Packing Messages and Optimizing Bootstrapping in GSW-FHE

Ryo Hiromasa^{1(✉)}, Masayuki Abe², and Tatsuaki Okamoto²

¹ Kyoto University, Kyoto, Japan

hiromasa@ai.soc.i.kyoto-u.ac.jp

² NTT Secure Platform Laboratories, Tokyo, Japan

{abe.masayuki,hiromasa}@lab.ntt.co.jp

Abstract. We construct the first fully homomorphic encryption (FHE) scheme that encrypts *matrices* and supports homomorphic *matrix* addition and multiplication. This is a natural extension of packed FHE and thus supports more complicated homomorphic operations. We optimize the bootstrapping procedure of Alperin-Sheriff and Peikert (CRYPTO 2014) by applying our scheme. Our optimization decreases the lattice approximation factor from $\tilde{O}(n^3)$ to $\tilde{O}(n^{2.5})$. By taking a lattice dimension as a larger polynomial in a security parameter, we can also obtain the same approximation factor as the best known one of standard lattice-based public-key encryption *without* successive dimension-modulus reduction, which was essential for achieving the best factor in prior works on bootstrapping of standard lattice-based FHE.

1 Introduction

Fully homomorphic encryption (FHE) allows us to evaluate any function over encrypted data by only using public information. This can be used, for example, to outsource computations to remote servers without compromising privacy. Since the breakthrough work by Gentry [12, 13], many different varieties of FHE have been proposed [5–8, 11, 17, 18]. To date, the fastest (and simplest) FHE based on the *standard* lattice assumption is the one by Gentry, Sahai, and Waters [17]. (hereafter, referred to as GSW-FHE). However, it is required to take heavy cost for evaluating a large number of ciphertexts. The way to deal with this issue is to *pack* multiple messages into one ciphertext.

Packing messages allows us to apply *single-instruction-multiple data* (SIMD) homomorphic operations to all encrypted messages. In the case where a remote server stores encrypted data and we want to retrieve certain data from this server, we first apply the equality function to every encrypted data. If the stored data have been packed into one ciphertext, we can do that by only one homomorphic evaluation of the equality function. Smart and Vercautren [25], for the first time, showed that applying the Chinese remainder theorem (CRT) to number fields partitions the message space of the Gentry’s FHE [12, 13] scheme into a vector of *plaintext slots*. On the standard lattice-based FHE schemes, Brakerski, Gentry, and Halevi [4] used the method of [22], which described a way to construct packed Regev’s encryption [23], to pack messages in the FHE variants [5–7] of [23]. In this paper,

we construct a matrix variant of [17] (whose security is also based on the standard lattice assumption) to implement SIMD homomorphic operations, and describe how to bring out the potential of our scheme: specifically optimizing *bootstrapping*.

The bootstrapping technique [12, 13] is currently the only way to go from limited amount of homomorphism to unlimited amount of homomorphism. The limited nature is caused by noise terms included in ciphertexts of all known FHE, which are needed to ensure security. Since homomorphic operations increases the noise level and the noise prevents us from correctly decrypting ciphertexts if the level increases too high, it is required to consider methods that reduce the noise. The bootstrapping technique is the one of such a methods, and achieved by homomorphically evaluating the decryption circuit of FHE.

There have recently been the significant progresses [1, 9] in improving the bootstrapping procedure on standard lattice-based FHE. Their progresses stem from the observation that noise terms in ciphertexts of GSW-FHE grow *asymmetrically*: for a parameter n (the dimension in the underlying lattice assumption), the noise of multiplication between two ciphertexts with noise size e_1 and e_2 grows to $e_1 + \text{poly}(n) \cdot e_2$. For example, if we want to multiply ℓ ciphertexts with the same noise size in *sequence*, the noise in the result increases by a factor of $\ell \cdot \text{poly}(n)$, which is in contrast to the noise blowup factor for all known FHE, $\text{poly}(n)^{\log \ell}$. To suppress the growth in noise from the bootstrapping procedure, the two recent developments [1, 9] tried to *sequentialize* the decryption circuit.

Brakerski and Vaikuntanathan [9] transformed the decryption circuit of [17] to a branching program by using the Barrington's theorem [2], and homomorphically evaluated the program. Since the Barrington's theorem can convert the decryption circuit to a polynomial length branching program, evaluating the program increases the noise by a factor of $\text{poly}(n)$. This procedure, however, has a significant drawback: the Barrington's theorem generates a branching program of *large* polynomial length. The scheme [9] also used a kind of *dimension leveraging* technique and successive dimension-modulus reduction to obtain the best approximation factor that is the same as standard lattice-based (plain) PKE.

Unlike most previous works, Alperin-Sheriff and Peikert [1] viewed the decryption as an arithmetic circuit. The decryption of all known standard lattice-based FHE consists of the inner product and rounding: for a ciphertext vector $\mathbf{c}$ and secret key vector $\mathbf{s}$, the decryption algorithm computes $\lfloor \langle \mathbf{c}, \mathbf{s} \rangle \rfloor_2 \in \{0, 1\}$ (where $\lfloor \cdot \rfloor_2$ is the rounding function introduced later). The authors observed that the inner product in the decryption can be expressed as a subset sum of the secret key elements. The subset sum can be computed only in the additive group, and the additive group is isomorphic to a group of cyclic permutations. The authors rewrote the inner product to the sequence of compositions of the cyclic permutations. Since this does not use the Barrington's theorem, the bootstrapping procedure of [1] can refresh ciphertexts faster and keep the noise growth in a *smaller* polynomial than that of [9], but the underlying security assumption was slightly stronger than that of [9]¹. In addition, the procedure of [1] was not fully

¹ By using successive dimension-modulus reduction, [1] can also obtain the same approximation factor as that of [9].

sequentialized, that is, there is a little room for sequentializing the decryption: the rounding.

1.1 Our Results

In this paper, we construct the first FHE scheme that encrypts matrices and supports homomorphic matrix operations. This is a natural extension of packed FHE and supports more complicated homomorphic operations. Using this scheme, we fully sequentialize and thus optimize the bootstrapping procedure of [1]. The result of the optimization is described in the following:

Theorem 1. *Our optimized bootstrapping scheme can be secure assuming the hardness of approximating the standard lattice problem to within the factor $\tilde{O}(n^{1.5}\lambda)$ on any n dimensional lattices.*

For 2^λ hardness, we need to take $n = \Omega(\lambda)$. Asymptotically minimal selection of $n = \tilde{O}(\lambda)$ leads to the approximation factor $\tilde{O}(n^{2.5})$ for the underlying worst-case lattice assumption, which is smaller than $\tilde{O}(n^3)$, the factor of [1]. Using a kind of dimension leveraging technique: selecting a larger dimension $n = \lambda^{1/\epsilon}$ for $\epsilon \in (0, 1)$, we can also obtain the best known approximation factor, $\tilde{O}(n^{1.5+\epsilon})$, without successive dimension-modulus reduction, which was essential for achieving the best factor in the prior works on bootstrapping of standard lattice-based FHE.

1.2 Our Techniques

Matrix GSW-FHE. The starting point of our scheme is the GSW-FHE scheme. In that scheme, a ciphertext of a plaintext $m \in \{0, 1\}$ is a matrix $\mathbf{C} \in \mathbb{Z}_q^{(n+1) \times N}$ such that $\mathbf{sC} = m \cdot \mathbf{sG} + \mathbf{e}$ for a secret key vector $\mathbf{s} \in \mathbb{Z}_q^{n+1}$, small noise vector $\mathbf{e} \in \mathbb{Z}^N$, and fixed matrix $\mathbf{G} \in \mathbb{Z}_q^{(n+1) \times N}$. A simple extension of the plaintext space from bits to binary vectors cannot yield plaintext-slot-wise addition and multiplication. Instead, we use matrices to store binary vectors in their diagonal entries. Actually, our construction even supports homomorphic matrix addition and multiplication that are richer than homomorphic plaintext-slot-wise operations.

Let $\mathbf{S} \in \mathbb{Z}_q^{r \times (n+r)}$ be a secret key matrix, $\mathbf{B} \in \mathbb{Z}_q^{n \times m}$ be a Learning with Errors (LWE) matrix such that $\mathbf{SB} \approx \mathbf{0}$, and $\mathbf{G} \in \mathbb{Z}^{(n+r) \times N}$ be a fixed matrix. To encrypt a square integer matrix $\mathbf{M} \in \{0, 1\}^{r \times r}$, the ciphertext $\mathbf{C} \in \mathbb{Z}^{(n+r) \times N}$ must be of the form $\mathbf{BR} + \mathbf{XG}$ for a matrix $\mathbf{X} \in \mathbb{Z}^{(n+r) \times (n+r)}$ such that $\mathbf{SX} = \mathbf{MS}$, and small random matrix $\mathbf{R} \in \mathbb{Z}^{m \times N}$. The ciphertext $\mathbf{C}$ satisfies $\mathbf{SC} = \mathbf{E} + \mathbf{MSG}$ for a small noise matrix $\mathbf{E} \in \mathbb{Z}^{r \times N}$. Homomorphic matrix addition is just matrix addition. For example, given two ciphertexts $\mathbf{C}_1$ and $\mathbf{C}_2$, it holds that

$$\mathbf{S}(\mathbf{C}_1 + \mathbf{C}_2) = (\mathbf{E}_1 + \mathbf{E}_2) + (\mathbf{M}_1 + \mathbf{M}_2)\mathbf{SG}.$$

Homomorphic matrix multiplication corresponds to a simple preimage sampling and matrix multiplication. For a matrix $\mathbf{C} \in \mathbb{Z}_q^{(n+r) \times N}$, let $\mathbf{G}^{-1}(\mathbf{C})$ be the function that outputs a matrix $\mathbf{X}' \in \mathbb{Z}_q^{N \times N}$ such that $\mathbf{G}\mathbf{X}' \equiv \mathbf{C} \pmod{q}$. If we let $\mathbf{X}'_2 \stackrel{R}{\leftarrow} \mathbf{G}^{-1}(\mathbf{C}_2)$, then it holds that

$$\begin{aligned} \mathbf{S}\mathbf{C}_1\mathbf{X}'_2 &= (\mathbf{E}_1 + \mathbf{M}_1\mathbf{S}\mathbf{G})\mathbf{X}'_2 \\ &= \mathbf{E}_1\mathbf{X}'_2 + \mathbf{M}_1\mathbf{E}_2 + \mathbf{M}_1\mathbf{M}_2\mathbf{S}\mathbf{G}. \end{aligned}$$

Now, the problem is how to construct a matrix $\mathbf{X}$ such that $\mathbf{S}\mathbf{X} = \mathbf{M}\mathbf{S}$. By construction, $\mathbf{S}$ includes an identity matrix: $\mathbf{S} = [\mathbf{I} \parallel \mathbf{S}']$ for a matrix $\mathbf{S}' \in \mathbb{Z}_q^{r \times n}$. The idea is to make $\mathbf{X}$ have $\mathbf{M}\mathbf{S}$ in its top rows and 0 below. This $\mathbf{X}$ clearly satisfies the condition, but cannot publicly be computed without knowing the secret key. We translate the resulting symmetric scheme to the asymmetric one by using the method similar to [3, 24]. In particular, let $\mathbf{M}_{(i,j)} \in \{0, 1\}^{r \times r}$ ($i, j = 1, \dots, r$) be the matrix with 1 in the (i, j) -th entry and 0 in the others. We first publish symmetric encryptions of $\mathbf{M}_{(i,j)}$ for all $i, j \in [r]$. A ciphertext for a plaintext matrix $\mathbf{M}$ is publicly computed by summing up all encryptions of $\mathbf{M}_{(i,j)}$ such that the (i, j) -th entry of $\mathbf{M}$ is equal to 1, and using $\mathbf{B}$ to randomize the sum. Since the public key includes the ciphertexts that encrypt partial information of the secret key, security of our scheme cannot directly be proven from the LWE assumption. The way to deal with this problem is to introduce a circular security assumption.

Optimizing Bootstrapping of [1]. For a dimension d and modulus q , let $\mathbf{c} \in \{0, 1\}^d$ be the $\ell - 1$ -th column of a binary GSW-FHE ciphertext under the secret key $\mathbf{s} \in \mathbb{Z}_q^d$. Since the decryption algorithm of GSW-FHE computes $\lfloor \langle \mathbf{c}, \mathbf{s} \rangle \rfloor_2$ ($\lfloor \cdot \rfloor_2$ is the rounding function that outputs 1 if the input is close to $q/4$ and 0 otherwise), and $\langle \mathbf{c}, \mathbf{s} \rangle = \sum_{i=1}^d c_i s_i = \sum_{i \in [d]: c_i=1} s_i$, the decryption can be viewed as a subset sum of $\{s_i\}_{i \in [d]}$. To bootstrap ciphertexts, we only have to be able to compute additions in $\mathbb{Z}_q$ homomorphically. The additive group $\mathbb{Z}_q^+$ is isomorphic to a group of cyclic permutations, where $x \in \mathbb{Z}_q^+$ corresponds to a cyclic permutation that can be represented by an indicator vector with 1 in the x -th position. The permutation matrix can be obtained from the cyclic rotation of the indicator vector. The addition in $\mathbb{Z}_q^+$ leads to the composition of the permutations (i.e., the multiplication of the corresponding permutation matrices), and the rounding function $\lfloor \cdot \rfloor_2 : \mathbb{Z}_q \rightarrow \{0, 1\}$ can be computed by summing the entries of the indicator vector corresponding to those values in $\mathbb{Z}_q$.

The bootstrapping procedure of [1] consists of two parts that compute an inner product and a rounding operation. The rounding checks equalities and computes summation. The matrix GSW-FHE scheme allows us to rewrite the bootstrapping procedure except for the summation as a *sequence* of homomorphic matrix multiplications, while the procedure of [1] computes only the inner product part as a sequence. Intuitively, our optimization use the matrix GSW-FHE scheme to *sequentialize* the bootstrapping procedure of [1]. The asymmetric noise growth property is more effective in estimating how much noise the procedure yields.

The inner product can be computed by compositions of cyclic permutations. The bootstrapping procedure of [1] represents elements in $\mathbb{Z}_q$ as cyclic permutations, and evaluates their compositions by the naive matrix multiplication algorithm on the ciphertexts that encrypt every elements in the permutation matrices. Instead of that, our bootstrapping procedure uses homomorphic matrix multiplication to directly evaluate the compositions. The rounding part tests for every value close to $q/4$ whether the output of the inner product part encrypts the permutation corresponding to the value, and sums their results (that are 0 or 1). Our procedure also use homomorphic matrix multiplication to realize the equality test. The result of the inner product is represented as an indicator vector, and encrypted component-wise in a SIMD encryption. The inner product equals to x if and only if its indicator vector has 1 in the x -th position. The homomorphic equality test between the inner product and x is computed by homomorphically permuting x -th slot to the first slot in the SIMD ciphertext. The result of the test is encrypted in the first slot. From the above, the bootstrapping procedure except for the summation can be represented as a sequence of $\tilde{O}(\lambda)$ homomorphic multiplications for a security parameter λ .

1.3 Related Work

Multilinear maps[10, 14, 15] are extensions of bilinear maps, and built from variants of FHE. The new multilinear maps construction of Gentry, Gorbunov, and Halevi [15] also starts from GSW-FHE. Recall that in GSW-FHE, a ciphertext of $m \in \{0, 1\}$ is a matrix $\mathbf{C} \in \mathbb{Z}_q^{(n+1) \times N}$ such that $\mathbf{sC} = m \cdot \mathbf{sG} + \mathbf{e}$ for a secret key vector $\mathbf{s} \in \mathbb{Z}_q^{(n+1)}$ and small noise vector $\mathbf{e} \in \mathbb{Z}^N$. That is, valid ciphertexts of GSW-FHE have the secret key as the *approximate eigenvector* and the message as the eigenvalue. The multilinear maps construction of [15] replaced the approximate eigenvector with the *approximate eigenspace* by increasing the dimension. In the construction, an encoding of $\mathbf{M} \in \mathbb{Z}^{r \times r}$ is a matrix $\mathbf{C} \in \mathbb{Z}_q^{N \times N}$ such that $\mathbf{SC} = \mathbf{E} + \mathbf{MS}$ for a random matrix $\mathbf{S} \in \mathbb{Z}_q^{r \times N}$ and small noise matrix $\mathbf{E} \in \mathbb{Z}^{r \times N}$. The approximate eigenspace is the matrix $\mathbf{S}$. To obtain the encoding $\mathbf{C}$, the construction samples a preimage of $\mathbf{MS} + \mathbf{E}$ for the function $f_{\mathbf{S}}(\mathbf{x}) = \mathbf{Sx} \bmod q$. In our scheme, a ciphertext $\mathbf{C} \in \mathbb{Z}_q^{N \times N}$ of $\mathbf{M} \in \mathbb{Z}^{r \times r}$ is a preimage of

$$\mathbf{BR} + \left(\frac{\mathbf{MS}}{0} \right) \mathbf{G}$$

for the function $f_{\mathbf{G}}$. Since the ciphertext $\mathbf{C}$ satisfies $(\mathbf{SG})\mathbf{C} = \mathbf{M}(\mathbf{SG}) + \mathbf{E}$ for a small noise matrix $\mathbf{E} \in \mathbb{Z}^{r \times N}$, the matrix $\mathbf{SG}$ can be seen as the approximate eigenspace.

1.4 Organization

In Section 2, we describe some preliminaries on the LWE assumption and sub-gaussian random variables. In Section 3, we present how to construct a matrix variant of [17]. In Section 4, we show that our scheme improves the bootstrapping procedure of [1].

2 Preliminaries

We denote the set of natural numbers by $\mathbb{N}$, the set of integers by $\mathbb{Z}$, the set of rational numbers by $\mathbb{Q}$, and the set of real numbers by $\mathbb{R}$. Let $\mathbb{G}$ be some group and $\mathcal{P}$ be some probability distribution, then we use $a \xleftarrow{U} \mathbb{G}$ to denote that a is chosen from $\mathbb{G}$ uniformly at random, and use $b \xleftarrow{R} \mathcal{P}$ to denote that b is chosen along $\mathcal{P}$. We take all logarithms to base 2, unless otherwise noted.

We assume that vectors are in column form and are written by using bold lower-case letters, e.g., $\mathbf{x}$, and the i -th element of a vector is denoted by x_i . We denote the ℓ_∞ norm (the maximum norm) of the vector $\mathbf{x}$ by $\|\mathbf{x}\|_\infty$, and the ℓ_2 norm (the Euclidean norm) of $\mathbf{x}$ by $\|\mathbf{x}\|_2$. The inner product between two vectors is denoted by $\langle \mathbf{x}, \mathbf{y} \rangle$. Matrices are written by using bold capital letters, e.g., $\mathbf{X}$, and the i -th column vector of a matrix is denoted by $\mathbf{x}_i$. For a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, we define the ℓ_∞ and ℓ_2 norms of $\mathbf{X}$ as $\|\mathbf{X}\|_\infty := \max_{i \in [n]} \{\|\mathbf{x}_i\|_\infty\}$ and $\|\mathbf{X}\|_2 := \max_{i \in [n]} \{\|\mathbf{x}_i\|_2\}$, respectively. For a matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$, the notation $\mathbf{X}^T \in \mathbb{R}^{n \times m}$ denotes the transpose of $\mathbf{X}$. For matrices $\mathbf{A} \in \mathbb{R}^{m \times n_1}$ and $\mathbf{B} \in \mathbb{R}^{m \times n_2}$, $[\mathbf{A} \parallel \mathbf{B}] \in \mathbb{R}^{m \times (n_1 + n_2)}$ denotes the concatenation of $\mathbf{A}$ with $\mathbf{B}$. When we refer to the $n \times n$ identity matrix, we denote it by $\mathbf{I}_n$.

2.1 Learning with Errors

The *learning with errors (LWE) assumption* was first introduced by Regev [23].

Definition 1 (DLWE). For a security parameter λ , let $n := n(\lambda)$ be an integer dimension, let $q := q(\lambda) \geq 2$ be an integer modulus, and let $\chi := \chi(\lambda)$ be an error distribution over $\mathbb{Z}$. $\text{DLWE}_{n,q,\chi}$ is the problem to distinguish the following two distributions: In the first distribution, a tuple $(\mathbf{a}_i, b_i)$ is sampled from uniform over $\mathbb{Z}_q^n \times \mathbb{Z}_q$; In the second distribution, $\mathbf{s} \xleftarrow{U} \mathbb{Z}_q^n$ and then a tuple $(\mathbf{a}_i, b_i)$ is sampled by sampling $\mathbf{a}_i \xleftarrow{U} \mathbb{Z}_q^n$, $e_i \xleftarrow{R} \chi$, and setting $b_i := \langle \mathbf{a}_i, \mathbf{s} \rangle + e_i \bmod q$. The $\text{DLWE}_{n,q,\chi}$ assumption is that $\text{DLWE}_{n,q,\chi}$ is infeasible.

Recall that GapSVP_γ is the promise problem to distinguish between the case in which the lattice has a vector shorter than $r \in \mathbb{Q}$, and the case in which all the lattice vectors are greater than $\gamma \cdot r$. SIVP_γ is the problem to find the set of short linearly independent vectors in a lattice. $\text{DLWE}_{n,q,\chi}$ has reductions to the standard lattice assumptions as follows. These reductions take χ to be a discrete Gaussian distribution $D_{\mathbb{Z},\alpha q}$ (that is centered around 0 and has parameter αq for some $\alpha < 1$), which is statistically indistinguishable from a B -bounded distribution (i.e., $\mathbb{E}[X] = 0$ and $|X| \leq B$) for an appropriate B .

Corollary 1 ([19–21, 23]). Let $q := q(n) \in \mathbb{N}$ be a power of primes $q := p^r$ or a product of distinct prime numbers $q := \prod_i q_i$ ($q_i := \text{poly}(n)$ for all i), and let $\alpha \geq \sqrt{n}/q$. If there exists an efficient algorithm that solves (average-case) $\text{DLWE}_{n,q,D_{\mathbb{Z},\alpha q}}$,

- there exists an efficient quantum algorithm that can solve $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ and $\text{SIVP}_{\tilde{O}(n/\alpha)}$ in the worst-case for any n -dimensional lattices.

- if in addition we have $q \geq \tilde{O}(2^{n/2})$, there exists an efficient classical algorithm that can solve $\text{GapSVP}_{\tilde{O}(n/\alpha)}$ in the worst-case for any n -dimensional lattices.

2.2 Subgaussian

A real random variable X is subgaussian with parameter s if for all $t \in \mathbb{R}$, its (scaled) moment generating function holds $\mathbb{E}[\exp(2\pi tX)] \leq \exp(\pi s^2 t^2)$. Any B -bounded (centered) random variable X is subgaussian with parameter $B \cdot \sqrt{2\pi}$. Subgaussian random variables have the following two properties that can be easily obtained from the definition of subgaussian random variables:

- Homogeneity: If the subgaussian random variable X has parameter s , then cX is subgaussian with parameter cs .
- Pythagorean additivity: For two subgaussian random variables X_1 and X_2 (that is independent from X_1) with parameter s_1 and s_2 , respectively, $X_1 + X_2$ is subgaussian with parameter $\sqrt{s_1^2 + s_2^2}$.

The above can be extended to vectors. A real random vector $\mathbf{x}$ is subgaussian with parameter s if for all real unit vectors $\mathbf{u}$, their marginal $\langle \mathbf{u}, \mathbf{x} \rangle$ is subgaussian with parameter s . It is clear from the definition that the concatenation of subgaussian variables or vectors, each of which has a parameter s and is independent of the prior one, is also subgaussian with parameter s . The homogeneity and Pythagorean additivity also hold from linearity of vectors. It is known that the euclidean norm of the subgaussian random vector has the following upper bound.

Lemma 1 ([26]). *Let $\mathbf{x} \in \mathbb{R}^n$ be a random vector that has independent subgaussian coordinates with parameter s . Then there exists a universal constant C such that $\Pr[\|\mathbf{x}\|_2 > C \cdot s\sqrt{n}] \leq 2^{-\Omega(n)}$.*

To suppress the growth in noise, Gentry et al. [17] made use of a procedure that decomposes a vector in binary representation. Alperin-Sheriff and Peikert [1] observed that instead of the decomposition procedure, using the following algorithm $\mathbf{G}^{-1}$ that samples a subgaussian random vector allows us to re-randomize errors in ciphertexts and tightly analyze the noise growth in [17]. Lemma 2 can be extended to matrices in the obvious way. Let $\mathbf{g}^T := (1, 2, 2^2, \dots, 2^{\lceil \log q \rceil - 1})$ and $\mathbf{G} := \mathbf{g}^T \otimes \mathbf{I}_n$.

Lemma 2 ([1], which is adapted from [20]). *There is a randomized, efficiently computable function $\mathbf{G}^{-1} : \mathbb{Z}_q^n \rightarrow \mathbb{Z}^{n \cdot \lceil \log q \rceil}$ such that for any $\mathbf{a} \in \mathbb{Z}_q^n$, $\mathbf{x} \stackrel{R}{\leftarrow} \mathbf{G}^{-1}(\mathbf{a})$ is subgaussian with parameter $O(1)$ and $\mathbf{a} = [\mathbf{G}\mathbf{x}]_q$.*

2.3 Homomorphic Encryption, Circular Security, and Bootstrapping

Here we describe the syntax of homomorphic encryption scheme to introduce a definition of circular security and the Gentry's bootstrapping theorem. Let $\mathcal{M}$ and $\mathcal{C}$ be the message and ciphertext space. A homomorphic encryption scheme consists of four algorithms, $\{\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval}\}$.

- $\text{KeyGen}(1^\lambda)$: output a public encryption key pk , a secret decryption key sk , and a public evaluation key evk .
- $\text{Enc}_{\text{pk}}(m)$: using a public key pk , encrypt a plaintext $m \in \mathcal{M}$ into a ciphertext $c \in \mathcal{C}$.
- $\text{Dec}_{\text{sk}}(c)$: using a secret key sk , recover the message encrypted in the ciphertext c .
- $\text{Eval}_{\text{evk}}(f, c_1, \dots, c_\ell)$: using the evaluation key evk , output a ciphertext $c_f \in \mathcal{C}$ that is obtained by applying the function $f: \mathcal{M}^\ell \rightarrow \mathcal{M}$ to $c_1, \dots, c_\ell$.

To prove the security of our construction, we introduce a special kind of circular security for a homomorphic encryption scheme.

Definition 2 (Circular security). *Let $\mathcal{K}$ be the key space defined by a security parameter λ . Let f be a function from $\mathcal{K}$ to $\mathcal{C}$. A homomorphic encryption scheme $\text{HE} = \{\text{KeyGen}, \text{Enc}, \text{Dec}, \text{Eval}\}$ is circular secure with respect to f if for all probabilistic polynomial-time adversary $\mathcal{A}$, the advantage of $\mathcal{A}$ in the following game is negligible in λ :*

1. A challenger computes $(\text{pk}, \text{sk}, \text{evk}) \xleftarrow{R} \text{KeyGen}(1^\lambda)$, and chooses a bit $b \xleftarrow{U} \{0, 1\}$.
2. Let $f_+ : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ be a function that computes $f_+(x, y) := x + y \in \mathcal{M}$. The challenger computes a challenge ciphertext c^* as follows and sends it to $\mathcal{A}$.

$$c^* := \begin{cases} \text{Eval}_{\text{evk}}(f_+, \text{Enc}_{\text{pk}}(0), f(\text{sk})) & \text{if } b = 0, \\ \text{Enc}_{\text{pk}}(0) \in \mathcal{C} & \text{otherwise.} \end{cases}$$

3. $\mathcal{A}$ outputs a guess $b' \in \{0, 1\}$.

The advantage of $\mathcal{A}$ is $\Pr[b = b'] - 1/2$.

In LWE-based FHE schemes, $\text{Eval}_{\text{evk}}(f_+, \text{Enc}_{\text{pk}}(0), f(\text{sk}))$ can be seen as a kind of ciphertexts that encrypt $f(\text{sk})$. This is why we call the above security notion circular security.

Gentry’s bootstrapping theorem states the way to go from limited homomorphism to unlimited homomorphism. This relates to augmented decryption circuits.

Definition 3 (Augmented decryption circuit). *Let $(\text{sk}, \text{pk}, \text{evk})$ be a tuple of keys generated appropriately, and $\mathcal{C}$ be the set of decryptable ciphertexts. Then the set of augmented decryption functions $\{f_{c_1, c_2}\}_{c_1, c_2 \in \mathcal{C}}$ is defined by*

$$f_{c_1, c_2}(x) = \overline{\text{Dec}_x(c_1)} \wedge \overline{\text{Dec}_x(c_2)}.$$

That is, the function uses its input as the secret key, decrypts c_1 and c_2 , and returns the NAND of the results.

Theorem 2 (Bootstrapping theorem [12, 13]). *A scheme that can evaluate the family of the augmented decryption circuits can be transformed into a “leveled” FHE scheme (in which KeyGen takes as additional input 1^L and we can only evaluate depth L circuits) with the same decryption circuit, ciphertext space, and public key.*

In addition, if the above scheme is weak circular secure (remains secure against an adversary that can obtain encryptions of the bits of the secret key), it can be “pure” FHE scheme (in which the number of homomorphic evaluations is unlimited).

3 Matrix GSW-FHE

We translate [17] to be able to encrypt a *matrix* and homomorphically compute *matrix* addition and multiplication. This is a natural extension of packed FHE schemes. In Section 3.1, we present our matrix FHE scheme. In Section 3.2, we discuss the relationship between our scheme and packed FHE schemes.

3.1 Construction

Let λ be the security parameter. Our scheme is parameterized by an integer lattice dimension n , an integer modulus q , and a distribution χ over $\mathbb{Z}$ that is assumed to be subgaussian, all of which depends on λ . We let $\ell := \lceil \log q \rceil$, $m := O((n+r) \log q)$, and $N := (n+r) \cdot \ell$. Let r be the number of bits to be encrypted, which defines the message space $\{0, 1\}^{r \times r}$. The ciphertext space is $\mathbb{Z}_q^{(n+r) \times N}$. Our scheme uses the rounding function $\lfloor \cdot \rfloor_2$ that for any $x \in \mathbb{Z}_q$, $\lfloor x \rfloor_2$ outputs 1 if x is close to $q/4$, and 0 otherwise. Recall that $\mathbf{g}^T = (1, 2, \dots, 2^{\ell-1})$ and $\mathbf{G} = \mathbf{g}^T \otimes \mathbf{I}_{n+r}$.

- **KeyGen**($1^\lambda, r$): Set the parameters n, q, m, ℓ, N , and χ as described above. Sample a uniformly random matrix $\mathbf{A} \xleftarrow{U} \mathbb{Z}_q^{n \times m}$, secret key matrix $\mathbf{S}' \xleftarrow{R} \chi^{r \times n}$, and noise matrix $\mathbf{E} \xleftarrow{R} \chi^{r \times m}$. Let $\mathbf{S} := [\mathbf{I}_r \parallel -\mathbf{S}'] \in \mathbb{Z}_q^{r \times (n+r)}$. We denote by $\mathbf{s}_i^T$ the i -th row of $\mathbf{S}$. Set

$$\mathbf{B} := \left(\frac{\mathbf{S}'\mathbf{A} + \mathbf{E}}{\mathbf{A}} \right) \in \mathbb{Z}_q^{(n+r) \times m}.$$

Let $\mathbf{M}_{(i,j)} \in \{0, 1\}^{r \times r}$ ($i, j = 1, \dots, r$) be the matrix with 1 in the (i, j) -th position and 0 in the others. For all $i, j = 1, \dots, r$, first sample $\mathbf{R}_{(i,j)} \xleftarrow{U} \{0, 1\}^{m \times N}$, and set

$$\mathbf{P}_{(i,j)} := \mathbf{B}\mathbf{R}_{(i,j)} + \left(\frac{\mathbf{M}_{(i,j)}\mathbf{S}}{\mathbf{0}} \right) \mathbf{G} \in \mathbb{Z}_q^{(n+r) \times N}.$$

Output $\text{pk} := (\{\mathbf{P}_{(i,j)}\}_{i,j \in [r]}, \mathbf{B})$ and $\text{sk} := \mathbf{S}$.

- **SecEnc_{sk}**($\mathbf{M} \in \{0, 1\}^{r \times r}$): Sample a random matrices $\mathbf{A}' \xleftarrow{U} \mathbb{Z}_q^{n \times N}$ and $\mathbf{E}' \xleftarrow{R} \chi^{r \times N}$, parse $\mathbf{S} = [\mathbf{I}_r \parallel -\mathbf{S}']$, and output the ciphertext

$$\mathbf{C} := \left[\left(\frac{\mathbf{S}'\mathbf{A}' + \mathbf{E}}{\mathbf{A}'} \right) + \left(\frac{\mathbf{M}\mathbf{S}}{\mathbf{0}} \right) \mathbf{G} \right]_q \in \mathbb{Z}_q^{(n+r) \times N}.$$

- $\text{PubEnc}_{\text{pk}}(\mathbf{M} \in \{0,1\}^{r \times r})$: Sample a random matrix $\mathbf{R} \xleftarrow{U} \{0,1\}^{m \times N}$, and output the ciphertext

$$\mathbf{C} := \mathbf{B}\mathbf{R} + \sum_{i,j \in [r]: \mathbf{M}[i,j]=1} \mathbf{P}_{(i,j)} \in \mathbb{Z}_q^{(n+r) \times N},$$

where $\mathbf{M}[i,j]$ is the (i,j) -th element of $\mathbf{M}$.

- $\text{Dec}_{\text{sk}}(\mathbf{C})$: Output the matrix $\mathbf{M} = (\lfloor \langle \mathbf{s}_i, \mathbf{c}_{j\ell-1} \rangle \rfloor_2)_{i,j \in [r]} \in \{0,1\}^{r \times r}$.
- $\mathbf{C}_1 \oplus \mathbf{C}_2$: Output $\mathbf{C}_{\text{add}} := \mathbf{C}_1 + \mathbf{C}_2 \in \mathbb{Z}_q^{(n+r) \times N}$ as the result of homomorphic addition between the input ciphertexts.
- $\mathbf{C}_1 \odot \mathbf{C}_2$: Output $\mathbf{C}_{\text{mult}} := \mathbf{C}_1 \mathbf{G}^{-1}(\mathbf{C}_2) \in \mathbb{Z}_q^{(n+r) \times N}$ as the result of homomorphic multiplication between the input ciphertexts.

Definition 4. We say that a ciphertext $\mathbf{C}$ encrypts a plaintext matrix $\mathbf{M}$ with noise matrix $\mathbf{E}$ if $\mathbf{C}$ is an encryption of $\mathbf{M}$ and $\mathbf{E} = \mathbf{S}\mathbf{C} - \mathbf{M}\mathbf{S}\mathbf{G} \pmod{q}$.

The following lemma states the correctness of our asymmetric encryption. Similar to this, the correctness of our symmetric encryption can be proven immediately.

Lemma 3. If a ciphertext $\mathbf{C}$ encrypts a plaintext matrix $\mathbf{M} \in \{0,1\}^{r \times r}$ with noise matrix $\mathbf{E}$ such that $\|\mathbf{E}\|_\infty < q/8$, then $\text{Dec}_{\text{sk}}(\mathbf{C}) = \mathbf{M}$.

Proof. We have

$$\begin{aligned} \mathbf{S}\mathbf{C} &= \mathbf{S} \left(\mathbf{B}\mathbf{R} + \sum_{i,j \in [r]: \mathbf{M}[i,j]=1} \mathbf{B}\mathbf{R}_{(i,j)} + \left(\frac{\mathbf{M}\mathbf{S}}{\mathbf{0}} \right) \mathbf{G} \right) \\ &= \mathbf{E}\mathbf{R} + \sum_{i,j \in [r]: \mathbf{M}[i,j]=1} \mathbf{E}\mathbf{R}_{(i,j)} + \mathbf{M}\mathbf{S}\mathbf{G} \\ &= \mathbf{E}\mathbf{R} + \sum_{i,j \in [r]: \mathbf{M}[i,j]=1} \mathbf{E}\mathbf{R}_{(i,j)} + [\mathbf{M}(\mathbf{g}^T \otimes \mathbf{I}_r) \parallel -\mathbf{M}\mathbf{S}'(\mathbf{g}^T \otimes \mathbf{I}_n)] \end{aligned}$$

Because of $\|\mathbf{E}(\mathbf{R} + \sum_{i,j \in [r]: \mathbf{M}[i,j]=1} \mathbf{R}_{(i,j)})\|_\infty < q/8$ and $2^{\ell-2} \in [q/4, q/2)$, for all $i, j = 1, \dots, r$, $\langle \mathbf{s}_i, \mathbf{c}_{j\ell-1} \rangle \approx q/4$ if $m_{i,j} = 1$, and $\langle \mathbf{s}_i, \mathbf{c}_{j\ell-1} \rangle \approx 0$ otherwise.

Security of SecEnc directly holds from $\text{DLWE}_{n,q,\chi}$. For a matrix $\mathbf{M} \in \{0,1\}^{r \times r}$, let $f_{\mathbf{M}}$ be a function from $\mathbb{Z}_q^{r \times (n+r)}$ to $\mathbb{Z}_q^{(n+r) \times N}$ such that for a matrix $\mathbf{S} \in \mathbb{Z}_q^{r \times (n+r)}$,

$$f_{\mathbf{M}}(\mathbf{S}) = \left(\frac{\mathbf{M}\mathbf{S}}{\mathbf{0}} \right) \mathbf{G} \in \mathbb{Z}_q^{(n+r) \times N}.$$

The security of PubEnc directly holds by $\text{DLWE}_{n,q,\chi}$ and assuming our scheme circular secure with respect to $f_{\mathbf{M}_{(i,j)}}$.

Lemma 4. Let $\mathbf{B}, \mathbf{M}_{(i,j)}, \mathbf{R}_{(i,j)}, \mathbf{P}_{(i,j)}$ ($i, j = 1, \dots, r$) be the matrices generated in KeyGen , and $\mathbf{R}$ be the matrix generated in PubEnc . For every $i, j = 1, \dots, r$, if our scheme is circular secure with respect to $f_{\mathbf{M}_{(i,j)}}$ and $\text{DLWE}_{n,q,\chi}$ holds, then the joint distribution $(\mathbf{B}, \mathbf{B}\mathbf{R}_{(i,j)}, \mathbf{P}_{(i,j)}, \mathbf{B}\mathbf{R})$ is computationally indistinguishable from uniform over $\mathbb{Z}_q^{(n+r) \times m} \times \mathbb{Z}_q^{(n+r) \times N} \times \mathbb{Z}_q^{(n+r) \times N} \times \mathbb{Z}_q^{(n+r) \times N}$.

We need to estimate the noise growth by the evaluation of homomorphic matrix addition and multiplication. Similar to [1], we employ the properties of subgaussian random variables for tight analysis. We collect the results of the estimation in the following lemma.

Lemma 5. *Let $\mathbf{S} \in \mathbb{Z}^{r \times (n+r)}$ be a secret key matrix. Let $\mathbf{C}_1 \in \mathbb{Z}_q^{(n+r) \times N}$ and $\mathbf{C}_2 \in \mathbb{Z}_q^{(n+r) \times N}$ be ciphertexts that encrypt $\mathbf{M}_1 \in \{0,1\}^{r \times r}$ and $\mathbf{M}_2 \in \{0,1\}^{r \times r}$ with noise matrices $\mathbf{E}_1 \in \mathbb{Z}^{r \times N}$ and $\mathbf{E}_2 \in \mathbb{Z}^{r \times N}$, respectively. Let $\mathbf{e}_{1,i}^T \in \mathbb{Z}^{1 \times N}$ ($i = 1, \dots, r$) be the i -th row vector of $\mathbf{E}_1$. Let $\mathbf{C}_{add} := \mathbf{C}_1 \oplus \mathbf{C}_2$ and $\mathbf{C}_{mult} \stackrel{R}{\leftarrow} \mathbf{C}_1 \odot \mathbf{C}_2$. Then, we have*

$$\begin{aligned} \mathbf{S}\mathbf{C}_{add} &= \mathbf{E}_{add} + (\mathbf{M}_1 + \mathbf{M}_2)\mathbf{S}\mathbf{G} \in \mathbb{Z}_q^{r \times N}, \\ \mathbf{S}\mathbf{C}_{mult} &= \mathbf{E}_{mult} + (\mathbf{M}_1\mathbf{M}_2)\mathbf{S}\mathbf{G} \in \mathbb{Z}_q^{r \times N}, \end{aligned}$$

where $\mathbf{E}_{add} := \mathbf{E}_1 + \mathbf{E}_2$ and $\mathbf{E}_{mult} := \mathbf{E} + \mathbf{M}_1\mathbf{E}_2$. In particular, $\mathbf{E}$ has in the i -th row the independent subgaussian entries with parameter $O(\|\mathbf{e}_{1,i}\|_2)$.

Proof. We can immediately prove the statements for $\mathbf{C}_{add}$. For $\mathbf{C}_{mult}$, we have

$$\begin{aligned} \mathbf{S}\mathbf{C}_{mult} &= \mathbf{S}\mathbf{C}_1\mathbf{G}^{-1}(\mathbf{C}_2) \\ &= (\mathbf{E}_1 + \mathbf{M}_1\mathbf{S}\mathbf{G})\mathbf{G}^{-1}(\mathbf{C}_2) \\ &= \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) + \mathbf{M}_1\mathbf{E}_2 + \mathbf{M}_1\mathbf{M}_2\mathbf{S}\mathbf{G}. \end{aligned}$$

From the subgaussian properties and Lemma 2, we can see that the i -th row entries of $\mathbf{E} := \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2)$ are independent subgaussian with parameter $O(\|\mathbf{e}_{1,i}\|_2)$.

Similar to the original GSW scheme, our scheme also has the asymmetric noise growth property, and thereby computing a polynomial length chain of homomorphic multiplications incurs the noise growth by a multiplicative polynomial factor. For ease of analyzing our optimized bootstrapping procedure described in the next section, we set the following corollary immediately proven from Lemma 5 and the properties of subgaussian random variables. This corollary includes the fixed ciphertext $\mathbf{G} \in \mathbb{Z}^{(n+r) \times N}$ of the message $\mathbf{I}_r$ with noise $\mathbf{0}$. This makes the noise in the output ciphertext subgaussian and independent from the noise in the input ciphertexts.

Corollary 2. *For $i = 1, \dots, k$, let $\mathbf{C}_i \in \mathbb{Z}^{(n+r) \times N}$ be a ciphertext that encrypts a message matrix $\mathbf{M}_i \in \{0,1\}^{r \times r}$ such that for a matrix $\mathbf{E} \in \mathbb{Z}^{r \times N}$, $\|(\mathbf{M}_i\mathbf{E})^T\|_2 \leq \|\mathbf{E}^T\|_2$ with noise matrix $\mathbf{E}_i \in \mathbb{Z}^{r \times N}$. Let*

$$\mathbf{C} \stackrel{R}{\leftarrow} \bigodot_{i=1}^k \mathbf{C}_i \odot \mathbf{G} = \mathbf{C}_1 \odot (\mathbf{C}_2 \odot (\dots (\mathbf{C}_{k-1} \odot (\mathbf{C}_k \odot \mathbf{G}))) \dots).$$

For $i = 1, \dots, k$, let $\mathbf{e}_i^T$ be a row vector of $\mathbf{E}_i$ whose norm is equal to $\|\mathbf{E}_i^T\|_2$, and $\mathbf{e}^T := [\mathbf{e}_1^T \parallel \mathbf{e}_2^T \parallel \dots \parallel \mathbf{e}_k^T] \in \mathbb{Z}^{1 \times kN}$. Then the noise matrix of $\mathbf{C}$ has in every row the independent subgaussian entries with parameter $O(\|\mathbf{e}\|_2)$.

Proof. The ciphertext $\mathbf{C}$ encrypts the message $\prod_{i=1}^k \mathbf{M}_i$ with noise $\mathbf{E}_1 \mathbf{X}_1 + \sum_{i=2}^k (\prod_{j=1}^{i-1} \mathbf{M}_j) \mathbf{E}_i \mathbf{X}_i$, where $\mathbf{X}_i$ is the matrix used in the evaluation of each $\odot$. By Lemma 5, the elements of $\mathbf{E}_1 \mathbf{X}_1$ in every row are independent and subgaussian with parameter $O(\|\mathbf{e}_1\|_2)$. Since we have $\|(\mathbf{M}_i \mathbf{E})^T\|_2 \leq \|\mathbf{E}^T\|_2$, $(\prod_{j=1}^{i-1} \mathbf{M}_j) \mathbf{E}_i \mathbf{X}_i$ has in its every row the independent subgaussian entries with parameter $O(\|\mathbf{e}_i\|_2)$. By the Pythagorean additivity of subgaussian random variables, $\mathbf{E}_1 \mathbf{X}_1 + \sum_{i=2}^k (\prod_{j=1}^{i-1} \mathbf{M}_j) \mathbf{E}_i \mathbf{X}_i$ has in every row the independent subgaussian entries with parameter $O(\|\mathbf{e}\|_2)$.

3.2 Relation to Packed FHE

The matrix GSW-FHE above is a natural extension of packed FHE. Plaintext slots in packed FHE correspond to diagonal entries of plaintext matrices in the matrix GSW-FHE scheme. It is easy to see that we can correctly compute homomorphic slot-wise addition and multiplication. In applications of packed FHE such as in [16], we may want to permute plaintext slots. This can be achieved by multiplying the encryptions of a permutation and its inverse from left and right. Security and correctness of the following algorithms clearly holds from Lemmas 4 and 5.

- **SwitchKeyGen**($\mathbf{S}, \sigma$): Given a secret key matrix $\mathbf{S} \in \mathbb{Z}_q^{r \times (n+r)}$ and a permutation σ , let $\Sigma \in \{0, 1\}^{r \times r}$ be a matrix corresponding to σ , and generate

$$\begin{aligned} \mathbf{W}_\sigma &\stackrel{R}{\leftarrow} \text{SecEnc}_{\mathbf{S}}(\Sigma), \\ \mathbf{W}_{\sigma^{-1}} &\stackrel{R}{\leftarrow} \text{SecEnc}_{\mathbf{S}}(\Sigma^T). \end{aligned}$$

Output the switch key $\text{ssk}_\sigma := (\mathbf{W}_\sigma, \mathbf{W}_{\sigma^{-1}})$.

- **SlotSwitch** $_{\text{ssk}_\sigma}(\mathbf{C})$: Take as input a switch key ssk_σ and a ciphertext $\mathbf{C}$, output

$$\mathbf{C}_{\sigma} \stackrel{R}{\leftarrow} \mathbf{W}_\sigma \odot (\mathbf{C} \odot (\mathbf{W}_{\sigma^{-1}} \odot \mathbf{G})),$$

where $\mathbf{G} \in \mathbb{Z}^{(n+r) \times N}$ is the fixed encryption of $\mathbf{I}_r$ with noise zero.

One nice feature of our plaintext-slot switching is that it does not suffer from the inconvenience of the security as in [4]: we do not have to use a larger modulus than the matrix GSW-FHE scheme. Brakerski et al. [4] made use of a larger modulus $Q = 2^\ell q$ to suppress noise growth when switching decryption keys, so the security of the plaintext-slot switching in [4] must have related to Q . The larger modulus leads the larger modulus-to-noise ratio. To obtain the same security level as the SIMD scheme of [4], it was required to select a larger dimension. As opposed to this, our plaintext-slot switching can use the same modulus as the matrix GSW-FHE scheme.

4 Optimizing Bootstrapping

We describe how to optimize the bootstrapping procedure of [1] by using our scheme. In Section 4.1, we present the optimized bootstrapping procedure outlined in Section 1.2, whose correctness and security are discussed in Section 4.2.

4.1 Optimized Procedure

Let Q be the modulus of the ciphertext to be refreshed. Using the dimension-modulus reduction technique [7, 9], we can publicly switch the modulus and the dimension to the arbitrary and possibly smaller ones $q, d = \tilde{O}(\lambda)$. Here, q has the form $q := \prod_{i=1}^t r_i$, where r_i are small and powers of distinct primes (and hence pairwise coprime). The following lemma allows us to choose a sufficiently large q that the correctness of the dimension-modulus reduction holds by letting it be the product of all maximal prime powers r_i bounded by $O(\log \lambda)$, and then there exists $t = O(\log \lambda / \log \log \lambda)$.

Lemma 6 ([1]). *For all $x \geq 7$, the product of all maximal prime powers $r_i \leq x$ is at least $\exp(3x/4)$.*

By CRT, $\mathbb{Z}_q^+$ is isomorphic to the direct product $\mathbb{Z}_{r_1}^+ \times \cdots \times \mathbb{Z}_{r_t}^+$. For all $i \in [t]$, $x \in \mathbb{Z}_{r_i}^+$ corresponds to a cyclic permutation that can be represented by a indicator vector with 1 in the x -th position. Let $\phi_i : \mathbb{Z}_q \rightarrow \{0, 1\}^r$ be the isomorphism of an element in $\mathbb{Z}_q$ into the cyclic permutation that corresponds to an element in $\mathbb{Z}_{r_i}$, where $r := \max_i \{r_i\}$.

Our optimized bootstrapping procedure consists of two algorithms, **BootKeyGen** and **Bootstrap**. The procedure can be used to refresh ciphertexts of all known standard LWE-based FHE. We achieve the input ciphertext $\mathbf{c} \in \{0, 1\}^d$ for **Bootstrap** from the dimension-modulus reduction and bit-decomposition of the ciphertext to be bootstrapped, and let $\mathbf{s} \in \mathbb{Z}_q^d$ be a secret key that corresponds to $\mathbf{c}$. This pre-processing is the same as that in [1], so see for further details.

- **BootKeyGen**($\mathbf{sk}, \mathbf{s}$): given a secret key $\mathbf{sk}$ for our scheme and the secret key $\mathbf{s} \in \mathbb{Z}_q^d$ for ciphertexts to be refreshed, output a bootstrapping key. For every $i \in [t]$ and $j \in [d]$, let $\pi_{\phi_i(s_j)}$ be the permutation corresponding to $\phi_i(s_j)$, and compute

$$\begin{aligned} \tau_{i,j} &\stackrel{R}{\leftarrow} \text{SecEnc}_{\mathbf{sk}}(\text{diag}(\phi_i(s_j))), \\ \text{ssk}_{i,j} &\stackrel{R}{\leftarrow} \text{SwitchKeyGen}(\mathbf{sk}, \pi_{\phi_i(s_j)}), \end{aligned}$$

where for a vector $\mathbf{x} \in \mathbb{Z}^r$, $\text{diag}(\mathbf{x}) \in \mathbb{Z}^{r \times r}$ is the square integer matrix that has $\mathbf{x}$ in its diagonal entries and 0 in the others. In addition, we generate hints to check the equality on packed indicator vectors. For every $i \in [t]$, and $x \in \mathbb{Z}_q$ such that $\lfloor x \rfloor_2 = 1$ ², generate

$$\text{ssk}_{\phi_i(x)} \stackrel{R}{\leftarrow} \text{SwitchKeyGen}(\mathbf{sk}, \pi_{\phi_i(x)}),$$

where $\pi_{\phi_i(x)}$ is the cyclic permutation that maps the $(x \bmod r_i)$ -th row to the first row in the matrix. To mask the first plaintext slot, generate an encryption of $(1, 0, \dots, 0)$:

$$\mathbf{P}_{(1,0,\dots,0)} \stackrel{R}{\leftarrow} \text{SecEnc}_{\mathbf{sk}}(\text{diag}((1, 0, \dots, 0))).$$

² Obviously, our procedure can work on not only the rounding function $\lfloor \cdot \rfloor_2$ but also some arbitrary functions $f : \mathbb{Z}_q \rightarrow \{0, 1\}$.

Output the bootstrapping key

$$\mathbf{bk} := \{(\tau_{i,j}, \mathbf{ssk}_{i,j}, \mathbf{P}_{(1,0,\dots,0)}, \mathbf{ssk}_{\phi_i(x)})\}_{i \in [t], j \in [d], x \in \mathbb{Z}_q : [x]_2 = 1}.$$

- **Bootstrap_{bk}**($\mathbf{c}$): Given a bootstrapping key $\mathbf{bk}$ and a ciphertext $\mathbf{c} \in \mathbb{Z}_q^d$, output the refreshed ciphertext $\mathbf{C}^*$. The decryption of all FHE based on the standard LWE computes $\lfloor \langle \mathbf{c}, \mathbf{s} \rangle \rfloor_2$. The algorithm **Bootstrap** consists of two phases that evaluate the inner product and rounding.

Inner Product: For every $i \in [t]$, homomorphically compute an encryption of $\phi_i(\langle \mathbf{c}, \mathbf{s} \rangle)$. Let $h := \min\{j \in [d] : c_j = 1\}$. For $i = 1, \dots, t$, set $\mathbf{C}_i^* := \tau_{i,h}$, and iteratively compute

$$\mathbf{C}_i^* \xleftarrow{R} \text{SlotSwitch}_{\mathbf{ssk}_{i,j}}(\mathbf{C}_i^*)$$

for $j = h + 1, \dots, d$ such that $c_j = 1$.

Rounding: For each $x \in \mathbb{Z}_q$ such that $[x]_2 = 1$, homomorphically check the equality between x and $\langle \mathbf{c}, \mathbf{s} \rangle$, and sum their results. The refreshed ciphertext is comuted as:

$$\mathbf{C}^* \xleftarrow{R} \bigoplus_{x \in \mathbb{Z}_q : [x]_2 = 1} \left(\bigodot_{i \in [t]} \left(\text{SlotSwitch}_{\mathbf{ssk}_{\phi_i(x)}}(\mathbf{C}_i^*) \right) \odot \mathbf{P}_{(1,0,\dots,0)} \right). \quad (1)$$

The post-processing is almost the same as that in [1] except for the way to extract a matrix ciphertext. When finishing the bootstrapping procedure, we have a ciphertext $\mathbf{C}^*$ that encrypts in the first slot the same plaintext as the ciphertext $\mathbf{c}$. A vector ciphertext like [5,6,8] can be obtained to just take the $\ell - 1$ -th column vector of $\mathbf{C}^*$, and a matrix ciphertext like [17] can be obtained by removing from the second row to the r -th row and from the $l + 1$ -th column to r -th column, and aggregating the remainders. We can utilize the key-switching procedure [5,8] for switching from $\mathbf{s}_1$ back to the original secret key $\mathbf{s}$. This requires us to assume circular security.

Our bootstrapping procedure is more time- and space- efficient than that of [1]. The procedure [1] encrypts every elements of the permutation matrices corresponding to the secret key elements, and homomorphically evaluates naive matrix multiplications to obtain encryptions of compositions of permutations. In our procedure, a permutation is encrypted in one ciphertext, and a composition is computed by two homomorphic multiplications. This makes our procedure time-efficient by roughly a $O(\log^2 \lambda)$ factor, and space-efficient by a $O(\log \lambda)$ factor.

4.2 Correctness and Security

From the security of our scheme, it is easy to see that our bootstrapping procedure can be secure by assuming the circular security and DLWE. Correctness holds as the following lemma.

Lemma 7. *Let sk be the secret key for our scheme. Let $\mathbf{c}$ and $\mathbf{s}$ be a ciphertext and secret key described in our bootstrapping procedure. Then, for $\text{bk} \xleftarrow{R} \text{BootKeyGen}(\text{sk}, \mathbf{s})$, the refreshed ciphertext $\mathbf{C}^* \xleftarrow{R} \text{Bootstrap}_{\text{bk}}(\mathbf{c})$ encrypts $\lfloor \langle \mathbf{s}, \mathbf{c} \rangle \rfloor_2 \in \{0, 1\}$ in the first slot.*

Proof. From Lemma 5 and group homomorphism of ϕ_i , $\mathbf{C}_i^*$ encrypts $\phi_i(\lfloor \langle \mathbf{s}, \mathbf{c} \rangle \rfloor_q)$. Since $\mathbb{Z}_q$ is isomorphic to $\mathbb{Z}_{r_1} \times \cdots \times \mathbb{Z}_{r_t}$ by CRT, $\bigodot_{i \in [t]} (\text{SlotSwitch}_{\text{ssk}_{\phi_i(x)}}(\mathbf{C}_i^*)) \odot \mathbf{P}_{(1,0,\dots,0)}$ encrypts 1 in the first slot if and only if $x = \langle \mathbf{s}, \mathbf{c} \rangle \bmod q$. Finally, $\mathbf{C}^*$ encrypts 1 if and only if $\lfloor \langle \mathbf{s}, \mathbf{c} \rangle \rfloor_2 = 1$.

Here, we let s be the Gaussian parameter. Recall that n is the LWE dimension, r is the number of encrypted bits, $\ell = \lceil \log Q \rceil$, $N = (n + r) \cdot \ell$, $t = O(\log \lambda / \log \log \lambda)$, $d = \tilde{O}(\lambda)$ and $q = \tilde{O}(\lambda)$. We estimate the noise growth by our optimized bootstrapping procedure.

Lemma 8. *For any ciphertext $\mathbf{c} \in \{0, 1\}^d$ described in our bootstrapping procedure, the noise in the refreshed ciphertext $\mathbf{C}^* \xleftarrow{R} \text{Bootstrap}_{\text{bk}}(\mathbf{c})$ has independent subgaussian entries with parameter $O(s\sqrt{n\ell dtq})$, except with probability $2^{-\Omega((n+r)\ell dt)}$ over the random choice of bk and Bootstrap .*

Proof. Since the parenthesized part before the additions in Eq. (1) can be broken down into a sequence of $O(dt)$ homomorphic multiplications, Corollary 2 and Lemma 1 tell us that the term has subgaussian noise with parameter $O(s\sqrt{Ndt})$, except with probability $2^{-\Omega(Ndt)}$. From the Pythagorean additivity of subgaussian random variables and $N = (n + r) \cdot \ell$, the noise in $\mathbf{C}^*$ are subgaussian with parameter $O(s\sqrt{(n + r)\ell dtq})$, and so $O(s\sqrt{n\ell dtq})$ by the fact $n > r$.

From the above lemma, we can see that our procedure refreshes ciphertexts with error growth by the $O(\sqrt{n\ell dtq})$ factor. Our scheme can evaluate its augmented decryption circuit by choosing a larger modulus than the final noise, and thus be pure FHE by the Gentry's bootstrapping theorem (Theorem 2) and the circular security assumption.

Theorem 3. *Our optimized bootstrapping scheme can be correct and secure assuming*

- the quantum worst-case hardness of approximating $\text{GapSVP}_{\tilde{O}(n^{1.5}\lambda)}$ and $\text{SIVP}_{\tilde{O}(n^{1.5}\lambda)}$,
- or the classical worst-case hardness of approximating $\text{GapSVP}_{\tilde{O}(n^2\lambda)}$

on any n dimensional lattice.

Proof. By Lemma 1, to rely on the quantum worst-case hardness, we choose $s = \Theta(\sqrt{n})$. From Lemma 8, for correctness we only have to select $Q = \tilde{\Omega}(n\lambda \log Q)$, which satisfies $Q = \tilde{O}(n\lambda)$. Since the LWE inverse error rate is $1/\alpha = Q/s = \tilde{O}(\sqrt{n}\lambda)$, the security of our bootstrapping scheme is reduced to $\text{GapSVP}_{\tilde{O}(n^{1.5}\lambda)}$ and $\text{SIVP}_{\tilde{O}(n^{1.5}\lambda)}$.

In the case of reducing to the classical hardness of the lattice problem, since $1/\alpha = \tilde{\Omega}(\lambda\sqrt{n}\log Q)$ and we must take $Q \approx 2^{n/2}$, the LWE inverse error rate satisfies $1/\alpha = \tilde{\Omega}(\lambda \cdot n)$. Therefore, the security of our optimized bootstrapping scheme is reduced to the classical hardness of $\text{GapSVP}_{\tilde{O}(n^2\lambda)}$.

Since all known algorithms that approximate GapSVP and SIVP on any n dimensional lattices to within a $\text{poly}(n)$ -factor run in time $2^{\Omega(n)}$, the 2^λ hardness requires us to choose $n = \Theta(\lambda)$. This makes the problems to which the security is reduced in the quantum case have the approximation factor $\tilde{O}(n^{2.5})$, which is smaller than $\tilde{O}(n^3)$, the one of [1]’s bootstrapping scheme. In the classical case, the LWE inverse error rate is $1/\alpha = \tilde{O}(n^2)$ and hence our approximation factor is $\tilde{O}(n^3)$. Furthermore, by selecting a larger dimension $n = \lambda^{1/\epsilon}$ for $\epsilon > 0$ (so at the cost of efficiency), the approximation factor can be $\tilde{O}(n^{1.5+\epsilon})$, which is comparable to the one of [9] and so the best known factor of standard lattice-based PKE. Consequently, our optimized bootstrapping scheme can be as secure as any other standard lattice-based PKE *without* successive dimension-modulus reduction, which is essential in all the known bootstrapping procedures [1, 9] provided recently.

Acknowledgments. We thank anonymous PKC reviewers for their helpful comments.

References

1. Alperin-Sheriff, J., Peikert, C.: Faster bootstrapping with polynomial error. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 297–314. Springer, Heidelberg (2014)
2. Barrington, D.A.M.: Bounded-Width polynomial-size branching programs recognize exactly those languages in NC^1 . In: STOC, pp. 1–5 (1986)
3. Barak, B.: Cryptography course - Lecture Notes, COS 433. Princeton University, Computer Science Department (2010). <http://www.cs.princeton.edu/courses/archive/spring10/cos433>
4. Brakerski, Z., Gentry, C., Halevi, S.: Packed ciphertexts in LWE-based homomorphic encryption. In: Kurosawa, K., Hanaoka, G. (eds.) PKC 2013. LNCS, vol. 7778, pp. 1–13. Springer, Heidelberg (2013)
5. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. In: ITCS, pp. 309–325 (2012)
6. 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)
7. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: FOCS, pp. 97–106 (2011)
8. 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)
9. Brakerski, Z., Vaikuntanathan, V.: Lattice-based FHE as secure as PKE. In: ITCS, pp. 1–12 (2014)
10. Coron, J.-S., Lepoint, T., Tibouchi, M.: Practical multilinear maps over the integers. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 476–493. Springer, Heidelberg (2013)
11. 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)

12. Gentry, C.: A Fully Homomorphic Encryption Scheme. PhD thesis, Stanford University (2009). <http://crypto.stanford.edu/craig>
13. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC, pp. 169–178 (2009)
14. Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 1–17. Springer, Heidelberg (2013)
15. Gentry, C., Gorbunov, S., Halevi, S.: Graph-Induced Multilinear Maps from Lattices. IACR Cryptology ePrint Archive **2014**, 645 (2014)
16. Gentry, C., Halevi, S., Smart, N.P.: Better bootstrapping in fully homomorphic encryption. In: Fischlin, M., Buchmann, J., Manulis, M. (eds.) PKC 2012. LNCS, vol. 7293, pp. 1–16. Springer, Heidelberg (2012)
17. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 75–92. Springer, Heidelberg (2013)
18. 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)
19. Micciancio, D., Mol, P.: Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 465–484. Springer, Heidelberg (2011)
20. Micciancio, D., Peikert, C.: Trapdoors for lattices: simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer, Heidelberg (2012)
21. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem. In: STOC, pp. 333–342 (2009)
22. Peikert, C., Vaikuntanathan, V., Waters, B.: A framework for efficient and composable oblivious transfer. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 554–571. Springer, Heidelberg (2008)
23. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC, pp. 84–93 (2005)
24. Rothblum, R.: Homomorphic encryption: from private-key to public-key. In: Ishai, Y. (ed.) TCC 2011. LNCS, vol. 6597, pp. 219–234. Springer, Heidelberg (2011)
25. 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)
26. Vershynin, R.: Introduction to the non-asymptotic analysis of random matrices. In: Eldar, Y.C., Kutyniok, G. (eds.) Compressed Sensing, Theory and Applications, ch. 5, pp. 210–268. Cambridge University Press (2012). <http://www-personal.umich.edu/romanv/papers/non-asymptotic-rmt-plain.pdf>