of noise and thereby remove the Decisional Small Polynomial Ratio (DSPR) assumption. Instead the scheme relies only on the standard lattice reductions as in [40]. However, as the authors also note, the YASHE scheme brings in large evaluation key and complicated key switching procedure. Specifically, the size of the evaluation keys are $\ell_{w,q}^3$ in which q is modulus, w is radix and $\ell = \lfloor \log q \rfloor$. In [1] the authors introduce a modification (YASHE') to their scheme to eliminate the problems of expensive

tensor product calculations and large evaluation keys. However, this modification re-introduces the DSPR assumption due to increase in noise.

Another modified LTV-FHE implementation was presented by Doröz et al. in [16]. The security of their scheme depends on the DSPR and R-LWE assumptions as in [31]. Their implementation uses the relinearization and modulus switching methods as in [31] to cope with noise. They use a ring structure to create a specialized modulus, i.e. select modulus as $q = p^L$ in which p is a prime number. With the specialization the authors managed to significantly reduce the evaluation key size by creating evaluation keys only for the first level and by reusing them for to other levels simply by computing modular reduction. This reduces the size of the evaluation keys significantly. They evaluate AES homomorphically and achieved amortized runtime of 51 seconds per block.

Motivated by the complex noise management techniques, e.g. *relinearization*, *modulus switching*, *bootstrapping*, of the earlier FHE schemes Gentry, Sahai and Waters [25] proposed a new scheme based on the approximate eigenvector problem. The system uses matrix additions and multiplications which makes it asymptotically faster. At first, they define the GSW scheme as a *somewhat homomorphic scheme* since for a depth L circuit with B -bounded parameters the noise grows with a double exponential B^{2^L} . To convert the scheme into a leveled FHE, they introduce a *Flatten* operation which decomposes the ciphertext entries into bits. The secret key is also kept in a special powers of two form. With these modifications, the noise performance is improved greatly. For a depth L circuit with B -bounded secret key entries and 1-bounded (flattened) ciphertexts, the error magnitude is at most $(N + 1)^L B$ for $N = \log(q)(n + 1)$. However, ciphertexts still take a considerable space $\Theta(n^2 \log(q)^2)$ and as noted by GSW [25] the scheme may not be as efficient as existing leveled schemes in practice.

Applications. The advancements in FHE constructions led many researchers to experiment on practicality of various homomorphic applications. Lauter et al. investigates evaluation of homomorphic averages, standard deviations, and logistical regressions, which are underlying tools for real-world applications, in [34]. They also investigate evaluation of genomic data algorithms like Pearson Goodness-of-Fit test, the D' and r^2 -measures of linkage disequilibrium, the Estimation Maximization (EM) algorithm for haplotyping, and the Cochran-Armitage Test for Trend homomorphically in [29]. Another medical application is presented by Bos et. al in [2]. They implement a cardiovascular disease prediction service homomorphically which patients send their health information in secret and an analysis is performed in the cloud server. Aside from the medical there are other FHE applications as well. Like in [17] Doröz et al. proposed a bandwidth efficient private information retrieval (PIR) scheme using the NTRU based homomorphic encryption scheme. Another study by Çetin et al. exposes the limitations of word-size arithmetic in FHE and suggests solutions [7], i.e. convergence based iterative division, comparison, zero check and thresholding algorithms. The same authors study homomorphic sorting algorithms with respect to multiplicative depth and proposes a new scheme that is more suitable for FHE implementations with low multiplicative depth; and homomorphic au-

to complete, in [8] and [6], respectively. In another work, Dai et al. implement existing schemes using GPU to achieve significant acceleration. In [14] they implemented sorting and achieved 12 – 41 times speedup over [8] depending on the sorting size. In [13] they implemented the PIR scheme of [17] and achieved 13–34 times speedup depending on the database size. Lastly, Dai et al. implemented AES and Prince block ciphers in [12] and achieved 7.3 and 1.28 seconds of runtime, respectively.

The Subfield Lattice Attack. To overcome the efficiency bottleneck of FHE numerous optimizations along with customized selections of parameters have been proposed in the literature, e.g. special form lattice dimensions and moduli permitting efficient number theoretical transforms and batching, assumptions regarding the distributions of noise and parameters, etc. Indeed, these optimizations yielded significant advances in performance sizes over the years, e.g. faster evaluation with smaller evaluation key.

On the downside many of these assumptions are still open to debate from a security point of view. A very recent work by Albrecht, Bai and Ducas [33] painfully demonstrated this fact. The authors exploit the presence of a subfield to solve the NTRU problem for large moduli q and show that when the NTRU parameters are chosen poorly then the DSPR problem is not as hard as believed thereby invalidating the underlying security assumption in LTV [31, 16] and YASHE' [1]. Thus, the asymptotic security of both schemes is significantly reduced. The authors suggest a more cautious choice of parameters to minimize the effect of the attack. Even with such a remedy, the lattice dimension needs to be increased to restore the projected security level. Moreover, the selection of parameters renders batching rather difficult further diminishing the performance of both schemes.

Our Contribution. In this work, we present a new leveled FHE scheme F-NTRU that is based on the Stehlé and Steinfeld variant of NTRU [40] and adopts Flattening noise management technique introduced in (GSW) [25]. In a nutshell, we are able to achieve the best of both implementations in [1], i.e. elimination of expensive evaluation keys and reduction of security assumptions. Specifically,

- We introduce the first NTRU based FHE implementation that uses the *Flattening* method of GSW. Our scheme uses the encryption methods of NTRU for ciphertext creation, converts them into a matrix structure and makes use of the *Flattening* noise management technique.
- Similar to the YASHE construction [1] our scheme uses a wide key distribution and hence only relies on standard lattice reductions as in [40] (no DSPR assumption). Thus, our scheme is *immune to the Subfield Lattice Attack* by Albrecht, Bai and Ducas [33].
- Using the *Flattening* technique, we are able to multiply ciphertexts with only linear increase in the size of the noise level. Our scheme does not use any expensive noise reduction techniques such as relinearization and the need for prohibitively large evaluation keys.

- Our construction does not require evaluation keys whereas the YASHE construction has evaluation keys whose size grows $\tilde{O}(L^4)$ where L represents the evaluation depth. This makes it impossible to use YASHE in deep evaluations. achieve significantly small ciphertext size compared to the YASHE construction for the same level of multiplicative depth support.
- We introduce a variant of the scheme that reduces the ciphertext size and improves the homomorphic evaluation speed at the expense of the DSPR assumption.
- We provide implementation results and a comparison to YASHE.

2 Overview of NTRU Based FHE Schemes and GSW

In this section we briefly summarize the schemes relevant to our construction. We also briefly summarize schemes that we compare against later.

2.1 Stehlé and Steinfeld’s NTRU Variant

In [40] Stehlé and Steinfeld introduced a modification to the NTRU [28] scheme to make the scheme secure under the assumed quantum hardness of the standard worst-case lattice problem. The scheme fixes the ring to $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$ where n is a power of 2 and chooses $q \leq \text{Poly}(n)$. Also, chooses a message space \mathbb{Z}_p . Using a discrete Gaussian distribution random samples f', g are chosen where the secret key becomes $f = pf' + 1$. Set public key $h = pf^{-1}g$. A message μ is encrypted by computing $c = hs + pe + \mu \pmod q$ where s, e are Gaussian distribution samples. To decrypt we simply compute $\mu = c \cdot f \pmod p$.

2.2 DHS Scheme

Recently, Doröz, Hu and Sunar implemented single-key variation of the multi-key LTV-FHE scheme. This scheme is introduced in 2012 by López-Alt, Tromer and Vaikuntanathan (LTV) [31] as a full-fledged, multi-key and leveled fully homomorphic encryption scheme which is based on the NTRU cryptosystem. The NTRU scheme is a public key cryptosystem that is introduced by Stehlé and Steinfeld [40]. New additions to the NTRU scheme are operations for noise and key control are relinearization and modulus switching. Another addition to the scheme is adding random noise to the public key for security.

The DHS implementation has the operational ring $R_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$ in which q is our modulus and n is the degree of the polynomial. The scheme uses a truncated Gaussian error distribution function χ for sampling polynomials. The truncated distribution χ is B -bounded which means coefficient sizes are between range $[-B, B]$. The implementation has following four primitive functions:

- **KeyGen.** Using the security parameter λ , we create a sequence of modulus for each level as $q_i = q^{d-i}$ in which q is a prime. Later, we sample polynomials

$g \in \chi$ and $f' \in \chi$ and compute secret key $f = 2f' + 1$ and public key $h = 2gf^{-1}$ in \mathbb{Z}_{q_0} . In the next step, we compute the evaluation keys $\zeta_\tau^{(0)}(x) = hs_\tau + 2e_\tau + 2^\tau f$ which $\{s_\tau, e_\tau\} \in \chi$ and $\tau = [0, \lceil \log q_0 \rceil]$. Computed evaluation keys are for level index 0. For the rest of the levels we recycle the evaluation keys using the ring structure. To use the evaluation keys for level index i , we simply compute $\zeta_\tau^{(i)}(x) = \zeta_\tau(x) \bmod q_i$. This reduces the memory requirement significantly.

- **Encrypt.** In order to encrypt a message bit μ for level i , we compute $c^{(i)} = h^{(i)}s + 2e + b$ which $\{s, e\} \in \chi$ and $h^{(i)} = h^{(0)} \bmod q_i$.
- **Decrypt.** In order to decrypt at level i , we compute $\mu = \lceil c^{(i)} f^{(i)} \rceil_{q_i} \pmod{2}$.
- **Evaluation.** Homomorphic evaluation requires relinearization and modulus switching after each multiplication or an addition operation. Since the additive operations do not create significant noise like multiplications, we apply the noise reduction techniques only after a multiplication. Relinearization is computed as $\tilde{c}^{(i)}(x) = \sum_\tau \zeta_\tau^{(i)}(x) \tilde{c}_\tau^{(i-1)}(x)$. The polynomials $\tilde{c}_\tau^{(i-1)}(x)$ have the form of $\tilde{c}^{(i-1)}(x) = \sum_\tau 2^\tau \tilde{c}_\tau^{(i-1)}(x)$. Relinearization is followed by modulus switching which we compute as $\tilde{c}^{(i)}(x) = \lfloor \frac{q_i}{q_{i-1}} \tilde{c}^{(i-1)}(x) \rfloor_2$. This reduces the noise level by $\log(q_i/q_{i-1})$ bits. In modulus switching we need to match the parity bits of the messages between the old and new moduli: $\lfloor \cdot \rfloor_2$.

Note that the DHS implementation uses cyclotomic polynomials $\Phi_m(x)$ to define the ring instead of $\langle x^n + 1 \rangle$. The degree of the polynomial is $n = \phi(m)$ where ϕ is the Euler's totient function. This setup allows batching of multiple messages into the same ciphertext polynomial and thereby enables parallel processing.

2.3 BLLN

The BLLN scheme was introduced by Bos et al. in [1]. BLLN is based of the scheme proposed by Stehlé and Steinfeld [40]. The authors create a scheme called YASHE by using the construction of [40] and modify it by applying the tensor product technique of [3] to curb the noise growth in multiplications. With this it becomes possible to use a high level of noise in encryptions with which the schemes becomes DSPR hard as in [40]. However, this brings large evaluation keys into the scheme and it makes difficult to use in practice: the evaluation key consists of $\ell^3 = \mathcal{O}((\log(q))^3)$ ciphertexts. To mitigate this problem, in the same reference the authors introduced another scheme called YASHE'. They discard the tensor product and decrease the size of the evaluation keys, i.e. ℓ ciphertexts. However, they have to reduce the noise levels on fresh ciphertexts which brings back the DSPR security assumption. In the following, we first give the Basic scheme and later explain YASHE and YASHE' constructions using the Basic scheme.

Basic. We set the ring as $R = \mathbb{Z}[X]/\langle x^n + 1 \rangle$, t as the plaintext modulus, χ_{err} and χ_{key} as the Gaussian distribution for sampling. The scheme has the following primitive functions:

- **ParamsGen.** For security parameter λ we create $n = n(\lambda)$, $q = q(\lambda)$, $\chi_{\text{key}} = \chi_{\text{key}}(\lambda)$ and $\chi_{\text{err}} = \chi_{\text{err}}(\lambda)$.
- **KeyGen.** We sample $f', g \in \chi_{\text{key}}$. Set secret key $f = tf' + 1$ and public key $h = tgf^{-1}$.
- **Encrypt.** To encrypt message μ , we sample $s, e \in \chi_{\text{err}}$ and compute $c = \lfloor q/t \rfloor \mu + e + hs$.
- **Decrypt.** To decrypt a message simply compute $\mu = \lfloor \frac{t}{q} fc \rfloor$ where $\lfloor \cdot \rfloor$ is rounding to the nearest.

There are 3 notations need explanation before we summarize the YASHE scheme. The authors use \otimes for the tensor product. The notation $P_{w,q}$ is used to convert a number to an array with powers of w , i.e. $P_{w,q}(x) \rightarrow (x, xw, xw^2, \dots, xw^{\ell_{w,q}-1})$ for $\ell_{w,q} = (\log_w q + 2)$. Last notation is $D_{w,q}$ which decomposes a value into its word sizes, i.e. $D_{w,q}(x) \rightarrow (x_0, x_1, \dots, x_{\ell_{w,q}-1})$ where $x = \sum_{i=0}^{\ell_{w,q}-1} x_i w^i$.

YASHE. The Basic scheme explained above is used to construct a leveled FHE scheme. The primitive operations of the scheme is as follows:

- **ParamsGen.** The parameters are selected the same way as in Basic scheme.
- **KeyGen.** The public and secret key pairs are selected using the Basic.KeyGen routine, i.e. $h, f \leftarrow \text{Basic.KeyGen}$. Sample $\mathbf{e}, \mathbf{s} \in \chi_{\text{err}}^{\ell_{w,q}^3}$. Compute the evaluation keys

$$\zeta = \left[f^{-1} P_{w,q}(D_{w,q}(f) \otimes D_{w,q}(f)) + \mathbf{e} + h\mathbf{s} \right] \in R^{\ell_{w,q}^3}.$$

- **Encrypt.** Encrypt message μ as in Basic scheme.
- **Decrypt.** Decrypt message as in Basic scheme.
- **KeySwitch.** Compute $\langle D_{w,q}(c), \zeta \rangle$ where c is a ciphertext.
- **Addition.** Addition of two ciphertexts c_1 and c_2 is $c = c_1 + c_2$.
- **Multiplication.** Multiplication of two ciphertexts is computed as

$$c = \left[\left[\frac{t}{q} P_{w,q}(c_1) \otimes P_{w,q}(c_2) \right] \right]$$

and later apply KeySwitch to c and output.

YASHE'. The primitive operations are as follows:

- **ParamsGen.** The parameters are selected the same way as in Basic scheme.
- **KeyGen.** The public and secret key pairs are selected using the Basic.KeyGen routine, i.e. $h, f \leftarrow \text{Basic.KeyGen}$. Sample $\mathbf{e}, \mathbf{s} \in \chi_{\text{err}}^{\ell_{w,q}}$. Compute the evaluation keys

$$\zeta = \left[f^{-1} P_{w,q}(D_{w,q}(f) \otimes D_{w,q}(f)) + \mathbf{e} + h\mathbf{s} \right] \in R^{\ell_{w,q}}.$$

- **Encrypt.** Encrypt message μ as in Basic scheme.
- **Decrypt.** Decrypt message as in Basic scheme.

- **KeySwitch.** Compute $\langle D_{w,q}(c), \zeta \rangle$ where c is a ciphertext.
- **Addition.** Addition of two ciphertexts c_1 and c_2 is $c = c_1 + c_2$.
- **Multiplication.** Multiplication of two ciphertexts is computed as

$$c = \left[\left[\begin{array}{c} t \\ - \\ q \end{array} c_1 c_2 \right] \right]$$

and later apply KeySwitch to c and output.

2.4 Impact of the Subfield Lattice Attack on LTV and YASHE'

On the downside many of these assumptions are still open to debate from a security point of view. A very recent work by Albrecht, Bai and Ducas [33] painfully demonstrated this fact. The authors exploit the presence of a subfield to solve the NTRU problem for large moduli q and show that when the NTRU parameters are chosen poorly it becomes possible to norm-down the NTRU public key h to a subfield yielding an easier lattice problem. Consequently, and any sufficiently good solution may be lifted to a short vector in the full NTRU lattice. The attack works when the secret key f is chosen from a narrow distribution, e.g. $\|f\| \leq \sqrt{q}$ and when the polynomial modulus is chosen such that a subfield of reasonable size exists. In this setting, Albrecht et al. show that the DSPR problem is not as hard as believed thereby invalidating the basic assumption in the LTV [31, 16] and YASHE' [1] schemes. Thus, the subfield lattice attack significantly diminishes the asymptotic security of both schemes.

Both LTV and YASHE' rely on the secret key from being sampled from a narrow distribution to support even a single homomorphic multiplication. This eliminates the possibility of sampling the key from a wider distribution. The adverse effect of the attack could be mitigated, by maximizing size of the subfield as recommended in [33]. Even then, the lattice dimension and parameters need to be increased to restore the projected security level of LTV and YASHE'. Another important side effect of the Subfield Lattice Attack is that it makes the selection of parameters that support *batching* rather difficult.

2.5 GSW Scheme

Gentry, Sahai and Waters [25] proposed this new scheme based on the hardness of the approximate eigenvector problem. The construction consists of simple matrix addition and multiplication operations to perform homomorphic addition and homomorphic multiplication. The advantage of the scheme is that it eliminates the need of relinearization, storage of evaluation keys and even modulus switching. As many other homomorphic schemes the security is based on the LWE problem.

We explain the scheme in four primitive functions in below. Before the explanations we should note some of the preliminaries that is used in the primitive functions. A vector $\vec{a} = (a_0, \dots, a_{k-1})$ is split into its bits using the function $\text{BitDecomp}(\vec{a}) = (a_{0,0}, \dots, a_{0,\ell-1}, \dots, a_{k-1,0}, \dots, a_{k-1,\ell-1})$. We are able to

reconstruct the elements from the bit representations using the inverse of the `BitDecomp` as $\text{BitDecomp}^{-1}(\vec{a}) = (\sum 2^j a_{0,j}, \dots, \sum 2^j a_{k-1,j})$. The most important function that the scheme uses is *flattening*. It keeps the ciphertexts bounded so that the noise increase after multiplicative operations are limited. We simply evaluate it as $\text{Flatten}(\vec{a}) = \text{BitDecomp}(\text{BitDecomp}^{-1}(\vec{a}))$. The last function is $\text{Powersof2}(\vec{a}) = (a_0, 2a_0, \dots, 2^{\ell-1}a_0, \dots, a_k, 2a_k, \dots, 2^{\ell-1}a_k)$ which multiplies the vector elements with powers of two. Using these functions, the encryption scheme is defined with the following primitives:

- **Setup.** We select λ as the security parameter and L as multiplicative depth. Then compute lattice dimension $n = n(\lambda, L)$, error distribution $\chi = \chi(\lambda, L)$, parameter $m = m(\lambda, L) = \mathcal{O}(n \log q)$. Also, set $\ell = \lceil \log q \rceil + 1$ and $N = (n + 1)\ell$.
- **KeyGen.** We sample $\vec{t} \leftarrow \mathbb{Z}_q^n$ and compute $\vec{s} \leftarrow (1, -t_1, -t_2, \dots, -t_n)$. Then, set the secret key $\vec{v} = \text{Powersof2}(\vec{s})$. The public key matrix is computed by first generating uniform matrix $B \leftarrow \mathbb{Z}_q^{m \times n}$ and error vector $\vec{e} \leftarrow \chi^m$. We set $(n + 1)$ column matrix A having $\vec{b} = B \cdot \vec{t} + \vec{e}$ as the first column and \vec{b} in rest of the columns as the public key.
- **Encrypt.** A message μ is encrypted by simply computing $C = \text{Flatten}(\mu I_N + \text{BitDecomp}(R \cdot A)) \in \mathbb{Z}_q^{N \times N}$. In the equation R is selected as a uniform matrix $R \in \{0, 1\}^{N \times m}$.
- **Decrypt.** Select a row of the matrix, i.e. C_i as the i -th row of matrix C . Compute $x_i \leftarrow \langle C_i, \vec{v} \rangle$ and the message as $\mu' = \lfloor x_i / v_i \rfloor$.

The beauty of the GSW scheme is that straightforward addition and multiplication of ciphertext matrices suffice to compute homomorphic additions and multiplications. First, let's observe that the construction holds the property $C \cdot \vec{v} = \mu \cdot \vec{v} + \vec{e}$, since secret key \vec{v} is the approximate eigenvector related to the ciphertext. Therefore, adding ciphertext matrices has the effect of adding the corresponding eigenvalues (messages): $(C_1 + C_2) \cdot \mathbf{v} = (\mu_1 + \mu_2) \cdot \mathbf{v} + (\mathbf{e}_1 + \mathbf{e}_2)$. Similarly the matrix product of the ciphertexts (in any order) multiplies the eigenvalues: $C_1 \cdot C_2 \cdot \mathbf{v} = C_1(\mu_2 \cdot \mathbf{v} + \mathbf{e}_2) = \mu_2(\mu_1 \cdot \mathbf{v} + \mathbf{e}_1) + C_1 \cdot \mathbf{e}_2 = \mu_1 \mu_2 \cdot \mathbf{v} + \mathbf{e}$.

3 Our proposal: F-NTRU Scheme

Here we propose a new scheme called F-NTRU that shares the goals of the GSW construction [25], i.e. no evaluation keys, no expensive relinearization operations, no modulus switching and simple homomorphic additions and multiplications. To this end we adopt the flattening approach of [25] and apply it to the NTRU variant, i.e. NTRU', by Stehlé and Steinfeld [40].

Preliminaries. We work in $R_q = \mathbb{Z}_q[x]/\langle x^n + 1 \rangle$. We adapt two functions from [25] to work in the polynomial setting as follows:

- **Bit-Decomposition:** The `BitDecomp` function takes a ciphertext polynomial $c(x)$ and splits them into binary polynomials and forms a vector as:

$$\vec{c}(x) = \text{BitDecomp}(c(x)) = [c_{\ell-1}(x) \quad c_{\ell-2}(x) \quad \dots \quad c_2(x) \quad c_1(x) \quad c_0(x)].$$

Here a polynomial $c_i(x)$ with index i represents the binary polynomial that is formed using the i^{th} bit index of the coefficients of $c(x)$. One may easily reconstruct the ciphertext $c(x)$ by simply computing:

$$c(x) = \sum_{i=0}^{\ell-1} 2^i \cdot c_i(x) \in R_q.$$

Note that when we are computing BitDecomp^{-1} , the elements in the vector do not necessarily have to be binary polynomials. It is possible that polynomials can contain coefficients that are not bits.

- **Flatten:** When we are performing arithmetic operations, i.e. addition and multiplication, in our scheme, the elements of a flattened ciphertext vector lose their binary form. These extra bits in the coefficients of the polynomials cause additional noise in subsequent arithmetic operations. To prevent this, we use **Flatten**. **Flatten** restructures all the elements of the ciphertext vector into binary polynomials. We evaluate the **Flatten** operation as follows:

$$\text{Flatten}(\vec{c}(x)) = \text{BitDecomp}(\text{BitDecomp}^{-1}(\vec{c}(x))).$$

Here in the equation BitDecomp^{-1} converts the ciphertext vector into a full ciphertext polynomial in R_q . Later, using **BitDecomp** we convert the ciphertext into a binary polynomial vector again. Basically, this method carries over the extra bits of an element in the vector from least significant to most significant and performs a modular reduction using q to prevent overflow in the overall scheme.

The F-NTRU Scheme. The primitive operations of F-NTRU are defined as follows:

- **KeyGen:** We use the key generation method of NTRU' to select our parameters. For a security parameter λ , we choose our message modulus as 2, modulus $q = q(\lambda)$, polynomial degree $n = n(\lambda)$ where n is power of 2. Also we set Gaussian distributions $\chi_{\text{err}} = \chi_{\text{err}}(\lambda)$ and $\chi_{\text{key}} = \chi_{\text{key}}(\lambda)$. Sample $g, f' \in \chi_{\text{key}}$ and set public key $h = 2gf^{-1}$ and secret key $f = 2f' + 1$.
- **Encrypt:** In order to encrypt a message μ , we create a vector with length $\ell = \log q$. Later, we fill the elements with encryptions of zeros using the NTRU' scheme:

$$\vec{c} = \{\text{Enc}_{\ell-1}(0), \text{Enc}_{\ell-2}(0), \dots, \text{Enc}_0(0)\} = \{c_{\ell-1}, c_{\ell-2}, \dots, c_0\},$$

where $\text{Enc}_i(0) = hs_i + 2e_i + 0$. We call this the *ciphertext vector*. We take the transpose of the vector to list the elements in row and apply **BitDecomp** to the ciphertexts to turn the vector into a $\ell \times \ell$ matrix:

$$c = \text{BitDecomp}(\vec{c}^T)$$

Then, we use the matrix to encrypt a message μ by evaluating:

$$C = \text{Flatten}(I_\ell \cdot \mu + c).$$

Here, in the equation I_ℓ is identity matrix with size ℓ .

- **Decrypt:** To decrypt the message we take the first row of the matrix and apply Inverse-Bit-Decomposition to form a NTRU' ciphertext:

$$\text{BitDecomp}^{-1}\{c_{(0,\ell-1)}, c_{(0,\ell-2)}, \dots, c_{(0,2)}, c_{(0,1)}, c_{(0,0)}\} = c_0.$$

Once the NTRU' ciphertext is constructed, we are able to decrypt the message using the secret key f for the NTRU' scheme as $\lfloor c_0 f \rfloor \bmod 2 = \mu$.

- **Eval.** The homomorphic XOR and AND operations are simply computed as matrix addition and multiplication operations, respectively followed by a Flatten operation, i.e.

$$C' = \text{Flatten}(C + \tilde{C}) \quad , \quad C' = \text{Flatten}(C \cdot \tilde{C}).$$

3.1 Correctness of Homomorphic Circuit Evaluation

The correctness of encryption/decryption trivially follows from the correctness of the NTRU' scheme. In this section we briefly demonstrate how the NTRU' ciphertext and associated F-NTRU ciphertext matrix forms are preserved thus allowing homomorphic evaluation. For clarity we use ciphertexts of sizes $\ell = 4$.

First note that $\text{BitDecomp}^{-1}(C)$ is actually a vector that contains the encryptions of message bits scaled by powers of 2:

$$\text{BitDecomp}^{-1}(\text{Flatten}(I_\ell \cdot \mu + c)) = [c_{\ell-1} + 2^{\ell-1} \cdot \mu, \dots, c_1 + 2^1 \cdot \mu, c_0 + 2^0 \cdot \mu]^\top.$$

This is the form we need to preserve throughout homomorphic evaluations for correctness.

Homomorphic XOR. A homomorphic XOR operation between two ciphertext matrices is computed as:

$$C' = \text{Flatten}(C + \tilde{C}) = \begin{bmatrix} c_{(3,3)} + \tilde{c}_{(3,3)} + \mu + \tilde{\mu} & c_{(3,2)} + \tilde{c}_{(3,2)} & c_{(3,1)} + \tilde{c}_{(3,1)} & c_{(3,0)} + \tilde{c}_{(3,0)} \\ c_{(2,3)} + \tilde{c}_{(2,3)} & c_{(2,2)} + \tilde{c}_{(2,2)} + \mu + \tilde{\mu} & c_{(2,1)} + \tilde{c}_{(2,1)} & c_{(2,0)} + \tilde{c}_{(2,0)} \\ c_{(1,3)} + \tilde{c}_{(1,3)} & c_{(1,2)} + \tilde{c}_{(1,2)} & c_{(1,1)} + \tilde{c}_{(1,1)} + \mu + \tilde{\mu} & c_{(1,0)} + \tilde{c}_{(1,0)} \\ c_{(0,3)} + \tilde{c}_{(0,3)} & c_{(0,2)} + \tilde{c}_{(0,2)} & c_{(0,1)} + \tilde{c}_{(0,1)} & c_{(0,0)} + \tilde{c}_{(0,0)} + \mu + \tilde{\mu} \end{bmatrix}.$$

When we apply $\text{BitDecomp}^{(-1)}$ to the rows of the addition matrix we obtain:

$$[(c_3 + \tilde{c}_3) + 8(\mu + \tilde{\mu}), (c_2 + \tilde{c}_2) + 4(\mu + \tilde{\mu}), (c_1 + \tilde{c}_1) + 2(\mu + \tilde{\mu}), (c_0 + \tilde{c}_0) + 1(\mu + \tilde{\mu})]$$

The ciphertext vector is still valid for the following two reasons:

1. The first part of the addition $c_i + \tilde{c}_i$ still holds an encryption of zero. The only difference is that it is noisier compared to a fresh encryption.

2. The second term in each entry is the message scaled by powers of two, i.e. $2^i \cdot (\mu + \tilde{\mu})$ as in a fresh ciphertext.

Homomorphic AND. A homomorphic AND operation between two ciphertext matrices is computed as:

$$C' = \text{Flatten}(C \cdot \tilde{C}) = \left[\vec{c}'_3, \vec{c}'_2, \vec{c}'_1, \vec{c}'_0 \right]^\top =$$

$$\begin{bmatrix} c_{(3,3)} + \mu & c_{(3,2)} & c_{(3,1)} & c_{(3,0)} \\ c_{(2,3)} & c_{(2,2)} + \mu & c_{(2,1)} & c_{(2,0)} \\ c_{(1,3)} & c_{(1,2)} & c_{(1,1)} + \mu & c_{(1,0)} \\ c_{(0,3)} & c_{(0,2)} & c_{(0,1)} & c_{(0,0)} + \mu \end{bmatrix} \cdot \begin{bmatrix} \tilde{c}_{(3,3)} + \tilde{\mu} & \tilde{c}_{(3,2)} & \tilde{c}_{(3,1)} & \tilde{c}_{(3,0)} \\ \tilde{c}_{(2,3)} & \tilde{c}_{(2,2)} + \tilde{\mu} & \tilde{c}_{(2,1)} & \tilde{c}_{(2,0)} \\ \tilde{c}_{(1,3)} & \tilde{c}_{(1,2)} & \tilde{c}_{(1,1)} + \tilde{\mu} & \tilde{c}_{(1,0)} \\ \tilde{c}_{(0,3)} & \tilde{c}_{(0,2)} & \tilde{c}_{(0,1)} & \tilde{c}_{(0,0)} + \tilde{\mu} \end{bmatrix}$$

We summarize the derivation of the rows of the product matrix as follows:

Row 0: In Table 1 we show columns of the vector $\vec{c}'_0 = \text{BitDecomp}(c'_0(x))$. In the last row of the table we evaluate $\text{BitDecomp}^{-1}(\vec{c}'_0)$. The last column contains the respective powers of 2 used in BitDecomp^{-1} for easy reference. The

$\vec{c}'_{(0,3)}$	$c_{(0,3)} \cdot (\tilde{c}_{(3,3)} + \tilde{\mu}) + c_{(0,2)} \cdot \tilde{c}_{(2,3)} + c_{(0,1)} \cdot \tilde{c}_{(1,3)} + (c_{(0,0)} + \mu) \cdot \tilde{c}_{(0,3)}$	8
$\vec{c}'_{(0,2)}$	$c_{(0,3)} \cdot \tilde{c}_{(3,2)} + c_{(0,2)} \cdot (\tilde{c}_{(2,2)} + \tilde{\mu}) + c_{(0,1)} \cdot \tilde{c}_{(1,2)} + (c_{(0,0)} + \mu) \cdot \tilde{c}_{(0,2)}$	4
$\vec{c}'_{(0,1)}$	$c_{(0,3)} \cdot \tilde{c}_{(3,1)} + c_{(0,2)} \cdot \tilde{c}_{(2,1)} + c_{(0,1)} \cdot (\tilde{c}_{(1,1)} + \tilde{\mu}) + (c_{(0,0)} + \mu) \cdot \tilde{c}_{(0,1)}$	2
$\vec{c}'_{(0,0)}$	$c_{(0,3)} \cdot \tilde{c}_{(3,0)} + c_{(0,2)} \cdot \tilde{c}_{(2,0)} + c_{(0,1)} \cdot \tilde{c}_{(1,0)} + (c_{(0,0)} + \mu) \cdot (\tilde{c}_{(0,0)} + \tilde{\mu})$	1
c'_0	$c_{(0,3)} \cdot \tilde{c}_3 + c_{(0,2)} \cdot \tilde{c}_2 + c_{(0,1)} \cdot \tilde{c}_1 + c_{(0,0)} \cdot \tilde{c}_0 + c_0 \cdot \tilde{\mu} + \tilde{c}_0 \cdot \mu + \mu \cdot \tilde{\mu}$	

Table 1. Derivation of Row 0 of product ciphertext

final ciphertext in Table 1 has the proper ciphertext form: The first part of the ciphertext $c_{(0,3)} \cdot \tilde{c}_3 + c_{(0,2)} \cdot \tilde{c}_2 + c_{(0,1)} \cdot \tilde{c}_1 + c_{(0,0)} \cdot \tilde{c}_0 + c_0 \cdot \tilde{\mu} + \tilde{c}_0 \cdot \mu$ is still an encryption of zero since ciphertexts c_i and \tilde{c}_i are encryptions of zero and their multiplication with a binary polynomial will result in zero encryptions as long as noise is contained.

Row 1: We apply the same arithmetic for the row number 1 and form a similar construction in Table 2. We achieve a similar derivation for second row in Table 2. The only difference is that now the message is scaled by 2.

We can now generalize the arithmetic for each row i as follows:

$$c'_i = \underbrace{\sum_{j=0}^{\ell-1} c_{(i,j)} \cdot \tilde{c}_j + c_i \cdot \tilde{\mu} + \tilde{c}_i \cdot \mu}_{\tilde{c}_i} + 2^i(\mu \cdot \tilde{\mu}). \quad (1)$$

As shown in Equation 1, after a matrix multiplication the ciphertext vector elements contains a (noisier) encryption of zero along with the message scaled by powers of 2, i.e. $\vec{c}' = \{\tilde{c}_{\ell-1} + 2^{\ell-1}\tilde{\mu}, \dots, \tilde{c}_1 + 2^1\tilde{\mu}, \tilde{c}_0 + 2^0\tilde{\mu}\}$ in which $\tilde{\mu} = \mu \cdot \tilde{\mu}$. Hence correctness holds throughout circuit evaluation as long as noise is contained.

$\vec{c}'_{(1,3)}$	$c_{(1,3)} \cdot (\tilde{c}_{(3,3)} + \tilde{\mu}) + c_{(1,2)} \cdot \tilde{c}_{(2,3)}$	$+(c_{(1,1)} + \mu) \cdot \tilde{c}_{(1,3)}$	$+c_{(1,0)} \cdot \tilde{c}_{(0,3)}$	8	
$\vec{c}'_{(1,2)}$	$c_{(1,3)} \cdot \tilde{c}_{(3,2)}$	$+c_{(1,2)} \cdot (\tilde{c}_{(2,2)} + \tilde{\mu})$	$+(c_{(1,1)} + \mu) \cdot \tilde{c}_{(1,2)}$	$+c_{(1,0)} \cdot \tilde{c}_{(0,2)}$	4
$\vec{c}'_{(1,1)}$	$c_{(1,3)} \cdot \tilde{c}_{(3,1)}$	$+c_{(1,2)} \cdot \tilde{c}_{(2,1)}$	$+(c_{(1,1)} + \mu) \cdot (\tilde{c}_{(1,1)} + \tilde{\mu})$	$+c_{(1,0)} \cdot \tilde{c}_{(0,1)}$	2
$\vec{c}'_{(1,0)}$	$c_{(1,3)} \cdot \tilde{c}_{(3,0)}$	$+c_{(1,2)} \cdot \tilde{c}_{(2,0)}$	$+(c_{(1,1)} + \mu) \cdot \tilde{c}_{(1,0)}$	$+c_{(1,0)} \cdot (\tilde{c}_{(0,0)} + \tilde{\mu})$	1
c_1	$\underbrace{c_{(1,3)} \cdot \tilde{c}_3 + c_{(1,2)} \cdot \tilde{c}_2 + c_{(1,1)} \cdot \tilde{c}_1 + c_{(1,0)} \cdot \tilde{c}_0 + c_1 \cdot \tilde{\mu} + \tilde{c}_1 \cdot \mu + 2(\mu \cdot \tilde{\mu})}_{\tilde{c}_1}$				

Table 2. Derivation of Row 1 of product ciphertext

4 Optimizations

In our scheme we are able to perform certain optimizations on the arithmetic operations to reduce computational complexity and the memory footprint. The two main optimizations we perform are:

Using High Radix Representations. The flattened ciphertext matrix takes up a large space, $\ell^2 = O(\log(q)^2)$. To mitigate we may use a higher radix representation, i.e. a larger radix 2^ω and group ω bits for each element of the ciphertext matrix. When we apply BitDecomp^{-1} , we use the powers of the chosen radix $(2^\omega)^i$ instead of powers of 2^i to reconstruct the ciphertext. With this approach we drastically reduce the matrix size: $\ell' = \ell/\omega$. Note that the number of zero encryptions of NTRU' should be equal to ℓ' as well. Using the high radix encoding the ciphertext size is reduced by ω times. In addition, the ciphertext matrix size is reduced by ω^2 which decreases the complexity of matrix multiplication significantly.

Matrix Multiplication. A straight schoolbook matrix multiplication takes $O(\ell^3)$ time and it may be reduced to $O(\ell^{2.374})$ time using Coppersmith-Winograd algorithm. However, to evaluate deep circuits we will need a large modulus and even with the Coppersmith-Winograd algorithm multiplication will be slow. Therefore, we change matrix multiplication into a matrix-vector multiplication as follows:

$$\text{BitDecomp}^{-1} \left(\text{BitDecomp}(\vec{c}) \cdot \text{BitDecomp}(\vec{\tilde{c}}) \right) = \text{BitDecomp}(\vec{c}) \cdot \vec{\tilde{c}}.$$

A schoolbook matrix-vector multiplication has $O(\ell^2)$ runtime which is faster than a matrix multiplication. The only downside of the algorithm is that a binary by a high radix polynomial multiplication should be implemented instead of a binary by binary polynomial multiplication. Although binary by high radix polynomial multiplication is slower, the time gap is closed in the higher level while computing the matrix-vector product for large ℓ values.

5 Security Analysis

Our scheme uses the NTRU' encryption scheme adopted from [40]. In F-NTRU, an attacker has access only to a ciphertext vector which holds NTRU' ciphertexts

build with *fresh* encryptions of zero. Therefore as long as NTRU' ciphertexts are secure, our scheme is also secure. According to the Stehlé and Steinfeld [40], their scheme is IND-CPA secure under the hardness of R-LWE assumption as long as the error has a *wide distribution* $\mathbb{D}_{\mathbb{Z}^n, \sigma}$, i.e. $\sigma > \sqrt{q} \cdot \text{poly}(n)$. However, in the DHS and YASHE' schemes such a high noise instantiations are not possible due to the noise growth in homomorphic multiplications. Therefore, both schemes use narrow error distributions and additionally assume hardness of the Decisional Small Polynomial Ratio (DSPR) Problem along with the R-LWE assumption.

Fortunately, our scheme has a better control on noise management. The (size of the) noise increases linearly with the multiplicative depth. Therefore, we are able to use a *wide error distribution* and achieve IND-CPA security as in the original scheme by Stehlé and Steinfeld in [40]. Thus the standard deviation σ_{key} of the discrete Gaussian distribution $\mathbb{D}_{\mathbb{Z}^n, \sigma_{\text{key}}}$ needs to be set as:

$$\sigma_{\text{key}} > 2n\sqrt{\log(8nq)} \cdot q^{1/2+\epsilon}$$

for $\epsilon > 0$, i.e. $\epsilon = 2^{-80}$. In this setting we are able to generate a public key $h = g/f$, i.e. $g \in \chi_{\text{key}}$ and $f' \in \chi_{\text{key}}$ ($f = 2f' + 1$) in which χ_{key} is truncated $\mathbb{D}_{\mathbb{Z}^n, \sigma_{\text{key}}}$, that is indistinguishable from a uniformly random distribution.

For R-LWE security we follow the settings in [31]. The noise parameters s and e are sampled from the distribution χ_{err} , a truncated Gaussian distribution $\mathbb{D}_{\mathbb{Z}^n, \sigma_{\text{err}}}$. The standard deviation σ_{err} of the error distribution has the following bound $\sigma_{\text{err}} > \sqrt{\log(n)}$. With this bound, we are able to add noise to our public key and keep it computationally indistinguishable from random uniform distribution.

As above mentioned, the distributions χ_{key} and χ_{err} are truncated Gaussian distributions. They are also defined as B -bounded¹ distributions which means the samples are selected from $[-B, B]$. The bound value B is selected as a function of the advantage ϵ in a distinguishing attack, e.g. $\epsilon = 2^{-80}$. The function, e.g. see [19], gives for a sample $x \in \chi$ the probability of x being larger than $k \cdot \sigma$ for a factor k is:

$$\text{Prob}_{x \leftarrow \mathbb{D}_{\mathbb{Z}^n, \sigma}}[|x| > k \cdot \sigma] = \text{erf}(k/\sqrt{2}).$$

Using the equation above, we select the factor k as $k(\epsilon) = \min\{k \mid \text{erf}(k/\sqrt{2}) < \epsilon\}$, and select the B -bound as $B = k(\epsilon) \cdot \sigma$.

Parameter Selection. The existing attacks rely on the lattice reduction algorithms and the best attack known in practice is BKZ 2.0 [9] by Chen and Nguyen. The quality of the security is measured by Hermite factor δ . In [30] Linder and Peikert derive a linear security estimate using the Hermite factor δ :

$$\log(T_{\text{BKZ}}(\delta)) = 1.8/\log(\delta) - 110.$$

¹ We get rid of the subscripts *key* and *err* for simplicity. We use B_{key} and B_{err} if they need to be specified.

In order to have at least 80-bit security, i.e. $T_{\text{BKZ}} = 2^{80}$, we need $\delta \leq 1.0066$. In [30] the Hermite factor δ is given as:

$$\delta(n, q, \sigma, \epsilon) = 2^{\log^2(q/\sigma \cdot \sqrt{2 \log(1/\epsilon)}) / (4n \log q)}.$$

Using these equations and by fixing the Hermite factor, i.e. $\delta = 1.0066$, we achieve the same security level for various values of n by changing q and σ .

6 Noise Analysis

In this section we analyze the noise performance of our scheme with homomorphic evaluations. In case of a homomorphic addition, the noise increases only a small percent compared to homomorphic multiplication. In our analysis we want to determine the depth of the circuit we can evaluate and still decrypt correctly with growing noise. This analysis includes average case and worst case scenarios for homomorphic multiplication that estimates the number of possible multiplicative levels. Since additions contribute minimally to noise growth we focus only on homomorphic multiplications.

For an element $a \in \mathbb{R}$ we define the Euclidean norm $\|a\| = \sqrt{\sum a_i^2}$ and the infinity norm $\|a\|_\infty = \max |a_i|$ for all possible values of i . In multiplication, we can bound the noise growth with the aid of the following Lemma.

Lemma 1 ([32, 31]). *In a ring $R = \mathbb{Z}[x]/\langle x^n + 1 \rangle$, for any two polynomials $a, b \in R$ we have the following norms $\|ab\| \leq \sqrt{n} \|a\| \cdot \|b\|$ and $\|ab\|_\infty \leq n \|a\|_\infty \cdot \|b\|_\infty$.*

Homomorphic Multiplication. We assume the evaluation circuit is arranged into a tree with levels of parallel multiplication gates that accept input the output ciphertexts from the previous level. Lets denote a ciphertext matrix that are at a certain multiplicative level i as $C^{(i)}$. Then, a ciphertext matrix at a multiplicative level $C^{(i)} = C^{(i-1)} \cdot \tilde{C}^{(i-1)}$. Lets recall that these ciphertext matrices are actually NTRU' ciphertext vectors with BitDecomp applied. We also denote the ciphertexts in the vector for multiplicative level i and row index j as $c_j^{(i)}$.

In the first multiplicative level, i.e. $i = 0$, ciphertext vectors are fresh encryptions with security paramters explained in Section 5. Basically, we have samples $g, f' \in \chi_{\text{key}}$ and $s, e \in \chi_{\text{err}}$ and ciphertexts $c_j^{(0)} = hs_j + 2e_j + 2^j \mu$ where μ is the message, $f = 2f' + 1$ and $h = 2gf^{-1}$ is public key. Lets recall that for ciphertext matrix multiplications we have NTRU' ciphertexts that hold the form given in Equation 1. The equation can be rewritten with level index to show the result of a multiplicative level as:

$$c_j^{(i)} = \sum_{k=0}^{\ell-1} c_{(j,k)} \cdot \tilde{c}_k^{(i-1)} + c_j^{(i-1)} \cdot \tilde{\mu} + \tilde{c}_j^{(i-1)} \cdot \mu + 2^j (\mu \cdot \tilde{\mu}).$$

We can simplify the equation, for noise evaluation, by replacing radix size polynomials $c_{(j,k)}$ with y_τ , choosing ciphertext vector index $j = 0$, and substituting

$y_{(i)}$ in place of all ciphertexts since they have the same noise level for the same multiplicative level i :

$$y_{(i)} = y_{(i-1)}y_\tau\ell + y_{(i-1)}\tilde{\mu} + y_{(i-1)}\mu + \mu\tilde{\mu}.$$

To be able to decrypt $y_{(i)}$ correctly, we need $\|y_{(i)}f\|_\infty < q/2$. Thus, we need

$$\|y_{(i)}f\|_\infty \leq \|y_{(i-1)}fy_\tau\ell\|_\infty + \|y_{(i-1)}f\tilde{\mu}\|_\infty + \|y_{(i-1)}f\mu\|_\infty + \|\mu\tilde{\mu}f\|_\infty.$$

Later, we expand $\|y_{(i-1)}f\|_\infty$ in terms of $\|y_{(i-2)}f\|_\infty$ and continue the process recursively. At the lowest level we have:

$$\|y_{(0)}f\|_\infty = \|c^{(0)}f\|_\infty \leq \|2gf^{(-1)}sf\|_\infty + \|2ef\|_\infty.$$

Since $\|g\|_\infty = \|f'\|_\infty = B_{\text{key}}$, $\|s\|_\infty = \|e\|_\infty = B_{\text{err}}$, the worst-case noise is equal to:

$$\begin{aligned} \|y_{(0)}f\|_\infty &\leq 2nB_{\text{key}}B_{\text{err}} + 2nB_{\text{err}}(2B_{\text{key}} + 1) \\ &\leq 2nB_{\text{err}}(3B_{\text{key}} + 1). \end{aligned}$$

For an arbitrary level i , the noise can be evaluated recursively by setting $B_i = \|y_{(i)}f\|_\infty$, $\|y_\tau\|_\infty = 2^\omega - 1$ and $\|\mu\|_\infty = 1$ as:

$$\begin{aligned} \underbrace{\|y_{(i)}f\|_\infty}_{B_i} &\leq \underbrace{\|y_{(i-1)}fy_\tau\ell\|_\infty}_{\ell n(2^\omega - 1)B_{(i-1)}} + \underbrace{\|y_{(i-1)}f\tilde{\mu}\|_\infty}_{nB_{(i-1)}} + \underbrace{\|y_{(i-1)}f\mu\|_\infty}_{nB_{(i-1)}} \\ B_i &\leq \ell n(2^\omega - 1)B_{(i-1)} + 2nB_{(i-1)} \\ B_i &\leq (\ell n(2^\omega - 1) + 2n)B_{(i-1)} \end{aligned}$$

In the noise analysis we exclude $\|\mu\tilde{\mu}f\|_\infty$ because it only occurs at the last level when we want to decrypt. The noise actually stems from the noisy zero encryptions, i.e. $\bar{c}_k = \sum_{j=0}^{\ell-1} c_{(j,k)} \cdot \tilde{c}_k^{(i-1)} + c_j^{(i-1)} \cdot \tilde{\mu} + \tilde{c}_j^{(i-1)} \cdot \mu$. In the final multiplicative level $\|\mu\tilde{\mu}f\|_\infty$ adds additional noise $\|\mu\tilde{\mu}2f'\|_\infty = 2n^2B_{\text{key}}$.

Average Case Noise Analysis. In the average case the noise accumulation will be much lower than what the bound given above predicts due to the Gaussian distribution. Using the worst case noise analysis, we can simply rewrite the average case noise analysis for the first multiplicative level as:

$$\begin{aligned} \|y_{(0)}f\|_\infty &\leq 2\sqrt{n}B_{\text{key}}B_{\text{err}} + 2\sqrt{n}B_{\text{err}}(2B_{\text{key}} + 1) \\ &\leq 2\sqrt{n}B_{\text{err}}(3B_{\text{key}} + 1), \end{aligned}$$

and for the remaining levels we obtain

$$\begin{aligned} B_i &\leq \sqrt{\ell}\sqrt{n}(2^\omega - 1)B_{(i-1)} + 2\sqrt{n}B_{(i-1)} \\ B_i &\leq (\sqrt{\ell n}(2^\omega - 1) + 2\sqrt{n})B_{(i-1)}. \end{aligned}$$

For the last stage we have additional noise from $\|\mu\tilde{\mu}f\|_\infty$ as $2nB_{\text{key}}$.

7 The F-NTRU' Variant

The F-NTRU construction manages to eliminate evaluation keys and relinearizations and the DSPR assumption. However, the ciphertext grows due to the matrix based construction compared to a standard NTRU based SHE constructions. The growth in the ciphertext size not only increases the memory footprint, but also increases the latency of homomorphic evaluations. With this motive we introduce a variant we call F-NTRU' that uses less noise, by sampling the secret key f' and g from a slightly narrower distribution, i.e. $\mathcal{O}(\text{poly}(n) \cdot q^{1/3})$ instead of $\mathcal{O}(\text{poly}(n) \cdot q^{1/2})$. From Section 5 remember that for F-NTRU we require $\sigma_{\text{key}} > 2n\sqrt{\log(8nq)} \cdot q^{1/2+\epsilon}$. For a discussion on the effect of noise distribution on the security of NTRU based schemes, and possible attacks enabled by the use of very narrow key distributions see [33].

Under this modification we can select smaller parameters. The security and noise analysis presented above still applies with the exception of the re-introduced DSPR assumption. Therefore we can still use the noise bounds and parameter derivations by simply setting the noise to the appropriate values, i.e. for f' and g we use $\sigma_{\text{key}} > 2n\sqrt{\log(8nq)} \cdot q^{1/3+\epsilon}$ and for e and s we use $\sigma_{\text{err}} = \sqrt{\log(n)}$.

8 Complexity

In Table 3 we summarize the asymptotic complexities of F-NTRU and YASHE schemes. We use $\ell = \log q/\omega$ for radix size ω . Note that our scheme does not require evaluation keys. However F-NTRU ciphertexts consist of ℓ polynomials. On the other hand YASHE requires ℓ^3 polynomials for evaluation keys. For homomorphic AND evaluation our scheme requires ℓ^2 polynomial multiplications whereas YASHE needs ℓ^2 polynomial multiplications followed by a costly key switching operation computed via ℓ^3 polynomial multiplications.

	F-NTRU/F-NTRU'	YASHE
Eval. Key Size	-	$\mathcal{O}(\ell^3 n \log q)$
Ciphertext Size	$\mathcal{O}(\ell n \log q)$	$\mathcal{O}(n \log q)$
AND Eval.	$\mathcal{O}(\ell^2)$	$\mathcal{O}(\ell^2)$
Key-Switching	-	$\mathcal{O}(\ell^3)$

Table 3. Comparison of F-NTRU and YASHE: homomorphic AND evaluation and Key Switching complexities are in terms of polynomial multiplications with $\ell = \log q/\omega$.

9 Parameter Selection

Using the earlier noise analysis we evaluate maximum modulus size $\log q_{\text{max}}$ for various values of polynomial degree n , i.e. $\log(n) \in [10, 14]$. Using the input pair $\log(n)$ and $\log q_{\text{max}}$ we evaluate maximum multiplicative depth L in Table 4.

Note that for similar sizes of q , the dimension n doubles for F-NTRU' to keep the Hermite at $\delta(n) = 1.0066$, however we can accommodate roughly % 50 more evaluation levels in F-NTRU' .

$\log(n)$	ω	F-NTRU			F-NTRU'		
		$\lfloor \log(q_{\max}) \rfloor$	L_{worst}	L_{avg}	$\lfloor \log(q_{\max}) \rfloor$	L_{worst}	L_{avg}
10	16		1	2		1	1
	8	197	2	3	117	1	2
	1		3	7		2	5
11	16		3	5		2	4
	8	359	5	8	210	3	6
	1		7	14		5	10
12	16		7	11		5	8
	8	677	10	16	389	7	12
	1		13	28		10	21
13	16		15	22		11	16
	8	1305	19	31	744	14	23
	1		26	52		19	40
14	16		29	43		22	32
	8	2555	36	59	1448	28	45
	1		48	97		37	74

Table 4. Supported multiplicative depth for various radix selections, i.e. $\omega = \{1, 8, 16\}$. The distinguishing advantage is set to 2^{-80} .

For comparison, we also tabulated parameter choices for selected depths for F-NTRU, F-NTRU' and YASHE as shown in Table 5.

L	F-NTRU	F-NTRU'	YASHE
5	(12,527)	(12,389)	(11,359)
10	(13,943)	(13,700)	(13,840)
20	(14,1807)	(14,1341)	(14,1705)
30	(15,2727)	(15,2024)	(14,2538)
40	(15,3608)	(15,2679)	(15,3550)
50	(15,4497)	(16,3422)	(15,4440)
100	(16,9227)	(17,7015)	(16,9378)

Table 5. Parameters $(\log(n), \log(q))$ to support depth L evaluations with $\omega = 16$.

10 Implementation Results

We implemented the F-NTRU and F-NTRU' schemes using Shoup's NTL library version 9.6.4 [38] compiled with the GMP 6.1 package. We run our experiments

on a server, which 125 GBs of RAM and Intel Xeon E5-2637v2 64-bit CPU at 3.5 Ghz, using a single thread. The timing results for homomorphic multiplication are summarized in Table 6. In the implementation we use cyclotomic polynomial Φ_m as modulus instead of $x^n + 1$. The timings given in the table are normalized by the number of message slots.

L	Message Slots	F-NTRU	F-NTRU'
5	256	52 msec	23 msec
10	512	260 msec	120 msec
20	1024	1.9 sec	780 msec
30	2048	8.0 sec	3.0 sec

Table 6. Homomorphic amortized multiplication times for radix selection $\omega = 16$. Distinguishing advantage is set to 2^{-80} .

Note in Table 4 that it is advantageous to select large radix sizes ω for faster evaluation. This decreases the run-time significantly, however, the depth is reduced somewhat since the noise is increased. As for ciphertext sizes, we are able to decrease it by a factor of ω . This gives a big advantage especially in large cases, e.g. $\log n \leq 13$.

The main advantage of the proposed F-NTRU and F-NTRU' schemes is that it eliminates costly evaluation keys and the slow relinearization operation that comes with it as demonstrated in Table 7. In addition the homomorphic evaluation is immensely simplified as we do no longer care about keeping track of the evaluation levels per ciphertext. Finally, we compare the ciphertext sizes also in Table 7. As in GSW [25] our constructions require larger ciphertexts for similar depth evaluations compared to the other schemes although this problem is somewhat alleviated in F-NTRU'. The ciphertext size may be further reduced with more aggressive choices of ω .

L	Evaluation Key		Ciphertext		
	F-NTRU/F-NTRU' $\omega = 16$	YASHE $\omega = 2$	F-NTRU $\omega = 16$	F-NTRU' $\omega = 16$	YASHE $\omega = 2$
5	0	3.86 TB	8.4 MB	4.6 MB	87 KB
10	0	478 TB	50 MB	30 MB	820 KB
20	0	n/a	389 MB	220 MB	3.3 MB
30	0	n/a	1.77 GB	970 MB	4.9 MB
40	0	n/a	3.10 GB	1.71 GB	13 MB

Table 7. Evaluation key and ciphertext sizes to support depth L evaluation.

11 Conclusion

We presented a new leveled FHE scheme F-NTRU based on a variant of NTRU and by making use of a recently proposed noise management technique. Our scheme manages to eliminate costly evaluation keys, and any need for key switching and relinearization operations. We rigorously analyzed the security and noise performance of the proposed scheme, and determined parameters needed for correct evaluation of circuits of arbitrary depth. Our scheme makes use of wide key distributions and does not suffer from the Subfield Lattice Attack as LTV and YASHE'. Unlike these schemes, our proposed scheme does not rely on the DSPR assumption.

We analyzed and compared our scheme to the only other NTRU based FHE scheme immune to the attack, i.e. YASHE. While our scheme features larger ciphertext sizes, unlike YASHE, our scheme does not suffer from costly evaluation keys and relinearization operations and therefore supports deep homomorphic evaluation.

12 Acknowledgements

We would like to thank Leo Ducas for clarifying the Subfield Lattice Attack.

References

1. Bos, J.W., Lauter, K., Loftus, J., Naehrig, M.: Cryptography and Coding: 14th IMA International Conference, IMACC 2013, Oxford, UK, December 17-19, 2013. Proceedings, chap. Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme, pp. 45–64. Springer Berlin Heidelberg, Berlin, Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-45239-0_4
2. Bos, J.W., Lauter, K., Naehrig, M.: Private predictive analysis on encrypted medical data. Tech. Rep. MSR-TR-2013-81 (September 2013), <http://research.microsoft.com/apps/pubs/default.aspx?id=200652>
3. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical gapSVP. IACR Cryptology ePrint Archive 2012, 78 (2012)
4. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: Fully homomorphic encryption without bootstrapping. Electronic Colloquium on Computational Complexity (ECCC) 18, 111 (2011)
5. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS. pp. 97–106 (2011)
6. Çetin, G.S., Dai, W., Doröz, Y., Sunar, B.: Homomorphic autocomplete. IACR Cryptology ePrint Archive 2015, 1194 (2015), <http://eprint.iacr.org/2015/1194>
7. Çetin, G.S., Doröz, Y., Sunar, B., Martin, W.J.: An investigation of complex operations with word-size homomorphic encryption. IACR Cryptology ePrint Archive 2015, 1195 (2015), <http://eprint.iacr.org/2015/1195>

8. Çetin, G.S., Doröz, Y., Sunar, B., Savas, E.: Depth optimized efficient homomorphic sorting. In: Progress in Cryptology - LATINCRYPT 2015 - 4th International Conference on Cryptology and Information Security in Latin America, Guadalajara, Mexico, August 23-26, 2015, Proceedings. pp. 61–80 (2015), http://dx.doi.org/10.1007/978-3-319-22174-8_4
9. Chen, Y., Nguyen, P.Q.: BKZ 2.0: Better lattice security estimates. In: ASIACRYPT. pp. 1–20 (2011)
10. Coron, J.S., Mandal, A., Naccache, D., Tibouchi, M.: Fully homomorphic encryption over the integers with shorter public keys. In: CRYPTO. pp. 487–504 (2011)
11. Coron, J.S., Naccache, D., Tibouchi, M.: Public key compression and modulus switching for fully homomorphic encryption over the integers. In: EUROCRYPT. pp. 446–464 (2012)
12. Dai, W., Doroz, Y., Sunar, B.: Accelerating ntru based homomorphic encryption using gpus. In: High Performance Extreme Computing Conference (HPEC), 2014 IEEE. pp. 1–6 (Sept 2014)
13. Dai, W., Dorz, Y., Sunar, B.: Accelerating swhe based pirs using gpus. In: Brenner, M., Christin, N., Johnson, B., Rohloff, K. (eds.) Financial Cryptography and Data Security, Lecture Notes in Computer Science, vol. 8976, pp. 160–171. Springer Berlin Heidelberg (2015), http://dx.doi.org/10.1007/978-3-662-48051-9_12
14. Dai, W., Sunar, B.: cuhe: A homomorphic encryption accelerator library. In: Pasalic, E., Knudsen, L.R. (eds.) Cryptography and Information Security in the Balkans, Lecture Notes in Computer Science, vol. 9540, pp. 169–186. Springer International Publishing (2016), http://dx.doi.org/10.1007/978-3-319-29172-7_11
15. van Dijk, M., Gentry, C., Halevi, S., Vaikuntanathan, V.: Fully homomorphic encryption over the integers. In: EUROCRYPT. pp. 24–43 (2010)
16. Doröz, Y., Hu, Y., Sunar, B.: Homomorphic aes evaluation using the modified ltv scheme. Designs, Codes and Cryptography pp. 1–26 (2015), <http://dx.doi.org/10.1007/s10623-015-0095-1>
17. Doröz, Y., Sunar, B., Hammouri, G.: Financial Cryptography and Data Security: FC 2014 Workshops, BITCOIN and WAHC 2014, Christ Church, Barbados, March 7, 2014, Revised Selected Papers, chap. Bandwidth Efficient PIR from NTRU, pp. 195–207. Springer Berlin Heidelberg, Berlin, Heidelberg (2014), http://dx.doi.org/10.1007/978-3-662-44774-1_16
18. ElGamal, T.: Advances in Cryptology: Proceedings of CRYPTO 84, chap. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms, pp. 10–18. Springer Berlin Heidelberg, Berlin, Heidelberg (1985), http://dx.doi.org/10.1007/3-540-39568-7_2
19. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. IACR Cryptology ePrint Archive 2012, 144 (2012)
20. Gentry, C.: A Fully Homomorphic Encryption Scheme. Ph.D. thesis, Stanford University (2009)
21. Gentry, C., Halevi, S.: Implementing Gentry’s fully-homomorphic encryption scheme. In: EUROCRYPT. pp. 129–148 (2011)
22. Gentry, C., Halevi, S., Smart, N.P.: Better bootstrapping in fully homomorphic encryption. IACR Cryptology ePrint Archive 2011/680 2011 (2011)
23. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. IACR Cryptology ePrint Archive Report 2011/566 (2011), <http://eprint.iacr.org/>
24. Gentry, C., Halevi, S., Smart, N.P.: Homomorphic evaluation of the aes circuit. Cryptology ePrint Archive, Report 2012/099 (2012), <http://eprint.iacr.org/>

25. Gentry, C., Sahai, A., Waters, B.: Advances in Cryptology – CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2013. Proceedings, Part I, chap. Homomorphic Encryption from Learning with Errors: Conceptually-Simpler, Asymptotically-Faster, Attribute-Based, pp. 75–92. Springer Berlin Heidelberg, Berlin, Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-40041-4_5
26. Goldwasser, S., Micali, S.: Probabilistic encryption & how to play mental poker keeping secret all partial information. In: Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing. pp. 365–377. STOC '82, ACM, New York, NY, USA (1982), <http://doi.acm.org/10.1145/800070.802212>
27. Halevi, S., Shoup, V.: HELib, homomorphic encryption library. Internet Source (2012)
28. Hoffstein, J., Pipher, J., Silverman, J.H.: Algorithmic Number Theory: Third International Symposium, ANTS-III Portland, Oregon, USA, June 21–25, 1998 Proceedings, chap. NTRU: A ring-based public key cryptosystem, pp. 267–288. Springer Berlin Heidelberg, Berlin, Heidelberg (1998), <http://dx.doi.org/10.1007/BFb0054868>
29. Lauter, K., López-Alt, A., Naehrig, M.: Progress in Cryptology - LATINCRYPT 2014: Third International Conference on Cryptology and Information Security in Latin America Florianópolis, Brazil, September 17–19, 2014 Revised Selected Papers, chap. Private Computation on Encrypted Genomic Data, pp. 3–27. Springer International Publishing, Cham (2015), http://dx.doi.org/10.1007/978-3-319-16295-9_1
30. Lindner, R., Peikert, C.: Better key sizes (and attacks) for lwe-based encryption. In: CT-RSA. pp. 319–339 (2011)
31. López-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In: Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing. pp. 1219–1234. STOC '12, ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2213977.2214086>
32. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. IACR Cryptology ePrint Archive 2012, 230 (2012)
33. Martin Albrecht, Shi Bai, L.D.: A subfield lattice attack on overstretched ntru assumptions: Cryptanalysis of some fle and graded encoding schemes. Cryptology ePrint Archive, Report 2016/127 (2016), <http://eprint.iacr.org/>
34. Naehrig, M., Lauter, K., Vaikuntanathan, V.: Can homomorphic encryption be practical? In: Proceedings of the 3rd ACM Workshop on Cloud Computing Security Workshop. pp. 113–124. CCSW '11, ACM, New York, NY, USA (2011), <http://doi.acm.org/10.1145/2046660.2046682>
35. Paillier, P.: Advances in Cryptology — EUROCRYPT '99: International Conference on the Theory and Application of Cryptographic Techniques Prague, Czech Republic, May 2–6, 1999 Proceedings, chap. Public-Key Cryptosystems Based on Composite Degree Residuosity Classes, pp. 223–238. Springer Berlin Heidelberg, Berlin, Heidelberg (1999), http://dx.doi.org/10.1007/3-540-48910-X_16
36. Rivest, R.L., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. Foundations of Secure Computation pp. 169–180 (1978)
37. Rivest, R.L., Shamir, A., Adleman, L.M.: A method for obtaining digital signatures and public-key cryptosystems. Commun. ACM 21(2), 120–126 (1978)
38. Shoup, V.: <http://www.shoup.net/ntl/>, NTL: A Library for doing Number Theory

39. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. IACR Cryptology ePrint Archive 2011, 133 (2011)
40. Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. Advances in Cryptology – EUROCRYPT '11 pp. 27–4 (2011)