

MatriGear: Accelerating Authenticated Matrix Triple Generation with Scalable Prime Fields via Optimized HE Packing

Hyunho Cha
Seoul National University
aiden132435@cml.snu.ac.kr

Intak Hwang
Seoul National University
intak.hwang@snu.ac.kr

Seonhong Min
Seoul National University
minsh@snu.ac.kr

Jinyeong Seo
Seoul National University
jinyeong.seo@snu.ac.kr

Yongsoo Song
Seoul National University
y.song@snu.ac.kr

Abstract—The SPDZ protocol family is a popular choice for secure multi-party computation (MPC) in a dishonest majority setting with active adversaries. Over the past decade, a series of studies have focused on improving its offline phase, where special additive shares called authenticated triples are generated. However, to accommodate recent demands for matrix operations in secure machine learning and big integer arithmetic in distributed RSA key generation, updates to the offline phase are required.

In this work, we propose a new protocol for the SPDZ offline phase, MatriGear, which improves upon the previous state-of-the-art construction, TopGear (Baum et al., SAC '19), and its variant for matrix triples (Chen et al., Asiacrypt '20). Our protocol aims to achieve a speedup in matrix triple generation and support for larger prime fields up to 4096 bits in size. To achieve this, we devise a variant of the BFV scheme and a new homomorphic matrix multiplication algorithm optimized for our purpose.

As a result, our protocol achieves about 3.6x speedup for generating scalar triples in a 1024-bit prime field and about 34x speedup for generating 128x128 matrix triples. In addition, we reduce the size of evaluation keys from 27.4 GB to 0.22 GB and the communication cost for MAC key generation from 816 MB to 16.6 MB.

1. Introduction

Secure Multi-Party Computation (MPC) enables a set of parties to jointly compute a function over their private inputs without revealing any information other than the result. Constructing MPC protocols has been a major research topic in the field of cryptography, and there has been a number of solutions optimized for various settings, such as assuming an honest or dishonest majority among the parties, and whether the corrupted parties behave passively or actively.

In a dishonest majority setting with an active adversary, the SPDZ protocol family [1, 2, 3, 4] has been a popular choice, as it provides active security against up to $n - 1$ corrupted parties. The SPDZ protocol is divided into two

phases: an online phase and an offline phase. In the online phase, parties jointly compute a target function based on additive shares of their private inputs, which are elements of some prime field \mathbb{Z}_p . Since the protocol assumes active security, the private inputs are distributed in a special form called authenticated shares to prevent deviation from the protocol. At the end of the online phase, the parties jointly verify the output by checking the validity of the authenticated shares to ensure the protocol was correctly executed.

In the offline phase, authenticated shares are generated in addition to authenticated triples, which are required for performing multiplication in the online phase. Multiplications in the online phase are processed using Beaver's trick [5], which requires additive shares of a , b , and c such that $ab = c \pmod{p}$. The authenticated shares of such triples (a, b, c) are called authenticated triples. The core idea of the initial SPDZ protocol [1] is to generate authenticated triples using homomorphic encryption (HE). The parties first jointly generate HE ciphertexts ct_a and ct_b for a and b , respectively. They then compute the encryption ct_c of c through the homomorphic multiplication of ct_a and ct_b . Finally, they collectively decrypt ct_c to obtain the additive share of c .

To achieve active security in the offline phase, it is necessary to validate the HE ciphertexts ct_a and ct_b to prevent the use of ill-formed ciphertexts that could be generated by malicious adversaries. This is done using a proof of plaintext knowledge (PoPK) protocol, which is a Schnorr-like zero-knowledge protocol that verifies the well-formedness of the HE ciphertexts. Since the PoPK protocol has been a main performance bottleneck in the offline phase, subsequent studies [2, 3] on the SPDZ protocol have focused on improving the efficiency of the PoPK protocol.

Recently, MPC protocols have been used in a wide range of applications. One important application is privacy-preserving machine learning [6, 7, 8, 9], which enables distrustful parties to perform machine learning tasks on their private data. Since machine learning tasks heavily utilize matrix operations, recent improvements [4, 10, 11] to the SPDZ protocol have introduced the concept of authenticated

matrix triples, which significantly reduce the overhead of matrix operations in both the online and offline phases. On the other hand, MPC protocols can also be applied to threshold cryptography, which has recently gained attention with NIST’s efforts toward standardizing threshold cryptographic schemes [12]. For example, several studies [13, 14] have investigated distributed RSA key generation, which requires arithmetic MPC over large moduli to multiply two large prime numbers in a distributed manner. Additionally, other studies [15, 16] propose MPC-friendly symmetric primitives based on the Legendre PRF, which similarly require multi-party computations over large prime fields.

1.1. Our Contribution

In this paper, we propose a new offline phase for SPDZ, which we denote as MatriGear, improving upon the previous state-of-the-art construction [3], TopGear, to address recent demands for matrix multiplication and large modulus fields.

Overview on TopGear. We begin by reviewing the SPDZ offline phase, as introduced in the studies most closely related to this work—TopGear [3] and its matrix triple variant [4]. The protocol consists of four main steps: ciphertext generation, homomorphic computation, proof generation, and distributed decryption.

- In the ciphertext generation step, each party samples and encrypts random matrices \mathbf{A}_i and \mathbf{B}_i using the BFV [17, 18] or BGV [19] homomorphic encryption schemes, which support integer arithmetic. Then, they aggregate ciphertexts and compute encryptions of $\mathbf{A} = \sum_i \mathbf{A}_i$ and $\mathbf{B} = \sum_i \mathbf{B}_i$.
- In the homomorphic computation step, the parties perform homomorphic matrix multiplication between encryptions of \mathbf{A} and \mathbf{B} to obtain an encryption of $\mathbf{C} = \mathbf{AB}$.
- In addition, the parties execute a proof of plaintext knowledge (PoPK) protocol to validate the well-formedness of the encryptions of \mathbf{A} and \mathbf{B} , thus preventing malicious behavior by the parties.
- Finally, the parties execute the distributed decryption protocol, and obtain additive shares \mathbf{C}_i of \mathbf{C} .

This work aims to improve the performance of ciphertext generation, homomorphic computation, and proof generation, which dominate both the communication and computation costs in the TopGear framework.

Generalized BFV. During ciphertext generation, we require HE schemes that support large prime fields, as authenticated triples in SPDZ are defined over such fields. However, the BGV and BFV schemes are not optimized for this setting. To address this, we devise a new variant of the BFV homomorphic encryption scheme, which we denote as generalized BFV. Our generalized BFV scheme can efficiently handle messages from large prime fields by employing a novel packing method for large prime fields, following the idea of [20].

Let p be the modulus of the message space and q be the modulus of the ciphertext space. In the ordinary BFV

scheme [17, 18], q must be set much larger than p , as the available depth for homomorphic multiplication is roughly determined by the factor $\log q / \log p$. However, both the computational complexity of homomorphic multiplication and the space complexity of the evaluation key scale approximately with $O(\log^2 q)$. Because of this, performance significantly degrades when handling messages from large moduli because of the huge increase in the ciphertext modulus q .

Our packing method effectively addresses this issue by reducing the overhead of homomorphic multiplication from $\log p$ to $\log b$, when $p = b^M + 1$. Then, the available multiplicative depth is now determined by the factor $\log q / \log b$, which allows for a much smaller q for the same field size p . As a result, we can use a ciphertext modulus with a bit length ten times smaller than the ordinary BFV scheme when p is a 1024-bit prime. This results in approximately a threefold speedup in generating authenticated triples, and a reduction in the evaluation key size by about 125 times.

Faster Homomorphic Matrix Multiplication. For the homomorphic computation step, previous work [4] uses the homomorphic matrix multiplication algorithm from [21]. To further improve performance, we develop a new homomorphic matrix multiplication algorithm that achieves asymptotic improvements in computational complexity.

For $d \times d$ matrix multiplication, the previous algorithm in [21] required $\tilde{O}(d^3 \cdot \log^2 q)$ complexity and $3d - 2$ evaluation keys. In this work, we reduce the computational complexity to $\tilde{O}(d^3 \cdot \log q)$ and the number of evaluation keys to d by introducing a new matrix encoding map called *shifted diagonal encoding*. This improvement works particularly well with our generalized BFV scheme, which effectively reduces the size of $\log q$. As a result, we achieve approximately a 30-fold speedup in computing 128×128 matrix multiplication for a field size $p \approx 2^{128}$.

PoPK for Generalized BFV. For the proof generation step, an efficient PoPK protocol optimized for the generalized BFV scheme is required. In previous PoPK protocols, there exists an exponential gap between the honest statement and the proven statement, known as soundness slack. Since soundness slack affects parameter settings, particularly the ciphertext modulus q , minimizing it is crucial for our protocol.

To address this, we design an efficient PoPK protocol for generalized BFV that achieves constant-sized soundness slack. In particular, we incorporate recent advancements in proof techniques by Kim et al. [22] and Hwang et al. [23]. Kim et al. [22] proposed a method for achieving constant-sized soundness slack for a special case of the BFV scheme. However, their method does not directly apply to our setting due to structural modifications in our generalized BFV ciphertexts. We observe that this issue can be resolved using the randomized packing method proposed by Hwang et al. [23]. By merging these two techniques, we achieve constant-sized soundness slack, thereby maximizing the performance benefits from the generalized BFV scheme.

1.2. Related Works

TopGear. The most relevant works to ours are TopGear [3] and its variant for matrix triples [4]. TopGear is designed to support a large number of parties, so its performance scales well as the number of parties grows. The core idea of TopGear is to run a PoPK protocol where all parties simultaneously participate to produce a single joint proof that validates all parties' HE ciphertexts. Additionally, the generation of authenticated triples is done in a similar manner, utilizing homomorphic multiplication. This contrasts with other SPDZ offline phases [1, 2], where proofs and authenticated triples are generated pairwise between parties.

LowGear. There is another line of work for the SPDZ offline phase, called LowGear [2, 11]. In the LowGear protocol, each party pairwise generates PoPK proofs and authenticated triples. This approach has advantages with a small number of parties, as it can replace costly homomorphic multiplications with cheaper linear homomorphic operations. However, its performance degrades rapidly as the number of parties increases. We expect that our generalized BFV scheme could also be applied to the LowGear protocol and its variants.

Pseudorandom Correlation Generator. Aside from homomorphic encryption, authenticated triples can be generated using the pseudorandom correlation generator [24], which enables the generation of correlated randomness with sub-linear communication costs. Abram and Scholl [25] proposed a protocol for authenticated triple generation based on the pseudorandom correlation generator. Their protocol requires a total communication cost of $O(n^4\sqrt{T})$ to generate T triples, compared to $O(n^2T)$ in the TopGear and LowGear protocols, where n is the number of parties. However, the generation of authenticated matrix triples is not addressed in the literature.

Matrix Multiplication. There has been numerous research on secure matrix multiplication using HE. Halevi and Shoup [26] introduced an efficient algorithm for homomorphic linear transformation of an encrypted vector, which can be trivially extended to matrix multiplication. On the other hand, [21, 27] studied the case where the entire matrix is packed into a single ciphertext, optimized for smaller dimensional matrix multiplication. More recently, Zheng et al. [28] showed that trace maps can be used to accelerate the matrix multiplication if the matrix dimension is even smaller. However, their method requires arithmetic over non-power-of-two cyclotomic rings, which is incompatible with our case.

2. Background

2.1. Notation

Let N be a power of two and q be an integer. We denote by $K = \mathbb{Q}[X]/(X^N + 1)$ the $(2N)$ -th cyclotomic

field, $R = \mathbb{Z}[X]/(X^N + 1)$ the ring of integers of K , and $R_q = R/qR$ the quotient ring of R . We identify an element $a = \sum_{0 \leq i < N} a_i \cdot X^i \in R_q$ with the vector of its coefficients $(a_0, \dots, a_{N-1}) \in \mathbb{Z}_q^N$. We use $\mathbb{Z} \cap (-q/2, q/2]$ as a representative of \mathbb{Z}_q for an integer q , and denote by $[a]_q$ the reduction of an integer a modulo q . For $a \in R$, we define $\|a\|_\infty$ as the ℓ^∞ -norm of its coefficient vector, and the canonical norm $\|a\|_\infty^{\text{can}}$ is defined as $\|(a(\zeta), a(\zeta^3), \dots, a(\zeta^{2N-1}))\|_\infty$, where ζ is a complex primitive $2N$ -th root of unity. For a real number r , $\lceil r \rceil$ denotes the nearest integer to r , rounding upwards in case of a tie. For a finite set S , we use $x \leftarrow \mathcal{U}(S)$ to denote the sampling x according to the uniform distribution over S .

2.2. Authenticated Shares in SPDZ

We review the definitions of authenticated share and authenticated triple used in the online phase of the SPDZ protocol [1]. Let n be the number of parties involved in the multi-party computation protocol, and let α be the global MAC key such that the i -th party holds an additive share α_i , where $\sum_{i=0}^{n-1} \alpha_i = \alpha \pmod{p}$. For an input $x \in \mathbb{Z}_p$, an authenticated share $\llbracket x \rrbracket$ is defined as $(x_i, m_i)_{0 \leq i < n}$, where $\sum_{i=0}^{n-1} x_i = x \pmod{p}$ and $\sum_{i=0}^{n-1} m_i = \alpha \cdot x \pmod{p}$. A triple $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$ is called an authenticated triple if $ab = c \pmod{p}$ holds. For a matrix $\mathbf{X} \in \mathbb{Z}_p^{d \times d}$, an authenticated share $\llbracket \mathbf{X} \rrbracket$ is defined as $(\mathbf{X}_i, \mathbf{M}_i)_{0 \leq i < n}$, where $\sum_{i=0}^{n-1} \mathbf{X}_i = \mathbf{X} \pmod{p}$ and $\sum_{i=0}^{n-1} \mathbf{M}_i = \alpha \cdot \mathbf{M} \pmod{p}$. A triple $(\llbracket \mathbf{A} \rrbracket, \llbracket \mathbf{B} \rrbracket, \llbracket \mathbf{C} \rrbracket)$ is called an authenticated matrix triple if $\mathbf{AB} = \mathbf{C} \pmod{p}$ holds.

We note that the security of the MAC is determined by the size of p . Therefore, large prime numbers, such as 128-bit primes, are employed in SPDZ-like protocols. In addition to MAC security, certain MPC applications intrinsically require large prime fields due to their underlying cryptographic assumptions. Examples include distributed RSA key generation [13, 14] and multi-party computation of the Legendre PRF [15, 16].

2.3. Discrete Gaussian Distribution

We define the n -dimensional spherical Gaussian function $\rho : \mathbb{R}^n \rightarrow (0, 1]$ as $\rho(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^\top \mathbf{x})$. In general, for a positive definite matrix $\Sigma \in \mathbb{R}^{n \times n}$, we define the elliptical Gaussian function $\rho_{\sqrt{\Sigma}} : \mathbb{R}^n \rightarrow (0, 1]$ as $\rho_{\sqrt{\Sigma}}(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^\top \Sigma^{-1} \mathbf{x})$. Let $\Lambda \subseteq \mathbb{R}^n$ be a lattice and $\mathbf{c} \in \mathbb{R}^n$. The discrete Gaussian distribution $\mathcal{D}_{\mathbf{c}+\Lambda, \sqrt{\Sigma}}$ is defined as a distribution over the coset $\mathbf{c} + \Lambda$, whose probability mass function is $\mathcal{D}_{\mathbf{c}+\Lambda, \sqrt{\Sigma}}(\mathbf{x}) = \rho_{\sqrt{\Sigma}}(\mathbf{x}) / \rho_{\sqrt{\Sigma}}(\mathbf{c} + \Lambda)$ for $\mathbf{x} \in \mathbf{c} + \Lambda$ where $\rho_{\sqrt{\Sigma}}(\mathbf{c} + \Lambda) := \sum_{\mathbf{v} \in \mathbf{c} + \Lambda} \rho_{\sqrt{\Sigma}}(\mathbf{v}) < \infty$. When $\Sigma = \sigma^2 \cdot \mathbf{I}_n$ for $\sigma > 0$ where \mathbf{I}_n is the n -dimensional identity matrix, then we substitute $\sqrt{\Sigma}$ by σ in the subscript and refer to σ as the width parameter. For a polynomial $a \in R$, we denote by $a \leftarrow \mathcal{D}_{\mathbf{c}+\Lambda, \sqrt{\Sigma}}$ if we sample its coefficient vector from $\mathcal{D}_{\mathbf{c}+\Lambda, \sqrt{\Sigma}}$. Below, we provide useful lemmas for tail bounds of discrete Gaussian distributions.

Definition 1 ([29, Definition 3.1]). For an n -dimensional lattice Λ and positive real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest s such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$.

Lemma 1 ([29, Lemma 3.3]). For any n -dimensional lattice Λ and $\varepsilon > 0$,

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(\Lambda)$$

where $\lambda_n(\Lambda)$ is the smallest real number $r > 0$ such that $\dim(\text{span}(\Lambda \cap \mathcal{B}(r))) = n$ and $\mathcal{B}(r)$ is the n -dimensional ball with radius r centered at the origin.

Lemma 2 ([30, Lemma 2.4]). Let $\Lambda \subseteq \mathbb{R}^n$ be any n -dimensional lattice. For any $0 < \varepsilon < 1/3$, $\sigma \geq \eta_\varepsilon(\Lambda)$, and any $\mathbf{c} \in \mathbb{R}^n$,

$$\Pr[\|\mathbf{x}\|_\infty > 6\sigma \mid \mathbf{x} \leftarrow \mathcal{D}_{\mathbf{c}+\Lambda, \sigma}] \leq n \cdot 2^{-161}$$

2.4. Ring Learning With Errors

Let χ and ψ be distributions over R . The ring learning with errors (RLWE) problem [31], denoted as $\text{RLWE}_{R, q, \chi, \psi}$, is to distinguish between the following two distributions.

$$\begin{aligned} &\{(a, b) \mid a \leftarrow \mathcal{U}(R_q), b \leftarrow \mathcal{U}(R_q)\} \\ &\{(a, b) \mid a \leftarrow \mathcal{U}(R_q), s \leftarrow \chi, e \leftarrow \psi, b = -a \cdot s + e \pmod{q}\} \end{aligned}$$

We refer to a pair $(c_0, c_1) \in R_q^2$ as an RLWE encryption of μ under s if $c_0 = -c_1 s + \mu + e \pmod{q}$ holds for a small error term $e \in R$. Then, it is decryptable by computing $c_0 + c_1 s \approx \mu \pmod{q}$.

2.5. Gadget Decomposition

Below, we provide the definition of gadget decomposition and external product, which are key building blocks for implementing nonlinear operations in RLWE-based homomorphic encryption schemes

Definition 2 (Gadget Decomposition). For a modulus q , a function $h : R_q \rightarrow R^L$ is called a gadget decomposition if there exists a fixed vector $\mathbf{g} = (g_0, g_1, \dots, g_{L-1}) \in R_q^d$ and a real $B > 0$ such that the followings hold for any $a \in R_q$ and its decomposition $h(a) = (b_0, b_1, \dots, b_{L-1})$.

$$\langle h(a), \mathbf{g} \rangle = \sum_{i=0}^{L-1} b_i \cdot g_i = a \pmod{q} \quad \text{and} \quad \|\mathbf{b}\|_\infty \leq B$$

We call \mathbf{g} a gadget vector and $B > 0$ a bound of h .

Definition 3 (External Product). For $a \in R_q$ and $\mathbf{u} = (u_0, \dots, u_{d-1}) \in R_q^L$, a binary operation external product $\square : R_q \times R_q^L \rightarrow R_q$ is defined as follows.

$$a \square \mathbf{u} := \langle h(a), \mathbf{u} \rangle \pmod{q}$$

By an abuse of notation, we write $a \square \mathbf{U} = (a \square \mathbf{u}_0, a \square \mathbf{u}_1) \in R_q^2$ for any $a \in R_q$ and $\mathbf{U} = [\mathbf{u}_0 \mid \mathbf{u}_1] \in R_q^{L \times 2}$.

The key-switching procedure is a major use case of the external product, which is used for implementing nonlinear

HE operations. Given an RLWE ciphertext $(c_0, c_1) \in R^2$, which is an encryption of μ under the secret key s , we can generate another RLWE ciphertext (c'_0, c'_1) with a different secret key s' such that $c_0 + c_1 s \approx c'_0 + c'_1 s' \pmod{q}$. This is achieved by computing $(c'_0, c'_1) = (c_0, 0) + c_1 \square \mathbf{U} \pmod{q}$, where $\mathbf{u}_0 = -s \cdot \mathbf{u}_1 + s' \cdot \mathbf{g} + \mathbf{e} \pmod{q}$ for some small error term \mathbf{e} . Then, (c'_0, c'_1) can be interpreted as the RLWE encryption of μ under the new secret key s' , switching the secret key from s to s' .

In a real-world implementation, a popular choice for a gadget decomposition is the RNS (Residue Number System) gadget decomposition [32]. Suppose that $q = \prod_{0 \leq i < L} q_i$ for some word-sized primes q_i 's, where $0 \leq i < L$. Then, a gadget decomposition and a gadget vector are defined as $h(a) = ([a]_{q_0}, \dots, [a]_{q_{L-1}})$ and $\mathbf{g} = (\tilde{q}_0, \dots, \tilde{q}_{L-1})$, where $\tilde{q}_i = q/q_i \cdot [(q/q_i)^{-1}]_{q_i}$. It follows that $\langle h(a), \mathbf{g} \rangle = a \pmod{q}$ and $\|h(a)\|_\infty < \max\{q_0, \dots, q_{L-1}\}$, so h becomes a gadget decomposition over R_q . The computation of the external product involves two steps: decomposition and inner product. The decomposition step computes $h(a)$, and the inner product step computes $\langle h(a), \mathbf{u} \rangle$. The complexity of the decomposition step is $\tilde{O}(LN)$ or $\tilde{O}(L^2 N)$ depending on the algorithm used, while the inner product computation requires $\tilde{O}(L^2 N)$. Thus, the overall complexity is $\tilde{O}(L^2 N)$. For a detailed analysis, we refer to [22].

3. Generalized BFV for Large Primes

In this section, we present our variant of the BFV scheme [17, 18] that can efficiently handle large prime plaintext moduli. For the ordinary BFV scheme, the plaintext space is defined as R_p , so the plaintext modulus p is usually set to be much smaller than the ciphertext modulus q , as the available multiplicative depth is roughly determined by the ratio $\log(q)/\log(p)$. Conversely, using large primes for plaintext moduli results in an increase in the ciphertext modulus q , but as noted in Section 2.5, this leads to a quadratic performance overhead, as the gadget level L scales with $O(\log q)$.

To overcome this problem, we devise a new variant of the BFV scheme that can accommodate messages from a large prime field, following the idea of the previous work by Bootland et al. [20]. The core idea is to design a variant of the BFV scheme with a plaintext space of $R_{X^D-b} = \mathbb{Z}[X]/(X^N + 1, X^D - b)$ for some $D \mid N$ and an integer b , instead of $R_p = \mathbb{Z}[X]/(X^N + 1, p)$ in the ordinary BFV scheme. Since $\mathbb{Z}[X]/(X^N + 1, X^D - b) = \mathbb{Z}[X]/(b^M + 1, X^D - b)$, where $M = N/D$, if we set $p = b^M + 1$, it follows that $R_{X^D-b} = \mathbb{Z}_p[X]/(X^D - b)$. This allows us to perform modular arithmetic over p using the plaintext space R_{X^D-b} . The advantage of this plaintext space is that the multiplicative depth is roughly determined by the ratio $\log q / \log b$, rather than $\log q / \log p$. Thus, by choosing a suitable b , we can construct an efficient HE scheme that can handle large prime plaintext moduli. For example, in our experiments, we set $b \approx 2^{16}$ for a prime $p \approx 2^{4096}$ with $M = 256$, which allows us to use a much smaller q compared to the ordinary BFV scheme.

We note that our variant of the BFV scheme can be viewed as a generalization of the ordinary BFV scheme, as the latter corresponds to the case when $M = 1$. Hence, we refer to it as the generalized BFV scheme. In the rest of this section, we provide more details about our generalized BFV scheme, including how we instantiate a SIMD-like packing method and investigate the automorphism group on R_{X^D-b} for implementing homomorphic rotation operations.

3.1. Packing Method

Although the previous work [20] utilized the plaintext space R_{X^D-b} , its packing method was designed for handling complex-valued messages. Thus, our first goal is to instantiate a HELib-like packing method [33] over $R_{X^D-b} = \mathbb{Z}_p[X]/(X^D-b)$ to support messages from large prime fields rather than complex numbers. More precisely, we aim to pack a vector $\mathbf{m} \in \mathbb{Z}_p^D$ into a single polynomial $\mu \in R_{X^D-b}$ using the Chinese Remainder Theorem (CRT), so that we can perform homomorphic operations over \mathbb{Z}_p in a SIMD manner. To achieve this, the polynomial X^D-b must fully split modulo p , so we investigate the conditions on p , D , and b required for X^D-b to split modulo p below. We note that all proofs in this subsection are deferred to Section B.

Lemma 3. *Let b , D , and M be positive integers such that $N = M \cdot D$, and $p = b^M + 1$ is a prime. If $2N$ divides $(p-1)$, then X^D-b splits in $\mathbb{Z}_p[X]$. Moreover, there exists an element $\xi \in \mathbb{Z}_p$ such that $\xi, \xi^{2M+1}, \dots, \xi^{2(D-1)M+1}$ are distinct roots for $X^D-b = 0 \pmod{p}$ for $0 \leq i < D$.*

From the above lemma, we obtain the following isomorphism relations due to the CRT.

$$\mathbb{Z}_p[X]/(X^D-b) \cong \prod_{i=0}^{D-1} \mathbb{Z}_p[X]/(X-\xi^{2iM+1}) \cong \mathbb{Z}_p^D$$

Now, a packing strategy for a vector in \mathbb{Z}_p^D into an element in R_{X^D-b} can be devised from the above relations. Briefly speaking, given a vector $\mathbf{m} \in \mathbb{Z}_p^D$, a degree $D-1$ polynomial $\mu \in R_{X^D-b}$ is computed, such that each evaluation value of μ at the roots $\{\xi^{2iM+1}\}_{0 \leq i < D}$ is each entry of the vector. Then, the homomorphic operations such as addition and multiplication over R_{X^D-b} directly translates into the component-wise arithmetic in \mathbb{Z}_p^D , so we refer to D as the number of slots. Finding such a polynomial μ can be efficiently done by adapting the Number Theoretic Transform (NTT) algorithm.

For the purpose of homomorphic matrix multiplication, we need to introduce an additional homomorphic operation over R_{X^D-b} called *rotation*. The rotation operation ρ on a vector $\mathbf{m} = (m_0, \dots, m_{D-1})$ is defined as $\rho(\mathbf{m}) := (m_1, \dots, m_{D-1}, m_0)$, which cyclically shifts the entries of the input vector. In the following, we demonstrate how to instantiate the rotation operation using an automorphism evaluation over the plaintext space R_{X^D-b} , following the strategy outlined in [33]. First, we investigate when a map $\theta_i : X \mapsto X^i$ becomes an automorphism over R_{X^D-b} .

For θ_i to be an automorphism, it should map the roots $\{\xi^{2jM+1}\}_{0 \leq j < D}$ of X^D-b modulo p to themselves. Applying θ_i to this set yields $\{\xi^{i(2jM+1)}\}_{0 \leq j < D}$. To ensure that $\{\xi^{i(2jM+1)}\}_{0 \leq j < D} = \{\xi^{2jM+1}\}_{0 \leq j < D}$, each $i(2jM+1)$ must be of the form $2jM+1$, which implies $i = 1 \pmod{2M}$. Furthermore, if $i = 1 \pmod{2M}$, the elements $\xi^{i(2jM+1)}$ will not overlap, making θ_i an automorphism. Below, we show that the set of automorphisms $\{X \mapsto X^{2iM+1} \mid 0 \leq i < D\}$ indeed forms a cyclic group.

Lemma 4. *Let $D, M \geq 2$ be power-of-2 integers such that $N/M = D$. Then, $G = \{2iM+1 \mid 0 \leq i < d\}$ is a cyclic multiplicative subgroup of \mathbb{Z}_{2N}^\times and $5^{M/2}$ is its generator.*

From the above lemma, the set $\{X \mapsto X^{2iM+1} \mid 0 \leq i < D\}$ forms a cyclic group generated by the automorphism $\varphi : X \mapsto X^{5^{M/2}}$. We can now describe our packing method with cyclic rotation as follows.

Theorem 1. *Let b , D , and M be positive integers such that $D = N/M$ and $p = b^M + 1$ is a prime. If $2N$ divides $p-1$, then X^D-b splits modulo p , and $R_{X^D-b} = \mathbb{Z}_p[X]/(X^D-b) \simeq \mathbb{Z}_p^D$. Moreover, there exists $\xi \in \mathbb{Z}_p$ such that the mapping $\tau : f \mapsto (f(\xi), f(\xi^{5^{M/2}}), \dots, f(\xi^{5^{(D-1)M/2}}))$, which maps $f \in \mathbb{Z}_p[X]/(X^D-b)$ to an element in \mathbb{Z}_p^D , is an isomorphism.*

This theorem implies that if we pack a vector in \mathbb{Z}_p^D using the map τ^{-1} , we can perform rotation ρ by the automorphism $\varphi : X \mapsto X^{5^{M/2}}$. To be precise, if a message $\mathbf{m} = (m_0, m_1, \dots, m_{D-1})$ is packed into a plaintext $\tau^{-1}(\mathbf{m}) = \mu(X)$, then $\varphi(\mu) = \mu(X^{5^{M/2}})$ corresponds to the rotated message $\rho(\mathbf{m}) = (m_1, \dots, m_{D-1}, m_0)$, as illustrated in Fig. 1.

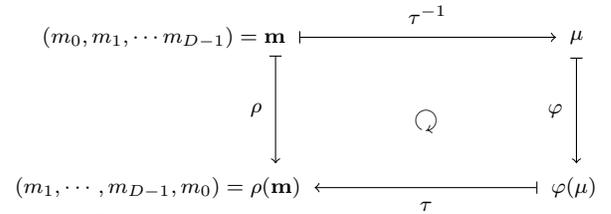


Figure 1: Rotation operation in R_{X^D-b}

We define the packing and unpacking algorithms for our scheme as follows. We note that τ and τ^{-1} can be computed with a time complexity of $O(D \log D)$ by modifying the NTT algorithm.

- **Pack**(\mathbf{m}) $\rightarrow \mu$: Given a message $\mathbf{m} = (m_0, m_1, \dots, m_{D-1}) \in \mathbb{Z}_p^D$, return a plaintext $\mu = \tau^{-1}(\mathbf{m}) \in R_{X^D-b}$.
- **Unpack**(μ) $\rightarrow \mathbf{m}$: Given a plaintext $\mu \in R_{X^D-b}$, return a message $\mathbf{m} = \tau(\mu) \in \mathbb{Z}_p^D$.

3.2. Scheme Description

We present our generalized BFV scheme that supports efficient arithmetic over large primes, built upon our new

packing method. Our construction basically follows the construction of [20].

Setup. The algorithms for setup are given as follows:

- **Setup**(1^λ) \rightarrow pp: Given a security parameter λ , choose the ring dimension N , the ciphertext modulus q , the secret key distribution χ , and the error distribution ψ . Choose a plaintext space R_{X^D-b} , and define $\Delta = \left\lfloor \frac{q}{X^D-b} \right\rfloor = \left\lfloor -\frac{q}{b^{M+1}}(X^{N-D} + bX^{N-2D} + \dots + b^{M-1}) \right\rfloor \in R$. Choose a gadget vector $\mathbf{g} \in R_q^L$ and a gadget decomposition $h : R_q \rightarrow R^L$, with bound B . Return public parameters pp = $(N, q, \chi, \psi, R_{X^D-b}, h, \mathbf{g}, B)$.
- **KeyGen**(pp) \rightarrow (sk, pk): Sample $s \leftarrow \chi$, $a \leftarrow \mathcal{U}(R_q)$ and $e \leftarrow \psi$, and set sk = $(1, s)$, and pk = $(-as + e, a) \in R_q^2$. Return a key pair (sk, pk).
- **RelinKeyGen**(sk) \rightarrow rlk: Given a secret key sk = $(1, s)$, sample $\mathbf{a} \leftarrow \mathcal{U}(R_q^L)$, and $\mathbf{e} \leftarrow \psi^L$. Return a relinearization key rlk = $(-s \cdot \mathbf{a} + s^2 \cdot \mathbf{g} + \mathbf{e}, \mathbf{a})$.
- **RotKeyGen**(sk, r) \rightarrow rtk $_r$: Given a secret key sk = $(1, s)$ and a rotation index r , sample $\mathbf{a} \leftarrow \mathcal{U}(R_q^L)$, and $\mathbf{e} \leftarrow \psi^L$. Return a rotation key rtk $_r$ = $(-s \cdot \mathbf{a} + \varphi^r(s) \cdot \mathbf{g} + \mathbf{e}, \mathbf{a})$.

While there is no need to generate the rotation key when the rotation index r is zero, rtk $_0$ will be used in the algorithm description for an easier notation. We remark that rotation with rtk $_0$ is essentially outputting the input ciphertext.

Encryption. To encrypt a plaintext $\mu \in R_{X^D-b}$, it needs to be represented as an element in R . There are multiple possible representations, but we want it to have a small norm to minimize noise growth during encryption. We observe that every $\mu \in R_{X^D-b}$ can have a representation $[\mu]_R \in R$ such that $\|[\mu]_R\|_\infty \leq b$, leveraging the fact that $X^D = b \pmod{X^D - b}$, and we use this representation for the encryption algorithm. The algorithms for encryption and decryption are given as follows.

- **Enc**(μ, \mathbf{r}) \rightarrow ct: Given a public key pk, a plaintext $\mu \in R$, and a randomness triple $\mathbf{r} = (r_0, r_1, r_2) \in R^3$, return ct = $r_0 \cdot \text{pk} + (\Delta \cdot \mu + r_1, r_2) \in R_q^2$.
- **Dec**(ct) \rightarrow μ : Given a ciphertext ct = (c_0, c_1) and a secret key sk = $(1, s)$, return a plaintext $\mu = \left\lfloor \frac{X^D-b}{q}(c_0 + c_1 \cdot s) \right\rfloor \pmod{X^D - b}$.

Homomorphic Operations. Below, we present algorithms for homomorphic operations. We incorporate a key-switching procedure for nonlinear homomorphic operations, such as multiplication and rotations.

- **Add**(ct, ct') \rightarrow ct'': Given ciphertexts ct, ct' $\in R_q^2$, return ct'' = ct + ct' (mod q).
- **PMul**(μ, ct) \rightarrow ct': Given a plaintext $\mu \in R_{X^D-b}$ and a ciphertext ct $\in R_q^2$, return ct' = $[\mu]_R \cdot \text{ct}$ (mod q).
- **Mul**(ct, ct') \rightarrow ct'': Given ciphertexts ct = (c_0, c_1) , ct' = (c'_0, c'_1) and a relinearization key rlk, let $d_0 = \left\lfloor \frac{X^D-b}{q} c_0 c'_0 \right\rfloor \pmod{q}$, $d_1 = \left\lfloor \frac{X^D-b}{q} (c_0 c'_1 + c_1 c'_0) \right\rfloor \pmod{q}$ and $d_2 = \left\lfloor \frac{X^D-b}{q} c_1 c'_1 \right\rfloor \pmod{q}$. Return ct'' = $(d_0, d_1) + d_2 \boxtimes \text{rlk}$ (mod q).

- **Rot**(ct, r) \rightarrow ct': Given a ciphertext ct = (c_0, c_1) , a rotation index r and the corresponding rotation key rtk $_r$, let $d_0 = \varphi^r(c_0)$, and $d_1 = \varphi^r(c_1)$. Return ct' = $(d_0, 0) + d_1 \boxtimes \text{rtk}_r \pmod{q}$.

Noise Analysis. For the noise analysis, we use the notion of invariant noise introduced in [20].

Definition 4 (Invariant Noise). *Let ct = (c_0, c_1) be a ciphertext with a secret key sk = $(1, s)$. For $\mu \in R$, we define the invariant noise $\text{Err}(\text{ct}, \mu)$ as an element $v \in K$ with the smallest canonical norm, which satisfies the following for some $I \in R$:*

$$\frac{X^D - b}{q}(c_0 + c_1 s) = \mu + v + (X^D - b)I$$

We prove the correctness of encryption and each homomorphic operation, along with the noise growth in Section B.4. We remark that the noise growth of homomorphic multiplication now depends on the size of b rather than p , as analyzed in Lemma 8.

Security Analysis. The secret, public, and evaluation key generation algorithms in the generalized BFV scheme are identical to those of the ordinary BFV scheme [17, 18]. Moreover, during encryption, both schemes generate random encryptions of zero in exactly the same manner; the only difference lies in their packing methods. Consequently, the generalized BFV scheme achieves semantic security under the hardness of the RLWE problem [31], along with an additional circular security assumption [34], similar to the ordinary BFV scheme.

4. Improved Matrix Multiplication

In this section, we present a new homomorphic square matrix multiplication algorithm that can be used for authenticated matrix triple generation. Our algorithm improves upon the previous approaches [26, 21] by introducing a novel matrix encoding method optimized for square matrix multiplication. Consequently, we reduce the overhead of key-switching in homomorphic matrix multiplication, thereby improving the performance by a factor of L .

Throughout this section, we use d to denote the matrix dimension, and the complexity of the algorithm is analyzed in terms of d . We begin by briefly reviewing the two previous methods [26] and [21] for homomorphic square matrix multiplication. We then demonstrate how our method improves performance over these approaches, together with a detailed analysis.

4.1. Previous Approach

Halevi-Shoup Method. In [26], Halevi and Shoup proposed a method to homomorphically compute a linear transformation on an encrypted vector. For a matrix $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}_p^{d \times d}$, it first defines a diagonal encoding map $\iota : \mathbb{Z}_p^{d \times d} \rightarrow (\mathbb{Z}_p^d)^d$ as

$$\iota : \mathbf{A} = (a_{i,j}) \mapsto ((a_{j,j+i})_{0 \leq j < d})_{0 \leq i < d}$$

where each index is computed modulo d . For simplicity, we denote $\iota(\mathbf{A}) = (\iota(\mathbf{A})_0, \dots, \iota(\mathbf{A})_{d-1})$ where each $\iota(\mathbf{A})_i$ is a vector in \mathbb{Z}_p^d . Then, for a vector $\mathbf{v} \in \mathbb{Z}_p^d$, we have

$$\mathbf{A}\mathbf{v} = \sum_{k=0}^{d-1} \iota(\mathbf{A})_k \odot \rho^k(\mathbf{v}) \quad (1)$$

where ρ^k denotes the left rotation of the elements of a vector by k , and \odot denotes the component-wise multiplication between vectors. To translate Eq. (1) into homomorphic operations, we need ciphertexts that encrypt $\iota(\mathbf{A})_0, \dots, \iota(\mathbf{A})_{d-1}$, \mathbf{v} , and perform homomorphic rotations $\text{Rot}(\cdot, k)$ on the encryption of \mathbf{v} to obtain an encryption of $\rho^k(\mathbf{v})$. Since each vector is encrypted into a single ciphertext, the number of slots is set to $D = d$. We also note that this method can be trivially extended to encrypted matrix-matrix multiplication by applying Eq. (1) d times, once for each column of the matrix.

Jiang et al. Method. In [21], Jiang et al. proposed an optimized algorithm for smaller-dimensional matrix multiplication, where each matrix is encrypted into a single ciphertext, so that the number of slots is $D = d^2$. For a matrix $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}_p^{d \times d}$, the flattening map $\sigma : \mathbb{Z}_p^{d \times d} \rightarrow \mathbb{Z}_p^{d^2}$ is defined as follows.

$$\sigma : \mathbf{A} = (a_{i,j}) \mapsto (a_{d \cdot i + j})_{0 \leq i, j < d}$$

In addition, it defines four $d^2 \times d^2$ permutation matrices $\mathbf{M}, \mathbf{N}, \mathbf{V}^{(k)}, \mathbf{W}^{(k)}$ with entries defined as follows.

$$\begin{aligned} \mathbf{M}_{d \cdot i + j, \ell} &= \begin{cases} 1 & \text{if } \ell = d \cdot i + j + [i + j]_d \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{N}_{d \cdot i + j, \ell} &= \begin{cases} 1 & \text{if } \ell = d \cdot [i + j]_d + j \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{V}_{d \cdot i + j, \ell}^{(k)} &= \begin{cases} 1 & \text{if } \ell = d \cdot i + j + [j + k]_d \\ 0 & \text{otherwise} \end{cases} \\ \mathbf{W}_{d \cdot i + j, \ell}^{(k)} &= \begin{cases} 1 & \text{if } \ell = d \cdot [i + k]_d + j \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Then, for any $d \times d$ matrices \mathbf{A}, \mathbf{B} , the following holds.

$$\sigma(\mathbf{A}\mathbf{B}) = \sum_{k=0}^{d-1} (\mathbf{V}^{(k)} \mathbf{M} \sigma(\mathbf{A})) \odot (\mathbf{W}^{(k)} \mathbf{N} \sigma(\mathbf{B})) \quad (2)$$

To translate Eq. (2) into homomorphic operations, we need ciphertexts that encrypt $\sigma(\mathbf{A})$ and $\sigma(\mathbf{B})$ and computes $\mathbf{V}^{(k)} \mathbf{M} \sigma(\mathbf{A}), \mathbf{W}^{(k)} \mathbf{N} \sigma(\mathbf{B})$, using the homomorphic linear transform algorithm by Halevi and Shoup.

4.2. Our Method

In both previous approaches, the main performance bottlenecks are the key-switching procedures in homomorphic rotations. Hence, we focus on reducing the overhead of key-switching by developing a new algorithm for homomorphic matrix multiplication. For this, we start by defining a new encoding map, called the *shifted diagonal encoding map*.

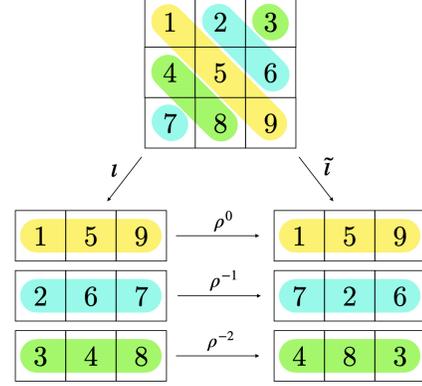


Figure 2: Comparison between diagonal encoding map ι and shifted diagonal encoding map $\tilde{\iota}$.

The shifted diagonal encoding map $\tilde{\iota} : \mathbb{Z}_p^{d \times d} \rightarrow (\mathbb{Z}_p^d)^d$ is defined as follows,

$$\tilde{\iota} : \mathbf{A} = (a_{i,j}) \mapsto ((a_{j-i,j})_{0 \leq j < d})_{0 \leq i < d}$$

or equivalently, $\tilde{\iota}_i(\mathbf{A}) = \rho^{-i}(\iota_i(\mathbf{A}))$, as illustrated in Fig. 2. We denote $\tilde{\iota}(\mathbf{A}) = (\tilde{\iota}(\mathbf{A})_0, \dots, \tilde{\iota}(\mathbf{A})_{d-1})$. Then, the multiplication between $\mathbf{A} = (a_{i,j})$ and $\mathbf{B} = (b_{i,j})$ can be written as

$$\iota(\mathbf{A} \cdot \mathbf{B})_i = \sum_{k=0}^{d-1} \iota(\mathbf{A})_k \odot \rho^i(\tilde{\iota}(\mathbf{B})_{i-k}) \quad (3)$$

since

$$\begin{aligned} \left(\sum_{k=0}^{d-1} \iota(\mathbf{A})_k \odot \rho^i(\tilde{\iota}(\mathbf{B})_{i-k}) \right)_j &= \sum_{k=0}^{d-1} \iota(\mathbf{A})_{k,j} \cdot \tilde{\iota}(\mathbf{B})_{i-k,j+i} \\ &= \sum_{k=0}^{d-1} a_{j,j+k} \cdot b_{j+k,j+i} \\ &= (\mathbf{A} \cdot \mathbf{B})_{j,j+i} = \iota(\mathbf{A} \cdot \mathbf{B})_{i,j}. \end{aligned}$$

The translation of Eq. (3) in terms of homomorphic operation is as follows. We assume that the number of slots is $D = d$, as in the Halevi-Shoup method. Thus, we encrypt $\text{Pack}(\iota(\mathbf{A})_i)$ into d ciphertexts, denoted as $\text{CT}_{\iota(\mathbf{A})} = (\text{ct}_{\iota(\mathbf{A})_0}, \dots, \text{ct}_{\iota(\mathbf{A})_{d-1}})$, and similarly encrypt $\text{Pack}(\tilde{\iota}(\mathbf{B})_i)$ into $\text{CT}_{\tilde{\iota}(\mathbf{B})} = (\text{ct}_{\tilde{\iota}(\mathbf{B})_0}, \dots, \text{ct}_{\tilde{\iota}(\mathbf{B})_{d-1}})$. A naive way of computing Eq. (3) is to first compute the encryption of $\rho^i(\iota(\mathbf{B})_{i-k})$ by applying $\text{Rot}(\cdot, i)$ on $\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-k}} = (b_{i-k,0}, b_{i-k,1})$, which is computed as follows.

$$\text{Rot}(\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-k}}, i) = (\varphi^i(b_{i-k,0}), 0) + \varphi^i(b_{i-k,1}) \square \text{rtk}_i$$

Then, we compute the sum of the product between $\text{ct}_{\iota(\mathbf{A})_k} = (a_{k,0}, a_{k,1})$ and $\text{Rot}(\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-k}}, i) = (b_{i-k,0}^{(i)}, b_{i-k,1}^{(i)})$ as follows.

$$\begin{aligned} \text{ct}_{\iota(\mathbf{A}\mathbf{B})_i} &= \sum_{k=0}^{d-1} \left(\left[\frac{X^D - b}{q} (a_{k,0} b_{i-k,0}^{(i)}, a_{k,0} b_{i-k,1}^{(i)} + b_{i-k,0}^{(i)} a_{k,1}) \right] \right. \\ &\quad \left. + \left[\frac{X^D - b}{q} a_{k,1} b_{i-k,1}^{(i)} \right] \square \text{rlk} \right) \pmod{q} \end{aligned}$$

This approach requires d^2 Mul operations and d^2 Rot op-

erations for each, resulting in a total of $2d^2$ key-switching operations.

Optimization via Lazy Key-switching. We then observe that the overhead of key-switching can be significantly reduced by employing the well-known optimization technique called *lazy key-switching*. To be precise, if we defer all the key-switching procedures for lazy key-switching, it can be rewritten as

$$(d'_0, d'_1, d'_2, d'_3) = \left[\frac{X^D - b}{q} \left(\sum_{k=0}^{d-1} \left(a_{k,0} \varphi^i(b_{i-k,0}), a_{k,1} \varphi^i(b_{i-k,0}), a_{k,0} \varphi^i(b_{i-k,1}), a_{k,1} \varphi^i(b_{i-k,1}) \right) \right) \right] \pmod{q}$$

so that the following holds.

$$d'_0 + d'_1 s + d'_2 \varphi^i(s) + d'_3 s \varphi^i(s) \approx \Delta \cdot \text{Pack}(\tilde{\iota}(\mathbf{AB})_i) \pmod{q}.$$

Then, we apply key-switching to the ciphertext (d'_0, d'_1, d'_2, d'_3) , changing the secret key from $(1, s, \varphi^i(s), s \varphi^i(s))$ to the original secret key $(1, s)$. More specifically, we first switch all occurrences of $\varphi^i(s)$ to s using the rotation key rtk_i , resulting in an overall secret key of $(1, s, s^2)$. Finally, we switch s^2 to s with the relinearization key rlk . In summary, we compute the following, where $(v_0, v_1) = d'_3 \boxtimes \text{rtk}_i$.

$$(d'_0, d'_1 + v_0) + v_1 \boxtimes \text{rlk} + d'_2 \boxtimes \text{rtk}_i \pmod{q}$$

Then the correctness can be sketched as follows.

$$\begin{aligned} & \langle (1, s), (d'_0, d'_1 + v_0) + v_1 \boxtimes \text{rlk} + d'_2 \boxtimes \text{rtk}_i \rangle \\ & \approx d'_0 + d'_1 \cdot s + s \cdot v_0 + s^2 \cdot v_1 + d'_2 \cdot \varphi^i(s) \\ & = d'_0 + d'_1 \cdot s + s \cdot (v_0 + s \cdot v_1) + d'_2 \cdot \varphi^i(s) \\ & \approx d'_0 + d'_1 \cdot s + d'_2 \cdot \varphi^i(s) + d'_3 \cdot s \varphi^i(s) \pmod{q} \end{aligned}$$

Consequently, we require 3 key-switching operations to homomorphically compute Eq. (3) for each $0 \leq i < d$. Therefore, a total of $3d$ key-switching operations are needed for the entire matrix multiplication, which is a reduction from the $2d^2$ key-switching operations required for naive computation. The full algorithm for our encrypted matrix multiplication is provided in Algorithm 1, which defines the following homomorphic operation.

- $\text{MatMul}(\text{CT}_{\iota(\mathbf{A})}, \text{CT}_{\tilde{\iota}(\mathbf{B})}) \rightarrow \text{CT}_{\iota(\mathbf{C})}$: Given ciphertexts $\text{CT}_{\iota(\mathbf{A})}$ and $\text{CT}_{\tilde{\iota}(\mathbf{B})}$, run Algorithm 1, and return the result $\text{CT}_{\iota(\mathbf{C})}$.

Noise Analysis. We analyze the noise growth of MatMul in Section B.5.

Complexity Analysis. Now, we provide a detailed complexity analysis of our method, together with a comparison to the previous method. For a fair comparison, we first explain the various optimization that can be made to previous approaches explained in Section 4.1. We first recall that a key-switching operation involves two steps: decomposition and inner product, as explained in Section 2.5.

In the method by Halevi and Shoup, it requires d^2 Rot operations and d^2 Mul operations for each. Similar to our

Algorithm 1 New Matrix Multiplication MatMul

Input: $\text{CT}_{\iota(\mathbf{A})}$, $\text{CT}_{\tilde{\iota}(\mathbf{B})}$, and $\text{evk} = \{\text{rlk}, \text{rtk}_0, \dots, \text{rtk}_{d-1}\}$.

Output: $\text{CT}_{\iota(\mathbf{C})}$ where $\mathbf{C} = \mathbf{AB}$.

```

1: for  $0 \leq i < d$  do
2:    $(d_0, d_1, d_2, d_3) \leftarrow (0, 0, 0, 0) \in R_q^4$ 
3:   for  $0 \leq k < d$  do
4:      $(a_0, a_1) \leftarrow \text{ct}_{\iota(\mathbf{A})_k}$ 
5:      $(b_0, b_1) \leftarrow \text{ct}_{\tilde{\iota}(\mathbf{B})_{i-k}}$ 
6:      $(d_0, d_1, d_2, d_3) \leftarrow (d_0, d_1, d_2, d_3) + (a_0 \varphi^i(b_0), a_0 \varphi^i(b_1), a_1 \varphi^i(b_0), a_1 \varphi^i(b_1))$ 
7:   end for
8:    $(d'_0, d'_1, d'_2, d'_3) \leftarrow \left\lfloor \frac{X^D - b}{q} (d_0, d_1, d_2, d_3) \right\rfloor \pmod{q}$ 
   // key-switching from  $\varphi^i(s)$  to  $s$ 
9:    $(u_0, u_1) \leftarrow d'_2 \boxtimes \text{rtk}_i \pmod{q}$ 
   // key-switching from  $s \cdot \varphi^i(s)$  to  $s^2$ 
10:   $(v_0, v_1) \leftarrow d'_3 \boxtimes \text{rtk}_i \pmod{q}$ 
   // key-switching from  $s^2$  to  $s$ 
11:   $(w_0, w_1) \leftarrow v_1 \boxtimes \text{rlk} \pmod{q}$ 
12:   $\text{ct}_{\iota(\mathbf{C})_i} \leftarrow (d'_0 + u_0 + w_0, d'_1 + u_1 + v_0 + w_1) \pmod{q}$ 
13: end for

```

matrix multiplication method, it can utilize the lazy key-switching technique. However, since the rotation index of the summation term is different, it can only compute the relinearization in a lazy manner. Instead, [26] proposed a technique called *hoisting*, which precomputes the decomposition of a ciphertext before applying several rotations. Using these optimizations, the number of decomposition becomes $2d$. However, the number of inner product operation still remains $d^2 + d$. Finally, it requires $d - 1$ rotation keys since it computes every rotation from 1 to $d - 1$.¹

In the method by Jiang et al., it requires d Mul operations and $2d$ linear transformations using Halevi-Shoup method, as $\mathbf{M}\sigma(\mathbf{A})$ and $\mathbf{N}\sigma(\mathbf{B})$ can be precomputed. A crucial observation made in [21] is that $\mathbf{V}^{(k)}$ and $\mathbf{W}^{(k)}$ has sparse diagonal entries. To be precise, $\iota(\mathbf{V}^{(k)})$ has only two nonzero vectors $\iota(\mathbf{V}^{(k)})_k$ and $\iota(\mathbf{V}^{(k)})_{k-d}$, and $\iota(\mathbf{W}^{(k)})$ has only one nonzero vector $\iota(\mathbf{W}^{(k)})_{k-d}$. Therefore, using the Halevi-Shoup method, it can compute the linear transforms $\mathbf{V}^{(k)}\mathbf{M}\sigma(\mathbf{A})$ and $\mathbf{W}^{(k)}\mathbf{N}\sigma(\mathbf{B})$ for all $0 \leq k < d$ with $2d$ and d key-switchings respectively. Moreover, it can further utilize hoisting and lazy key-switching. This reduces the total number of decomposition to 3, and the number of inner products to $3d+1$. Finally, since it needs to compute rotation indices of the form $k, k-d, k \cdot d$ for every $0 \leq k < d$, the total number of rotation keys required is $3d - 3$. We note that the method by Jiang et al. requires two multiplicative depths, unlike other methods that only require one.

In contrast, our method performs only three key-switchings at the end of the summation. Therefore, we require just $3d$ decompositions and inner products, achieving a quadratic speedup compared to the Halevi-Shoup method. Similar to the Halevi-Shoup method, we require $d - 1$ rotation keys and one multiplication depth, which is lower than the method by Jiang et al.

We summarize the complexity analysis in Table 1. For

1. There are some methods to reduce the number of rotation keys, such as BSGS(Baby-Step, Giant-Step) algorithms. Since these algorithms degrade performance as a trade-off, we omit them from the comparison here.

a fair comparison, we also analyze the complexity of inner products and decompositions in terms of d . For both the Halevi-Shoup method and our method, the number of slots D is set to d and $D = O(N)$, so both decomposition and inner products have a complexity of $\tilde{O}(dL^2)$, as analyzed in Section 2.5, where L is the level of gadget decomposition. Meanwhile, for the method by Jiang et al., the number of slots is set to d^2 , so both decomposition and inner products have a complexity of $\tilde{O}(d^2L^2)$. We compute the complexity of decompositions and inner products based on this analysis in Table 1. Note that the complexity of basic ring arithmetic over R_q is $\tilde{O}(d^3L)$, equal for all methods.

In the previous method, the inner product is the most dominant factor, with a total complexity of $\tilde{O}(d^3L^2)$. However, our method reduces this complexity to $\tilde{O}(d^2L^2)$, making it quadratic in d . As a result, ring arithmetic now becomes the dominant factor in our method with complexity $\tilde{O}(d^3L)$, which improves the total complexity by a factor of L from $\tilde{O}(d^3L^2)$.

	Jiang et al. [21]	Halevi-Shoup [26]	Ours
Decomp.	$\tilde{O}(d^2L^2)$	$\tilde{O}(d^2L^2)$	$\tilde{O}(d^2L^2)$
Inner Product	$\tilde{O}(d^3L^2)$	$\tilde{O}(d^3L^2)$	$\tilde{O}(d^2L^2)$
Ring Arithmetic	$\tilde{O}(d^3L)$	$\tilde{O}(d^3L)$	$\tilde{O}(d^3L)$
Depth	1 PMu1 + 1 Mu1	1 Mu1	1 Mu1
# Rotation Key	$3d - 3$	$d - 1$	$d - 1$

TABLE 1: Comparison of number of basic operations required for multiplying two $d \times d$ matrices.

4.3. Batched Matrix Multiplication

As previously mentioned, our algorithm focused on the case where the matrix dimension d equals the number of slots D . However, in practice, the number of slots can exceed the matrix dimension. Hence, we propose batching techniques that can handle multiple matrices within a single ciphertext, inspired by the batched multiplication in [21]. With this technique, we can perform multiple matrix multiplications at once in batched form.

Let us elaborate on this in detail. Suppose that we are given ℓ matrices $\mathbf{A}^{(0)}, \dots, \mathbf{A}^{(\ell-1)}$ of size $d \times d$, where $\ell = D/d$. Then, the diagonal encoding map ι and shifted diagonal encoding map $\tilde{\iota}$ in Algorithm 1 can be naturally extended to the batched versions $\iota^{(\ell)} : (\mathbb{Z}_p^{d \times d})^\ell \rightarrow (\mathbb{Z}_p^D)^d$ and $\tilde{\iota}^{(\ell)} : (\mathbb{Z}_p^{d \times d})^\ell \rightarrow (\mathbb{Z}_p^D)^d$ as follows.

$$\begin{aligned} \iota^{(\ell)} : (\mathbf{A}^{(0)}, \dots, \mathbf{A}^{(\ell-1)}) &\mapsto \left((a_{i,j+i}^{(0)}, \dots, a_{i,j+i}^{(\ell-1)})_{0 \leq j < d} \right)_{0 \leq i < d} \\ \tilde{\iota}^{(\ell)} : (\mathbf{A}^{(0)}, \dots, \mathbf{A}^{(\ell-1)}) &\mapsto \left((a_{j-i,j}^{(0)}, \dots, a_{j-i,j}^{(\ell-1)})_{0 \leq j < d} \right)_{0 \leq i < d} \end{aligned}$$

Now, using the analogous logic we used in the proof of Eq. (3), we obtain the following equality.

$$\begin{aligned} &\iota(\mathbf{A}^{(0)}\mathbf{B}^{(0)}, \dots, \mathbf{A}^{(\ell-1)}\mathbf{B}^{(\ell-1)})_i \\ &= \sum_{k=0}^{d-1} \iota(\mathbf{A}^{(0)}, \dots, \mathbf{A}^{(\ell-1)})_k \odot \rho^{ik} \left(\tilde{\iota}(\mathbf{B}^{(0)}, \dots, \mathbf{B}^{(\ell-1)})_{i-k} \right) \end{aligned}$$

Then, we can compute ℓ matrix multiplications at once by calling a single instance of Algorithm 1 utilizing the equation above. This way, only $3d$ decomposition and inner product operations with d rotation keys are sufficient to perform ℓ matrix multiplications of dimension $d \times d$.

5. Protocol Specification

In this section, we present protocols for generating authenticated scalar and matrix triples. Our protocol is based on the TopGear [3] framework and its variant [4] for matrix triples. We begin by explaining how we modify the existing PoPK protocol to support the generalized BFV scheme. Then, we introduce a protocol for authenticated triple generation, incorporating all the optimization techniques discussed in Sections 3 and 4.

5.1. Proof of Plaintext Knowledge

Π_{PoPK}

This protocol is run between n parties, where each of them acts as both a prover and a verifier simultaneously. The protocol is parameterized by integers U, V , and $\text{Flag} \in \{\text{Diag}, \perp\}$. A public parameter pp and an encryption key pk are pre-distributed to each party.

Sampling phase

- 1) If $\text{Flag} = \perp$, party P_ℓ samples a message $\mathbf{x}_i^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$ for $0 \leq i < U$. If $\text{Flag} = \text{Diag}$, party P_ℓ samples a message $x_i^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p)$ and sets $\mathbf{x}_i^{(\ell)} = (x_i^{(\ell)}, \dots, x_i^{(\ell)}) \in \mathbb{Z}_p^D$ for $0 \leq i < U$.
- 2) Party P_ℓ samples a randomness $\mathbf{r}_i^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_1}^3$ and a plaintext $\mu_{\mathbf{x}_i^{(\ell)}} \leftarrow \text{RandPack}(\mathbf{x}_i^{(\ell)}, s_1)$ for $0 \leq i < U$.
- 3) Party P_ℓ computes a ciphertext $\text{ct}_i^{(\ell)} = \text{Enc}(\mu_{\mathbf{x}_i^{(\ell)}}, \mathbf{r}_i^{(\ell)})$, broadcasts it, and sets $\text{ct}_i = \sum_{\ell=0}^{n-1} \text{ct}_i^{(\ell)}$ for $0 \leq i < U$.

Commitment phase

- 1) If $\text{Flag} = \perp$, party P_ℓ samples a pseudo-message $\mathbf{y}_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$ for $0 \leq j < V$. If $\text{Flag} = \text{Diag}$, party P_ℓ samples a pseudo-message $y_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p)$ and sets $\mathbf{y}_j^{(\ell)} = (y_j^{(\ell)}, \dots, y_j^{(\ell)}) \in \mathbb{Z}_p^D$ for $0 \leq j < V$.
- 2) Party P_ℓ samples a pseudo-randomness $\mathbf{s}_j^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sqrt{U+1} \cdot \sigma_2}^3$ and a pseudo-plaintext $\mu_{\mathbf{y}_j^{(\ell)}} \leftarrow \text{RandPack}(\mathbf{y}_j^{(\ell)}, \sqrt{U+1} \cdot s_2)$ for $0 \leq j < V$.
- 3) Party P_ℓ computes a commitment $\text{comm}_j^{(\ell)} = \text{Enc}(\mu_{\mathbf{y}_j^{(\ell)}}, \mathbf{s}_j^{(\ell)})$, and broadcasts it for $0 \leq j < V$.

Challenge phase

- 1) Parties call $\mathcal{F}_{\text{Rand}}$ to obtain a challenge $w_{i,j}$ for $0 \leq i < U$ and $0 \leq j < V$.
- 2) If $\text{Flag} = \perp$, each challenge is sampled from $\{1, X, X^2, \dots, X^{2^{N-1}}\}$. If $\text{Flag} = \text{Diag}$, each challenge is sampled from $\{1, X^D, X^{2D}, \dots, X^{2^{N-D}}\}$.

Response phase

- 1) Party P_ℓ computes $\mu_{\mathbf{z}_j^{(\ell)}} = \mu_{\mathbf{y}_j^{(\ell)}} + \sum_{i=0}^{U-1} w_{i,j} \cdot \mu_{\mathbf{x}_i^{(\ell)}}$, $\mathbf{t}_j^{(\ell)} = \mathbf{s}_j^{(\ell)} + \sum_{i=0}^{U-1} w_{i,j} \cdot \mathbf{r}_i^{(\ell)}$, and broadcasts a response $\text{resp}_j^{(\ell)} = (\mu_{\mathbf{z}_j^{(\ell)}}, \mathbf{t}_j^{(\ell)})$ for $0 \leq j < V$.

Verification phase

- 1) The parties compute $\text{comm}_j = \sum_{\ell=0}^{n-1} \text{comm}_j^{(\ell)}$, $\mathbf{t}_j = \sum_{\ell=0}^{n-1} \mathbf{t}_j^{(\ell)}$, and $\mu_{\mathbf{z}_j} = \sum_{\ell=0}^{n-1} \mu_{\mathbf{z}_j^{(\ell)}}$ for $0 \leq j < V$.

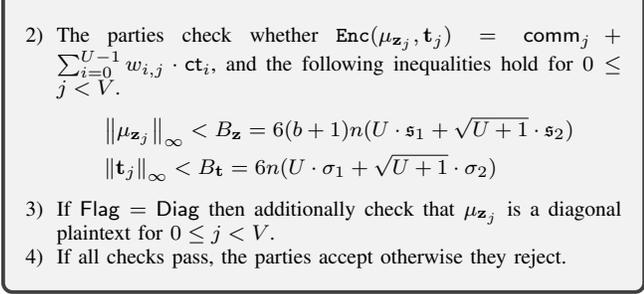


Figure 3: Protocol for proof of plaintext knowledge

A proof of plaintext knowledge (PoPK) protocol is essential for achieving adaptive security in the generation of authenticated triples, as it proves the well-formedness of the input ciphertexts. Its basic structure is a 3-move sigma protocol consisting of five phases: sampling, commitment, challenge, response, and verification. We provide a simplified overview of the PoPK protocol. In the sampling phase, a prover generates a ciphertext $ct = \text{Enc}(\mu_x, r)$ with the input message x by sampling a plaintext μ_x and encryption randomness r . In the commitment phase, the prover generates random masks μ_y and s for the plaintext μ_x and encryption randomness r , and sends a commitment $\text{comm} = \text{Enc}(\mu_y, s)$ to the verifier. After that, in the challenge phase, the verifier replies with a random challenge w . Subsequently, the prover responds with $\mu_z = \mu_y + w \cdot \mu_x$ and $t = s + w \cdot r$ to the verifier. Finally, the verifier runs verification algorithms to determine the validity of the ciphertext ct .

In the actual protocol, several modifications are added. First, since there are a total of n parties involved, the protocol is adapted to an n -prover version. Secondly, due to the limited size of the challenge sets, the sigma protocol must be repeated several times to achieve a reasonable security level, and this repetition count is denoted as V . Finally, to utilize the amortization technique in [3], there are usually multiple ciphertexts as input, with the number of input ciphertexts denoted by U . A detailed description of our PoPK protocol Π_{PoPK} is presented in Fig. 3.

Simulatability. We first demonstrate how we achieve simulatability with our generalized BFV scheme. In the PoPK protocol, the response $\text{resp} = (\mu_z, t)$ may contain partial information about the plaintext μ_x and the encryption randomness r , which should be kept secret to hide the input message x . Therefore, it must be simulatable without knowledge of the input message x to ensure that no information about the input messages is leaked. Achieving simulatability incurs parameter overhead, because the proven upper bounds for encryption noise are determined by μ_z and t . This results in an increase in the ciphertext modulus q , which increases the overall communication cost compared to the semi-honest setting. This overhead can be measured by the size ratio between (μ_z, t) and (μ_x, r) , referred to as the *soundness slack*. Thus, achieving a small soundness slack is crucial for overall performance.

In [3], the noise flooding method is used to achieve simulatability, which incurs exponential soundness slack for both r and μ_x . The exponential soundness slack for

r is reduced to a constant factor in [22] by introducing the Hint-MLWE problem, which proves that if both r and s are sampled from a discrete Gaussian distribution, t is computationally indistinguishable from a simulated one. For μ_x , the exponential soundness slack is partially addressed in [4] for BFV ciphertexts when $p \mid q$, i.e., $\Delta = q/p$.

However, in our generalized BFV scheme, the scaling factor Δ is of the form $\frac{q}{X^d - b}$, making the approach in [4] inapplicable. To achieve simulatability for the plaintext part with constant overhead, we adapt the recently proposed technique by Hwang et al. [23], called *randomized packing*. In [23], it is used to achieve simulatability in the proof of opening knowledge protocol for the Ajtai commitment scheme. However, we observe that the same technique can be utilized in the PoPK protocol for our generalized BFV scheme. The randomized packing is defined as follows.

- **RandPack**(m, s) $\rightarrow \mu$: Given a message $m \in \mathbb{Z}_p^D$ and a positive real $s > 0$, return $\mu \in R$, which is sampled from $\mathcal{D}_{[\text{Pack}(m)]_R + P\mathbb{Z}^N, sP}$, where $P \in \mathbb{Z}^{N \times N}$ is the negacyclic matrix corresponding to $X^D - b$, and $[\text{Pack}(m)]_R \in \mathbb{Z}^N$ is interpreted as a coefficient vector.

We first note that $\text{Unpack}(\text{RandPack}(x, s)) = x$ is preserved for any message $x \in \mathbb{Z}_p^D$, allowing **RandPack** to be used in place of **Pack** when packing messages into plaintexts. The core idea of **RandPack** is to randomize the output plaintext so that it follows a discrete Gaussian distribution. By generating both μ_x and μ_y using randomized packing, the convolution lemma from [35] can be applied to simulate $\mu_z = \mu_y + w \cdot \mu_x$, similar to the Hint-MLWE-based analysis for $t = s + w \cdot r$ in [22]. We also note that the size of the randomized plaintext is determined by the width parameter s in **RandPack**. Since we can set similarly sized width parameters for μ_x and μ_y , this allows to decrease the soundness slack for μ_x from an exponential scale to a constant scale.

Handling Diagonal Plaintexts. In the PoPK protocol for authenticated triples, an additional step is required to prove a special structure of plaintexts known as diagonal plaintexts, where each slot component contains the same element. This case is marked in Fig. 3 by setting $\text{Flag} = \text{Diag}$. For an ordinary BFV scheme, diagonal plaintexts are simply constant polynomials, but in our generalized BFV scheme, the structure of diagonal plaintexts changes. Specifically, they must be polynomials in X^D to ensure invariance under the automorphism φ , which implies that all slots consist of the same value. Our PoPK protocol must preserve this structure when proving knowledge of diagonal plaintexts. To achieve this, we sample the mask μ_y from diagonal plaintexts and set the challenge w to be a power of X^D . This modification is illustrated in Fig. 3. We note that this modification effectively reduces communication costs compared to the previous PoPK protocol, because we can set the repetition count V much lower due to the increased size of the challenge set. In the previous PoPK, the challenge set for diagonal plaintexts was limited to $\{0, 1\}$, with a size of 2. In contrast, our PoPK expands this to $\{1, X^D, \dots, X^{2N-D}\}$, with a size of $2M$. This larger challenge set allows for

a reduced repetition count V while maintaining the same security level.

Security. We prove the security of our PoPK protocol Π_{PoPK} in Section B.6, following the definition of security for n -party PoPK described in [3].

5.2. Generation of Authenticated Triples

Π_{Triple}

Initialize: All parties invoke $\mathcal{F}_{\text{KeyGen}}$ to obtain a shared public key pk , and evaluation key $\text{evk} = \{\text{rlk}, \text{rtk}_1, \dots, \text{rtk}_{D-1}\}$. Then, each party P_ℓ does the followings for $0 \leq \ell < n$:

- 1) Party P_ℓ invokes the sampling phase of Π_{PoPK} with parameter $U = 1$, $V = \lceil (\lambda_{\text{Snd}} + 2) / \log_2(2M) \rceil$, and $\text{Flag} = \text{Diag}$ to obtain a ciphertext $\text{ct}_\alpha^{(\ell)}$, and an additive share $\alpha^{(\ell)} \in \mathbb{Z}_p$.
- 2) All parties compute $\text{ct}_\alpha = 2 \cdot \sum_{\ell=0}^{n-1} \text{ct}_\alpha^{(\ell)}$, and run the remaining phases of Π_{PoPK} .

Scalar Triples: Parties run this protocol to generate T random authenticated scalar triples in \mathbb{Z}_p in one invocation, where $D \mid T$.

- 1) Party P_ℓ invokes the sampling phase of Π_{PoPK} with parameter $U = 2T/D$, $V = \lceil (\lambda_{\text{Snd}} + 2) / \log_2(2N) \rceil$, and $\text{Flag} = \perp$ to obtain ciphertexts $\text{ct}_{\mathbf{a}_i}^{(\ell)}$, $\text{ct}_{\mathbf{b}_i}^{(\ell)}$, and additive shares $\mathbf{a}_i^{(\ell)}, \mathbf{b}_i^{(\ell)} \in \mathbb{Z}_p^D$ for $0 \leq i < T/D$, and run the remaining phases of Π_{PoPK} .
- 2) All parties compute $\text{ct}_{\mathbf{a}_i} = 2 \cdot \sum_{\ell=0}^{n-1} \text{ct}_{\mathbf{a}_i}^{(\ell)}$, $\text{ct}_{\mathbf{b}_i} = 2 \cdot \sum_{\ell=0}^{n-1} \text{ct}_{\mathbf{b}_i}^{(\ell)}$ for $0 \leq i < T/D$.
- 3) All parties compute $\text{ct}_{\mathbf{c}_i} \leftarrow \text{Mul}(\text{ct}_{\mathbf{a}_i}, \text{ct}_{\mathbf{b}_i})$, $\text{ct}_{\alpha \mathbf{a}_i} \leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\mathbf{a}_i})$, $\text{ct}_{\alpha \mathbf{b}_i} \leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\mathbf{b}_i})$, and $\text{ct}_{\alpha \mathbf{c}_i} \leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\mathbf{c}_i})$ for $0 \leq i < T/D$.
- 4) Parties run Π_{DDec} with ciphertexts $\text{ct}_{\mathbf{c}_i}, \text{ct}_{\alpha \mathbf{a}_i}, \text{ct}_{\alpha \mathbf{b}_i}, \text{ct}_{\alpha \mathbf{c}_i}$ so that each P_ℓ obtains additive shares $\mathbf{c}_i^{(\ell)}, (\alpha \mathbf{a}_i)^{(\ell)}, (\alpha \mathbf{b}_i)^{(\ell)}, (\alpha \mathbf{c}_i)^{(\ell)}$ for $0 \leq i < T/D$.

Matrix Triples: Parties run this protocol to generate T random authenticated matrix triples in $\mathbb{Z}_p^{D \times D}$ in one invocation.

- 1) Party P_ℓ invokes the sampling phase of Π_{PoPK} with parameters $U = 2DT$, $V = \lceil (\lambda_{\text{Snd}} + 2) / \log_2(2N) \rceil$, and $\text{Flag} = \perp$ to obtain ciphertexts $\text{CT}_{\iota(\mathbf{A}_i)}^{(\ell)}$, $\text{CT}_{\tilde{\iota}(\mathbf{B}_i)}^{(\ell)}$, and additive shares $\iota(\mathbf{A}_i)^{(\ell)}, \tilde{\iota}(\mathbf{B}_i)^{(\ell)}$ for $0 \leq i < T$, and runs the remaining phases of Π_{PoPK} .
- 2) All parties compute $\text{CT}_{\iota(\mathbf{A}_i)} = 2 \cdot \sum_{\ell=0}^{n-1} \text{CT}_{\iota(\mathbf{A}_i)}^{(\ell)}$, $\text{CT}_{\tilde{\iota}(\mathbf{B}_i)} = 2 \cdot \sum_{\ell=0}^{n-1} \text{CT}_{\tilde{\iota}(\mathbf{B}_i)}^{(\ell)}$ for $0 \leq i < T$.
- 3) All parties compute $\text{CT}_{\iota(\mathbf{C}_i)} \leftarrow \text{MatMul}(\text{CT}_{\iota(\mathbf{A}_i)}, \text{CT}_{\tilde{\iota}(\mathbf{B}_i)})$ for $0 \leq i < T$.
- 4) All parties compute the followings for $0 \leq i < T$ and $0 \leq j < D$,

$$\begin{aligned} \text{ct}_{\alpha \cdot \iota(\mathbf{A}_i)_j} &\leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\iota(\mathbf{A}_i)_j}) \\ \text{ct}_{\alpha \cdot \tilde{\iota}(\mathbf{B}_i)_j} &\leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\tilde{\iota}(\mathbf{B}_i)_j}) \\ \text{ct}_{\alpha \cdot \iota(\mathbf{C}_i)_j} &\leftarrow \text{Mul}(\text{ct}_\alpha, \text{ct}_{\iota(\mathbf{C}_i)_j}) \end{aligned}$$

and set $\text{CT}_{\alpha \cdot \iota(\mathbf{A}_i)} = (\text{ct}_{\alpha \cdot \iota(\mathbf{A}_i)_j})_{0 \leq j < D}$, $\text{CT}_{\alpha \cdot \tilde{\iota}(\mathbf{B}_i)} = (\text{ct}_{\alpha \cdot \tilde{\iota}(\mathbf{B}_i)_j})_{0 \leq j < D}$, and $\text{CT}_{\alpha \cdot \iota(\mathbf{C}_i)} = (\text{ct}_{\alpha \cdot \iota(\mathbf{C}_i)_j})_{0 \leq j < D}$.

- 5) Parties run Π_{DDec} with ciphertexts $\text{CT}_{\iota(\mathbf{C}_i)}, \text{CT}_{\alpha \cdot \iota(\mathbf{A}_i)}, \text{CT}_{\alpha \cdot \tilde{\iota}(\mathbf{B}_i)}, \text{CT}_{\alpha \cdot \iota(\mathbf{C}_i)}$ so that each P_ℓ obtains additive shares $\iota(\mathbf{C}_i)^{(\ell)}, (\alpha \cdot \iota(\mathbf{A}_i))^{(\ell)}, (\alpha \cdot \tilde{\iota}(\mathbf{B}_i))^{(\ell)}, (\alpha \cdot \iota(\mathbf{C}_i))^{(\ell)}$ for $0 \leq i < T$.

Figure 4: Protocol for authenticated triples generation

We present our authenticated triple generation protocol Π_{Triple} in Fig. 4, which incorporates the generalized BFV scheme discussed in Section 3, the optimized homomorphic matrix multiplication algorithm in Section 4, and our PoPK

protocol Π_{PoPK} tailored for the generalized BFV scheme. Compared to prior work [4], which utilizes the ordinary BFV scheme and the matrix multiplication algorithm by Jiang et al. [21], we have modified the protocol to accommodate these changes. For matrix triples, we assume that the matrix dimension $d = D$, but it can be easily generalized to support smaller dimensions such that $d \mid D$ using the batching technique described in Section 4.3. In this case, the parameter U is adjusted to $2dT$, where $T \mid D/d$, and the number of rotation keys is reduced to $d - 1$. We omit the security proof for Π_{Triple} as it is identical to the proofs in [3, 4], except for the noise bound analysis.

Key Generation. For an authenticated triple generation protocol, previous work [3, 4] assumes the existence of an ideal key generation functionality, deferring detailed instantiation as an independent research topic, such as in [36]. Similarly, we also assume the existence of such functionality $\mathcal{F}_{\text{KeyGen}}$ for our triple generation protocol, as described in Fig. 5. In the functionality $\mathcal{F}_{\text{KeyGen}}$, it not only generates the shared public key pk but also the shared evaluation key evk , which includes the relinearization key rlk for homomorphic multiplication and the rotation key rtk_r for homomorphic rotation. As analyzed in Section 4, our functionality generates $d - 1$ rotation keys, whereas prior work [4] generates $3d - 3$ rotation keys for matrix multiplications, as it utilizes the algorithm in [21]. Thus, our new matrix multiplication algorithm reduces the number of rotation keys by one-third.

Elimination of Reshare. In our protocol, we eliminate all occurrences of the reshare functionality by adopting the approach from [4]. In [1, 3], the reshare functionality is utilized, which converts a ciphertext into secret shares of its message and a fresh encryption of it. This optimization allows for more compact parameter settings where the ciphertext modulus supports just one multiplicative depth, resulting in reduced communication costs for scalar triple generation. However, in [4], this functionality is not used for matrix triple generation, because it rather degrades performance, given that the matrix multiplication algorithm [21] used in [4] consumes multiple depths.

Although our matrix multiplication algorithm consumes only one multiplicative depth, we choose not to use the reshare functionality because the overhead for additional multiplicative depth is particularly low in our generalized BFV scheme. As analyzed in Section 3, noise growth in our generalized BFV scheme is proportional to b , rather than p as in the ordinary BFV scheme. This difference significantly reduces the increase in ciphertext modulus q due to the additional depth. Based on this observation, we completely eliminate all occurrences of the reshare functionality in our protocol by setting the ciphertext modulus to support the additional multiplicative depth, which provides better performance.

Noise Bound. In the protocol Π_{Triple} , authenticated scalar and matrix triples are generated through homomorphic operations and then distributed to each party via the distributed decryption protocol described in Fig. 6. To ensure the correctness of the protocol, we need to analyze an upper bound

for ciphertext noise to guarantee that distributed decryption correctly outputs the results. Among all the ciphertexts, the noise of $\text{CT}_{\alpha, \iota(\mathbf{C}_i)}$ is the most dominant. Therefore, we focus on analyzing the noise upper bound for these ciphertexts.

We provide a detailed analysis in Section B.7. The noise for $\text{CT}_{\alpha, \iota(\mathbf{C}_i)}$ is bounded as follows, which becomes the upper bound B_{Dec} for ciphertext noise in Π_{Triple} .

$$\|\text{Err}(\text{ct}_{\alpha, \iota(\mathbf{C}_i)_j}, \iota(\mathbf{C}_i)_j)\|_{\infty}^{\text{can}} \lesssim \frac{d(b+1)^3}{q} \cdot 8N^4 \sqrt{3N} \cdot B_{\mathbf{z}} = B_{\text{Dec}}$$

To ensure correct decryption in Π_{DDec} , the condition $B_{\text{Dec}} \cdot 2^{\lambda_{\text{DDec}}} < 1/2$ must hold, as stated in Lemma 6. This implies that the ciphertext modulus q should satisfy the following inequality: $q \gtrsim d(b+1)^3 \cdot 16N^4 \sqrt{3N} \cdot B_{\mathbf{z}} \cdot 2^{\lambda_{\text{DDec}}}$.

6. Experimental Results

We present the concrete parameter settings and experimental results for our new authenticated triple generation protocol. We implemented our generalized BFV scheme and new homomorphic matrix multiplication algorithm on top of the Lattigo library [37]. Our source code is available at an anonymous GitHub link: <https://github.com/SNUCP/matrix-triple>. All experiments were conducted on a machine with an Intel(R) Xeon(R) Platinum 8268 CPU at 2.90 GHz and 378 GB RAM, using a single thread. We first provide the concrete parameter settings used in our experiments, and then analyze the costs for setup, communication, and computation.

6.1. Parameter Setting

For the parameters of Π_{Triple} , we first set the security parameters λ_{Snd} and λ_{Sim} to 128, and λ_{DDec} to 80, as in the previous work [4]. Next, we assume the matrix dimension is given as $d = 128$ for the comparison with the previous work [4]. When the number of slots D exceeds $d = 128$, we use the batched matrix multiplication algorithm described in Section 4.3. For the parameter T , we set $T = D/d$ for matrix triples, which is the smallest values for T , then U is set to $2dT = 2D$.

For the parameters of Π_{PoPK} , we set the values of \mathfrak{s}_1 , \mathfrak{s}_2 , σ_1 , and σ_2 as follows so that they satisfy the conditions for simulatability.

$$\begin{aligned} \mathfrak{s}_1 &= \sqrt{2} \cdot \frac{b+1}{b-1} \cdot B_{\eta}, & \mathfrak{s}_2 &= \sqrt{2V} \cdot \frac{b+1}{b-1} \cdot B_{\eta} \\ \sigma_1 &= 2\sqrt{2} \cdot B_{\eta}, & \sigma_2 &= 2\sqrt{2V} \cdot B_{\eta} \end{aligned}$$

The upper bound B_{η} for the smoothing parameter is set to $\sqrt{\frac{\ln(2N(1+2^{\lambda_{\text{Sim}}}))}{\pi}}$, following Lemma 1.

For the generalized BFV parameters, we prepare three sets: Params I, II, and III, as shown in Table 2. Params I is optimized for performance, Params II is optimized for setup cost, and Params III corresponds to ordinary BFV for comparison with our generalized BFV. For Params I and II, we choose parameters $b \approx 2^{64}$ for Params I and $b \approx 2^{16}$

for Params II, for each field size $\log p$, so that $p = b^M + 1$ is a prime number satisfying the condition in Lemma 3. Then, we set the ciphertext modulus to be $q \geq d(b+1)^3 \cdot 16N^4 \sqrt{3N} \cdot B_{\mathbf{z}} \cdot 2^{\lambda_{\text{DDec}}}$ following the analysis in Section 5.2. For the key-switching procedure, we use the RNS gadget decomposition described in Section 2.5. Specifically, we set $q = \prod_{0 \leq i < L} q_i$, where q_i 's are prime numbers, each sized between 50 and 60 bits. Finally, the ring degree N is chosen to ensure the hardness of RLWE when the key distribution is $\chi = \mathcal{U}(R_3)$. According to the estimation from [38], Params I provides 118-bit security, and Params II, III provide 128-bit security.

6.2. Setup Cost

We first analyze the setup cost for authenticated triple generation, which is required for the initialization phase of Π_{Triple} . In this phase, each party receives a shared public key pk and an evaluation key $\text{evk} = \text{rlk}, \text{rtk}_1, \dots, \text{rtk}_{d-1}$ from the key generation functionality $\mathcal{F}_{\text{KeyGen}}$. Thus, we measure the total setup cost by the combined size of pk and evk .

The public key pk consists of a pair of R_q elements, so its size is calculated as $2N/8 \cdot \log q$ bytes. For rlk and each rtk_i , which are pairs of R_q^L elements, their size is computed as $2LN/8 \cdot \log q$ bytes each. In Table 3, we present the total setup cost for each parameter set, together with the setup cost of the previous work [4]. For the previous work [4], we estimate the cost by assuming it generates $3d - 3$ rotation keys, as it employs the matrix multiplication algorithm from [21]. Params I provides about 19 times smaller setup cost, and Params II offers 125 times smaller setup cost compared to the previous work [4]. We attribute these reductions in setup costs to both our generalized BFV scheme and new matrix multiplication algorithm. As analyzed in Section 3, the generalized BFV scheme allows for a much smaller ciphertext modulus q relative to a large plaintext modulus p . This not only reduces $\log q$, but also decreases L and N . Additionally, the new matrix multiplication algorithm requires only $d - 1$ rotation keys, compared to $3d - 3$ in the previous work. These improvements collectively result in a significant reduction in setup costs. In particular, for Params II, we minimized the size of $\log q$ by choosing a smaller b compared to Params I. This results in about 6.6 times smaller key sizes compared to Params I, which may be useful for less powerful machines, such as IoT devices.

6.3. Communication Cost

Next, we analyze the communication cost of Π_{Triple} when generating 128×128 dimensioned matrix triples, which is measured by the size of data sent by each party. First, we examine the communication cost of the PoPK protocol Π_{PoPK} , which is frequently used as a sub-protocol in Π_{Triple} . In the sampling phase, each party broadcasts U elements in R_q^2 , resulting in a size of $\frac{2UN}{8} \cdot \log q$ bytes. For the commitment phase, each party broadcasts V elements in R_q^2 , leading to a size of $\frac{2VN}{8} \cdot \log q$ bytes.

$\log p$	Params I					Params II					Params III		
	$\log q$	N	L	D	b	$\log q$	N	L	D	b	$\log q$	$N = D$	L
128	426	2^{14}	7	2^{13}	$2^{64} - 3072$	230	2^{13}	4	2^{10}	$2^{16} - 196$	686	2^{15}	12
256				2^{12}	$2^{64} - 64$				2^9	$2^{16} - 22$	1202	2^{16}	20
512				2^{11}	$2^{64} - 428$				2^8	$2^{16} - 72$	2231	2^{17}	37
1024				2^{10}	$2^{64} - 8$				2^7	$2^{16} - 28$	4279	2^{18}	71
2048				2^9	$2^{64} - 22$				2^6	$2^{16} - 190$	-		
4096				2^8	$2^{64} - 56$				2^5	$2^{16} - 288$			

TABLE 2: Parameter sets for benchmarks

Params I	Params II	Prev [4]
1.45 GB	0.22 GB	27.4 GB

TABLE 3: Setup costs for each parameter.

	Params I	Params II	Prev [4]
Init	121 MB	16.6 MB	816 MB
Matrix triple	13.6 MB	29.4 MB	12.5 MB

TABLE 4: Communication costs for each parameter

Finally, in the response phase, each party broadcasts V pairs in R , bounded by B_z and B_t respectively, resulting in a size of $\frac{VN}{8} \cdot (\log B_t + \log B_z)$ bytes. To summarize, each invocation of Π_{PoPK} requires a communication cost of $\frac{2(U+V)N}{8} \cdot \log q + \frac{VN}{8} \cdot (\log B_t + \log B_z)$ bytes.

Now, we analyze the communication cost in Π_{Triple} based on the analysis for Π_{PoPK} . In the initialization phase of Π_{Triple} , where it collectively generates the encryption ct_α of the MAC key α , Π_{PoPK} is invoked with parameters $U = 1$ and $V = \lceil (\lambda_{\text{Snd}} +) / 2M \rceil$, where $M = N/D$. For the matrix triple phase of Π_{Triple} , Π_{PoPK} is invoked with parameters $U = 2dT$ and $V = \lceil (\lambda_{\text{Snd}} +) / 2N \rceil$, and it ends with the distributed decryption protocol Π_{DDec} for $2U$ ciphertexts, adding a communication cost of $\frac{4UN}{8} \cdot \log q$ bytes.

In Table 4, we summarize the communication costs for the initialization and matrix triple phases. For the matrix triple phase, we report the amortized cost, which is obtained by dividing the total cost by the number T of batched matrices, as done in previous work [4]. For the costs of the previous work [4], we estimate the cost of initialization based on our analysis, as they did not provide data for the initialization phase, while the cost for the matrix triple phase is directly taken from the analysis in [4].

For the initialization phase, Params I achieves about 6.7 times smaller costs, and Params II achieves about 49 times smaller costs compared to the previous work. This saving is particularly meaningful for the initialization phase, where communication overhead dominates over computational overhead since it does not perform any nonlinear homomorphic operations. We attribute these savings to our optimization for diagonal plaintexts in Π_{PoPK} , which reduces the size of V by a factor of $\log M$. Since the parameter U is fixed at 1 during the initialization phase, V becomes the most dominant factor in determining the communication cost. Consequently, our optimization leads to a significant reduction in communication cost during the setup phase.

For matrix triples, our approach incurs higher costs compared to the previous work since, in the matrix triple phase, U is usually set to be much larger than V . In this case, the generalized BFV scheme rather increases the cost since the overhead for distributed decryption is multiplied by M . However, we note that this does not lead to significant performance degradation in the matrix triple phase, as computational overhead is the most dominant factor due to homomorphic matrix multiplications. For example, in the benchmarks of [4], the communication overhead typically accounts for about one percent of the total elapsed time in a single-thread setting.

6.4. Computation Cost

At last, we analyze the computation cost for Π_{Triple} , especially the cost of homomorphic operations in generating scalar and matrix triples.

Scalar Triples. For scalar triples, we present the effect of the generalized BFV scheme when generating authenticated triples for large prime fields, extending the limits of field size to 4096 bits. Since authenticated triple generation requires only two Mul operations, we measure the performance of Mul for each parameter set in Table 5. For the unit cost of Mul , which generates D scalar triples, Params I and II dominate Params III, which corresponds to the ordinary BFV scheme. This is because the asymptotic complexity follows $\tilde{O}(L^2N)$, as analyzed in Section 2.5, and the generalized BFV scheme offers much smaller values for N and L .

When it comes to the amortized cost, which is the cost per single triple, Params I achieves about a 1.14 times speedup for $\log p = 128$, and this gap gradually increases to a 3.6 times speedup for $\log p = 1024$ as $\log p$ grows. This is because the complexity of the amortized cost follows $\tilde{O}(L^2M)$, where $M = N/D$. For the ordinary BFV scheme, doubling the size of $\log p$ results in doubling L , while M is fixed at 1, which quadruples the amortized cost due to the L^2 factor. In contrast, for the generalized BFV scheme, we can double the size of $\log p$ by doubling M while keeping L fixed, which only doubles the amortized cost since the complexity is linear in M . This effect is well demonstrated in Table 5, where the gap between the generalized BFV and ordinary BFV schemes widens as $\log p$ increases. Although we could not report benchmark results for Params III when $\log p = 2048$ and 4096 due to parameter blow-up, we

$\log p$	Params I		Params II		Params III	
	Unit	Amortized	Unit	Amortized	Unit	Amortized
128	56.62ms	6.91 μ s	12.72ms	12.22 μ s	0.25s	7.85 μ s
256	56.58ms	13.81 μ s	12.76ms	24.91 μ s	1.34s	20.52 μ s
512	56.53ms	27.60 μ s	12.75ms	49.79 μ s	7.62s	58.15 μ s
1024	56.49ms	55.17 μ s	12.74ms	99.57 μ s	53.22s	203.03 μ s
2048	56.46ms	110.27 μ s	12.74ms	199.03 μ s	-	-
4096	56.33ms	220.05 μ s	12.75ms	398.49 μ s	-	-

TABLE 5: Benchmark results for MuL. The amortized costs are obtained from dividing the unit costs by the number of slots D .

$\log p$	Params I		Params II	
	Unit	Amortized	Unit	Amortized
128	47.77s	0.75s	12.21s	1.53s
256	47.64s	1.49s	12.31s	3.08s
512	48.34s	3.02s	12.31s	6.16s
1024	48.45s	6.06s	12.30s	12.30s
2048	47.83s	11.06s	-	-
4096	48.23s	24.12s	-	-

TABLE 6: Benchmark results for MatMuL when $d = 128$. The amortized costs are obtained from dividing the unit costs by the number T of batched matrices.

d	Params I	Prev [4]	Speedup
128	0.99s	34.25s	34.5 \times
256	6.84s	207.14s	30.28 \times
512	48.62s	1459.63s	30.02 \times

TABLE 7: Benchmark results for authenticated matrix triple generations when $\log p = 128$.

expect Params I to achieve about a 7 times speedup for $\log p = 2048$ and a 14 times speedup for $\log p = 4096$ compared to Params III, following the complexity analysis.

Matrix Triples. For matrix triples, we measure the performance of MatMuL and the generation of authenticated matrix triples, which is accomplished by additionally multiplying the encryption ct_α of the MAC key. Thanks to the generalized BFV scheme, the performance of MatMuL scales linearly with the field size $\log p$ as demonstrated in Table 6.

To demonstrate the effect of the new matrix multiplication algorithm, we compare the performance of authenticated matrix triple generation with the previous work [4] for the same matrix dimensions and field size in Table 7. The benchmark results for the previous work are taken from Table 3 in [4], and all costs are measured in an amortized sense, dividing the total cost by the number of batched matrices. As shown in Table 7, our method achieves an approximately 30 to 34 times improvement in each dimension compared to previous work. We attribute this speedup to the complexity improvements in our algorithm, which reduce the total complexity from $\tilde{O}(d^3 L^2)$ to $\tilde{O}(d^3 L)$. This improvement particularly synergizes well with our generalized BFV scheme, which further reduces the value of L for the

same field size $\log p$. As a result, our approach achieves a significant performance boost over the previous work by leveraging both the generalized BFV scheme and the new matrix multiplication algorithm.

Acknowledgments

This work was partly supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No. RS-2023-00211649) and Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. RS-2024-00399491, Development of Privacy-Preserving Multiparty Computation Techniques for Secure Multiparty Data Integration).

References

- [1] I. Damgård, V. Pastro, N. Smart, and S. Zakarias, “Multiparty computation from somewhat homomorphic encryption,” in *Annual Cryptology Conference*. Springer, 2012, pp. 643–662.
- [2] M. Keller, V. Pastro, and D. Rotaru, “Overdrive: Making spdz great again,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2018, pp. 158–189.
- [3] C. Baum, D. Cozzo, and N. P. Smart, “Using topgear in overdrive: a more efficient zkpk for spdz,” in *International Conference on Selected Areas in Cryptography*. Springer, 2019, pp. 274–302.
- [4] H. Chen, M. Kim, I. Razenshteyn, D. Rotaru, Y. Song, and S. Wagh, “Maliciously secure matrix multiplication with applications to private deep learning,” in *International Conference on the Theory and Application of Cryptology and Information Security*, 2020, pp. 31–59.
- [5] D. Beaver, “Efficient multiparty protocols using circuit randomization,” in *Advances in Cryptology—CRYPTO’91: Proceedings 11*. Springer, 1992, pp. 420–432.
- [6] P. Mohassel and Y. Zhang, “Secureml: A system for scalable privacy-preserving machine learning,” in *2017 IEEE symposium on security and privacy (SP)*. IEEE, 2017, pp. 19–38.
- [7] P. Mohassel and P. Rindal, “Aby3: A mixed protocol framework for machine learning,” in *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, 2018, pp. 35–52.
- [8] A. Dalskov, D. Escudero, and M. Keller, “Secure evaluation of quantized neural networks,” *Proceedings on Privacy Enhancing Technologies*, 2020.
- [9] S. Wagh, D. Gupta, and N. Chandran, “Secureml: 3-party secure computation for neural network training,” *Proceedings on Privacy Enhancing Technologies*, 2019.

- [10] J. Mono and T. Güneysu, “Implementing and optimizing matrix triples with homomorphic encryption,” in *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security*, 2023, pp. 29–40.
- [11] P. Reisert, M. Rivinius, T. Krips, and R. Küsters, “Overdrive lowgear 2.0: Reduced-bandwidth mpc without sacrifice,” in *Proceedings of the 2023 ACM Asia Conference on Computer and Communications Security*, 2023, pp. 372–386.
- [12] L. T. Brandão and R. Peralta, “Nist first call for multi-party threshold schemes,” National Institute of Standards and Technology, Tech. Rep., 2023.
- [13] J. Burkhardt, I. Damgård, T. K. Frederiksen, S. Ghosh, and C. Orlandi, “Improved distributed rsa key generation using the miller-rabin test,” in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, 2023, pp. 2501–2515.
- [14] M. Chen, R. Cohen, J. Doerner, Y. Kondi, E. Lee, S. Rosefield, and A. Shelat, “Multiparty generation of an rsa modulus,” in *Annual International Cryptology Conference*, 2020, pp. 64–93.
- [15] L. Grassi, C. Rechberger, D. Rotaru, P. Scholl, and N. P. Smart, “Mpc-friendly symmetric key primitives,” in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 2016, pp. 430–443.
- [16] W. Beullens, L. Dodgson, S. Faller, and J. Hesse, “The 2hash oprf framework and efficient post-quantum instantiations,” in *EUROCRYPT*, 2025.
- [17] Z. Brakerski, “Fully homomorphic encryption without modulus switching from classical gapsvp,” in *Annual cryptology conference*. Springer, 2012, pp. 868–886.
- [18] J. Fan and F. Vercauteren, “Somewhat practical fully homomorphic encryption,” *Cryptology ePrint Archive*, 2012.
- [19] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, “(leveled) fully homomorphic encryption without bootstrapping,” in *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, 2012, pp. 309–325.
- [20] C. Bootland, W. Castryck, I. Iliashenko, and F. Vercauteren, “Efficiently processing complex-valued data in homomorphic encryption,” *Journal of Mathematical Cryptology*, vol. 14, no. 1, pp. 55–65, 2020.
- [21] X. Jiang, M. Kim, K. Lauter, and Y. Song, “Secure outsourced matrix computation and application to neural networks,” in *Proceedings of the 2018 ACM SIGSAC conference on computer and communications security*, 2018, pp. 1209–1222.
- [22] D. Kim, D. Lee, J. Seo, and Y. Song, “Toward practical lattice-based proof of knowledge from hint-mlwe,” in *Annual International Cryptology Conference*. Springer, 2023, pp. 549–580.
- [23] I. Hwang, J. Seo, and Y. Song, “Concretely efficient lattice-based polynomial commitment from standard assumptions,” in *Annual International Cryptology Conference*. Springer, 2024, pp. 414–448.
- [24] E. Boyle, G. Couteau, N. Gilboa, Y. Ishai, L. Kohl, and P. Scholl, “Efficient pseudorandom correlation generators: Silent ot extension and more,” in *Annual International Cryptology Conference*, 2019, pp. 489–518.
- [25] D. Abram and P. Scholl, “Low-communication multiparty triple generation for spdz from ring-lpn,” in *IACR International Conference on Public-Key Cryptography*. Springer, 2022, pp. 221–251.
- [26] S. Halevi and V. Shoup, “Faster homomorphic linear transformations in HELIB,” in *Annual International Cryptology Conference*. Springer, 2018, pp. 93–120.
- [27] J. Jang, Y. Lee, A. Kim, B. Na, D. Yhee, B. Lee, J. H. Cheon, and S. Yoon, “Privacy-preserving deep sequential model with matrix homomorphic encryption,” in *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, 2022, pp. 377–391.
- [28] X. Zheng, H. Li, and D. Wang, “A new framework for fast homomorphic matrix multiplication,” *Designs, Codes and Cryptography*, pp. 1–23, 2025.
- [29] D. Micciancio and O. Regev, “Worst-case to average-case reductions based on gaussian measures,” *SIAM Journal on Computing*, vol. 37, no. 1, pp. 267–302, 2007.
- [30] W. Banaszczyk, “Inequalities for convex bodies and polar reciprocal lattices in \mathbb{R}^n ,” *Discrete & Computational Geometry*, vol. 13, pp. 217–231, 1995.
- [31] V. Lyubashevsky, C. Peikert, and O. Regev, “On ideal lattices and learning with errors over rings,” in *Proceedings of the 29th Annual international conference on Theory and Applications of Cryptographic Techniques*, 2010, pp. 1–23.
- [32] S. Halevi, Y. Polyakov, and V. Shoup, “An improved rms variant of the bfv homomorphic encryption scheme,” in *Cryptographers’ Track at the RSA Conference*, 2019, pp. 83–105.
- [33] S. Halevi and V. Shoup, “Bootstrapping for helib,” in *Annual International conference on the theory and applications of cryptographic techniques*. Springer, 2015, pp. 641–670.
- [34] D. Micciancio and V. Vaikuntanathan, “Sok: Learning with errors, circular security, and fully homomorphic encryption,” in *IACR International Conference on Public-Key Cryptography*. Springer, 2024, pp. 291–321.
- [35] C. Peikert, “An efficient and parallel gaussian sampler for lattices,” in *Annual Cryptology Conference*. Springer, 2010, pp. 80–97.
- [36] D. Rotaru, N. P. Smart, T. Tanguy, F. Vercauteren, and T. Wood, “Actively secure setup for spdz,” *Journal of Cryptology*, vol. 35, no. 1, p. 5, 2022.
- [37] C. V. Mouchet, J.-P. Bossuat, J. R. Troncoso-Pastoriza, and J.-P. Hubaux, “Lattigo: A multiparty homomorphic encryption library in go,” in *Proceedings of the 8th Workshop on Encrypted Computing and Applied Homomorphic Cryptography*, 2020, pp. 64–70.
- [38] M. R. Albrecht, R. Player, and S. Scott, “On the concrete hardness of learning with errors,” *Journal of Mathematical Cryptology*, vol. 9, no. 3, pp. 169–203, 2015.

Appendix A. Deferred Protocols

$\mathcal{F}_{\text{KeyGen}}$

- 1) Upon receiving start from all honest parties, sample $(\text{sk}, \text{pk}) \leftarrow \text{KeyGen}(\text{pp})$, $\text{rlk} \leftarrow \text{RelinKeyGen}(\text{sk})$, $\text{rtk}_r \leftarrow \text{RotKeyGen}(\text{sk}, r)$ for $1 \leq r < D$, send $\text{pk}, \text{evk} = \{\text{rlk}, \text{rtk}_1, \dots, \text{rtk}_{D-1}\}$ to the adversary, and store $\text{sk} = (1, s)$.
- 2) Receive the shares of secret key $s^{(\ell)}$ from the adversary.
- 3) Construct a complete set of shares $(s^{(0)}, \dots, s^{(\ell-1)})$ consistent with the adversary’s choice and s .
- 4) Send pk, evk and $s^{(\ell)}$ to each honest party P_ℓ .

Figure 5: The ideal functionality for key generation

Π_{DDec}

This protocol is run between n parties, where each P_ℓ holds a shared ciphertext $\text{ct} = (c_0, c_1)$ and a secret key share $s^{(\ell)}$ for $0 \leq \ell < n$. At the end of the protocol, each party obtains an additive share of the message \mathbf{m} of ct .

- 1) Party P_ℓ samples $e^{(\ell)} \leftarrow \mathcal{U}(R_{B_{\text{Dec}} \cdot 2^{\lambda_{\text{DDec}}}})$, $\mathbf{x}^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$, and sets $\mu_{\mathbf{x}}^{(\ell)} \leftarrow \text{Pack}(\mathbf{x}^{(\ell)})$.
- 2) Party P_ℓ computes $d^{(\ell)} = c_1 s^{(\ell)} + e^{(\ell)} + \Delta \mu_{\mathbf{x}}^{(\ell)} \pmod{q}$, and

distributes it.

3) Party P_ℓ computes $\mu = \left\lfloor \frac{X^D - b}{q} (c_0 + \sum_{\ell=0}^{n-1} d^{(\ell)}) \right\rfloor \pmod{X^D - b}$. If $\ell = 0$, P_0 sets $\mathbf{m}^{(0)} = \text{Unpack}(\mu) - \mathbf{x}^{(0)} \pmod{p}$. Otherwise, P_ℓ sets $\mathbf{m}^{(\ell)} = -\mathbf{x}^{(\ell)} \pmod{p}$.

Figure 6: Protocol for distributed decryption

Appendix B. Deferred Proofs

B.1. Proof of Lemma 3

Proof. From the relation $p = b^M + 1$, it follows that the order of b modulo p is $2M$. Therefore, there exists a primitive root g modulo p such that $g^{\frac{p-1}{2M}} = b$. Then, $\xi := g^{\frac{p-1}{2N}}$ is a root of the equation $X^D - b = 0$ modulo p . Moreover, ξ^{2iM+1} is also a root since $\xi^{(2iM+1)D} = \xi^{2iN+D} = \xi^D$ for any $i \in \mathbb{Z}$. Since the order of ξ is N , $\{\xi^{2iM+1}\}_{0 \leq i < D}$ are distinct and they are essentially the roots of the given equation. \square

B.2. Proof of Lemma 4

Proof. For any two elements $2iM + 1, 2jM + 1 \in G$, $(2iM + 1)(2jM + 1) = 2(2ijM + i + j)M + 1 \in G$ holds, so G is closed under multiplication. Also, G contains the identity $1 \in G$. For an element $2iM + 1 \in G$, it has a multiplicative inverse $x \in \mathbb{Z}_{2N}^\times$ such that $(2iM + 1)x = 1 \pmod{2N}$, then $(2iM + 1)x = x = 1 \pmod{2M}$ as $2M$ divides $2N$, so $x \in G$. Then each element in G has an inverse in G . Therefore, G is a subgroup of $\mathbb{Z}_{2N}^\times = \langle 5, -1 \rangle$. Since $-1 \notin G$ as $M \geq 2$, G is a (cyclic) subgroup of $\langle 5 \rangle$. As $|G| = D$, it follows that G is generated by $5^{M/2}$. \square

B.3. Proof of Theorem 1

Proof. By Lemmas 3 and 4, $\xi \in \mathbb{Z}_p$ such that $\xi, \xi^{5^{M/2}}, \dots, \xi^{5^{(D-1)M/2}}$ are distinct roots of $X^D - b$ modulo p . Then, $(X^D - b) = (X - \xi) \dots (X - \xi^{5^{(D-1)M/2}}) \pmod{p}$ and each linear factors are mutually coprime in $\mathbb{Z}_p[X]$. By the Chinese remainder theorem, $\mathbb{Z}_p[X]/(X^D - b) \simeq \mathbb{Z}_p[X]/(X - \xi) \times \dots \times \mathbb{Z}_p[X]/(X - \xi^{5^{(D-1)M/2}})$. Note that $\mathbb{Z}_p[X]/(X - \xi^{5^{iM/2}})$ is isomorphic to \mathbb{Z}_p by the isomorphism $f \mapsto f(\xi^{5^{iM/2}})$, it follows that $R_{X^D - b} = \mathbb{Z}_p[X]/(X^D - b) \simeq \mathbb{Z}_p^d$ with an isomorphism $\tau : f \mapsto (f(\xi), f(\xi^{5^{M/2}}), \dots, f(\xi^{5^{(D-1)M/2}}))$. \square

B.4. Noise Bound of Generalized BFV Scheme

We analyze the noise growth of each algorithm under the secret key distribution $\chi = \mathcal{U}(R_3)$ and the error distribution $\psi = \mathcal{D}_\sigma$. The noise growth during the encryption algorithm is as follows.

Lemma 5 (Encryption Noise). *Let $\mu \in R$, $\mathbf{r} \in R^3$ and $\text{ct} \leftarrow \text{Enc}(\text{pk}, \mu, \mathbf{r})$, then the following holds with very high probability*

$$\|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} < \frac{b+1}{q} \left(\|\mu\|_\infty \sqrt{3N^3} + \|\mathbf{r}\|_\infty \left(\left(\sqrt{\frac{18}{\pi}} \sigma + 2 \right) \sqrt{N^3} + N \right) \right)$$

Proof. Let $\mathbf{r} = (r_0, r_1, r_2)$ and $\text{ct} = (c_0, c_1)$, then the following holds for some $I \in R$ and $e_{rd} = q/(X^D - b) - \Delta$.

$$\frac{X^D - b}{q} (c_0 + c_1 s) = \mu + \frac{X^D - b}{q} (e_{rd} \cdot \mu + r_0 \cdot e + r_1 + r_2 \cdot s) + (X^D - b)I$$

Following the analysis in [20], we get $\|e\|_\infty^{\text{can}} < \sqrt{18/\pi} \sigma \sqrt{N}$, $\|s\|_\infty^{\text{can}} < 2\sqrt{6N}$, and $\|e_{rd}\|_\infty^{\text{can}} < \sqrt{3N}$ with very high probability. Also, it holds that $\|\mathbf{r}\|_\infty^{\text{can}} \leq N\|\mathbf{r}\|_\infty$ and $\|\mu\|_\infty^{\text{can}} \leq N\|\mu\|_\infty$. Thus, the given inequality holds with very high probability. \square

The following lemma shows that the decryption works correctly if the canonical norm of the invariant noise is less than $1/2$.

Lemma 6 (Decryption Bound [20]). *Let ct be a ciphertext and $\mu \in R$ be a plaintext. If $\|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} < 1/2$, then $\text{Dec}(\text{sk}, \text{ct}) = \mu \pmod{X^D - b}$.*

The noise growth of each homomorphic operation is as follows.

Lemma 7 (Addition Noise [20]). *Let ct and ct' be ciphertexts with plaintexts μ and μ' , respectively. For $\text{ct}'' \leftarrow \text{Add}(\text{ct}, \text{ct}')$, it holds that*

$$\|\text{Err}(\text{ct}'', \mu + \mu')\|_\infty^{\text{can}} \leq \|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} + \|\text{Err}(\text{ct}', \mu')\|_\infty^{\text{can}}$$

Lemma 8 (Multiplication Noise [20]). *Let ct and ct' be ciphertexts with plaintexts μ and μ' , respectively. Let rlk be a relinearization key. For $\text{ct}'' \leftarrow \text{Mul}(\text{rlk}, \text{ct}, \text{ct}')$, the following holds with very high probability.*

$$\begin{aligned} \|\text{Err}(\text{ct}'', \mu\mu')\|_\infty^{\text{can}} &\leq (b+1)\sqrt{2N^2 + 3N} (\|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} + \|\text{Err}(\text{ct}', \mu')\|_\infty^{\text{can}}) \\ &\quad + \frac{b+1}{q} \sqrt{4N^3/3 + 2N^2 + 3N} + \frac{b+1}{q} \sigma_{BN} \sqrt{\frac{3L}{2\pi}} \\ &\quad + 3\|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} \cdot \|\text{Err}(\text{ct}', \mu')\|_\infty^{\text{can}} \end{aligned}$$

Lemma 9 (Rotation Noise). *Let ct be a ciphertext with a plaintext μ . Let rtk_r be a rotation key with a rotation index r . For $\text{ct}' \leftarrow \text{Rot}(\text{rtk}_r, \text{ct})$, the following holds with very high probability*

$$\|\text{Err}(\text{ct}', \varphi^r(\mu))\|_\infty^{\text{can}} \leq \|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} + \frac{b+1}{q} \sigma_{BN} \sqrt{\frac{3L}{2\pi}}$$

Proof. Let $\text{ct} = (c_0, c_1)$, $\text{ct}' = (c'_0, c'_1)$, and $\text{rtk}_r = (-s \cdot \mathbf{a} + \varphi^r(s) \cdot \mathbf{g} + \mathbf{e}, \mathbf{a})$. Then, it holds that $c'_0 + c'_1 s = \varphi^r(c_0) + \varphi^r(c_1) \cdot \varphi^r(s) + \langle h(c'_1), \mathbf{e} \rangle \pmod{q}$. Following the analysis of noise growth during key switching as detailed in [20], the given inequality holds. \square

B.5. Noise Bound of MatMul

Lemma 10 (Noise Growth of Algorithm 1). *For the input encrypted matrices $\text{CT}_{\iota(\mathbf{A})}$ and $\text{CT}_{\tilde{\iota}(\mathbf{B})}$, and the output*

encrypted matrix $\text{CT}_{\iota(\mathbf{C})}$ of Algorithm 1, let $e_{\mathbf{A}}$ and $e_{\mathbf{B}}$ be reals such that

$$\|\text{Err}(\text{ct}_{\iota(\mathbf{A})_i}, \iota(\mathbf{A})_i)\|_{\infty}^{\text{can}} \leq e_{\mathbf{A}} \text{ and } \|\text{Err}(\text{ct}_{\tilde{\iota}(\mathbf{B})_i}, \tilde{\iota}(\mathbf{B})_i)\|_{\infty}^{\text{can}} \leq e_{\mathbf{B}}$$

for all $0 \leq i < d$. Then, it follows that

$$\|\text{ct}_{\iota(\mathbf{C})_i}\|_{\infty}^{\text{can}} \leq d(b+1)\sqrt{2N^2+3N}(e_{\mathbf{A}}+e_{\mathbf{B}})+3de_{\mathbf{A}}e_{\mathbf{B}} \\ + \frac{b+1}{q}\sqrt{\frac{4N^3}{3}+4N^2+3N} + \frac{b+1}{q}\sigma BN\sqrt{\frac{3(N+2)L}{2\pi}}$$

for all $0 \leq i < d$.

Proof. Note that the computation algorithm remains the same for every $\text{ct}_{\iota(\mathbf{C})_i}$, we estimate the noise for a single ciphertext $\text{ct}_{\iota(\mathbf{C})_i}$ without loss of generality. First, the noise growth from a single ciphertext tensoring of $\text{ct}_{\iota(\mathbf{A})_j}$ and $\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-j}}$ is

$$(b+1)\sqrt{2N^2+3N}(\|\text{ct}_{\iota(\mathbf{A})_j}\|_{\infty}^{\text{can}} + \|\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-j}}\|_{\infty}^{\text{can}}) \\ + 3\|\text{ct}_{\iota(\mathbf{A})_j}\|_{\infty}^{\text{can}} \cdot \|\text{ct}_{\tilde{\iota}(\mathbf{B})_{i-j}}\|_{\infty}^{\text{can}}$$

Therefore, after the for loop (lines 3-6), the noise is bounded by

$$d\left((b+1)\sqrt{2N^2+3N}(e_{\mathbf{A}}+e_{\mathbf{B}})+3e_{\mathbf{A}}e_{\mathbf{B}}\right).$$

Now, note that the tensoring ciphertext after the for loop is encrypted under the key $(1, s, \varphi^i(s), s\varphi^i(s))$ as specified, the modulus changing performed in line 7 introduces the rounding noise bounded as follows.

$$\frac{b+1}{q}\sqrt{4N^3/3+4N^2+3N}$$

Finally, we analyze the noise from the key-switching. Let e_1, e_2 and e_3 denote that error from key-switching operations in lines 8-10, respectively, i.e., $u_0 + u_1 \cdot s = d'_2 \cdot \varphi^i(s) + e_1 \pmod{q}$, $v_0 + v_1 \cdot s = d'_3 \cdot \varphi^i(s) + e_2 \pmod{q}$ and $w_0 + w_1 \cdot s = v_1 \cdot s^2 + e_3 \pmod{q}$. Then, the error from key-switching is $e_1 + e_2 \cdot s + e_3$ since the following holds.

$$d'_0 + u_0 + w_0 + s \cdot (d'_1 + u_1 + v_0 + w_1) \\ = d'_0 + d'_1 \cdot s + d'_2 \cdot \varphi^i(s) + d'_3 \cdot s\varphi^i(s) + e_1 + e_2 \cdot s + e_3$$

Since the variance of $e_1 + e_2 \cdot s + e_3$ is equal to $(N+2) \cdot L(\sigma BN)^2/24\pi$, the canonical norm of the key-switching is bounded by $(b+1)/q\sigma BN\sqrt{3(N+2)L/2\pi}$. Therefore, the final noise bound can be computed as follows.

$$d(b+1)\sqrt{2N^2+3N}(e_{\mathbf{A}}+e_{\mathbf{B}})+3de_{\mathbf{A}}e_{\mathbf{B}} \\ + \frac{b+1}{q}\sqrt{4N^3/3+4N^2+3N} + \frac{b+1}{q}\sigma BN\sqrt{\frac{3(N+2)L}{2\pi}}$$

We refer to Appendix A of [20] for a detailed analysis of the rounding noise. \square

B.6. Security Proof of Π_{PoPK}

Theorem 2. *The n -party PoPK-protocol Π_{PoPK} defined in Fig. 3 satisfies the following three properties:*

- **Correctness:** *If all parties P_ℓ , with inputs sampled from the sampling phase in Π_{PoPK} , follow the protocol honestly, then the protocol ends with acceptance with high probability, which is at least $1 - 2^{-128}$.*

- **Knowledge Soundness:** *Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ be a tuple of PPT algorithms and let $\varepsilon \in [0, 1)$. Consider the following game:*

- 1) \mathcal{A}_1 takes no input and outputs $C \subset [n]$, $\{(\text{ct}_i^{(\ell)})_{0 \leq i < U}\}_{\ell \in C}$, and $\text{stat}_{\mathcal{A}_1}$.
- 2) Sample $(\text{ct}_i^{(\ell)}, \mu_{\mathbf{x}_i}^{(\ell)}, \mathbf{r}_i^{(\ell)})_{0 \leq i < U}$ from the sampling phase for each $\ell \notin C$.
- 3) Compute $((\text{comm}_j^{(\ell)})_{0 \leq j < V}, \text{stat}^{(\ell)})$ from the commitment phase for each $\ell \notin C$.
- 4) \mathcal{A}_2 on input $\text{stat}_{\mathcal{A}_1}, (\text{ct}_i^{(\ell)})_{0 \leq i < U}, (\text{comm}_j^{(\ell)})_{0 \leq j < V}$ for $i \notin C$, outputs $\text{stat}_{\mathcal{A}_2}$ and $(\text{comm}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \in C$.
- 5) Choose a uniformly random $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$ and compute $(\text{resp}_j^{(\ell)})_{0 \leq j < V}$ from the response phase for $\ell \notin C$.
- 6) \mathcal{A}_3 on input $\text{stat}_{\mathcal{A}_2}, (w_{i,j}), (\text{resp}_j^{(\ell)})$ for $0 \leq i < U, 0 \leq j < V$, and $\ell \notin C$, outputs $(\text{resp}_j^{(\ell)})$ for $0 \leq j < V$ and $\ell \in C$.
- 7) \mathcal{A} wins the game if the verification phase ends with acceptance.

Suppose \mathcal{A} wins the game with probability $\delta > \varepsilon$. Then, there exists a PPT algorithm \mathcal{E} which for any fixed output of \mathcal{A}_1 , honestly generated inputs given by $\{(\text{ct}_i^{(\ell)}, \mu_{\mathbf{x}_i}^{(\ell)}, \mathbf{r}_i^{(\ell)})_{0 \leq i < U}, (\text{comm}_j^{(\ell)})_{0 \leq j < V}, \text{stat}^{(\ell)}\}_{\ell \notin C}$, and black-box access to $\mathcal{A}_2, \mathcal{A}_3$, outputs $\{(\mu_{\mathbf{x}_i}^{(\ell)}, \mathbf{r}_i^{(\ell)})_{0 \leq i < U}\}_{\ell \in C}$ such that

$$\text{ct}_i^{(\ell)} = \text{Enc}(\mu_{\mathbf{x}_i}^{(\ell)}, \mathbf{r}_i^{(\ell)}), \|\mu_{\mathbf{x}_i}^{(\ell)}\|_{\infty} < 2NB_{\mathbf{z}}, \|\mathbf{r}_i^{(\ell)}\|_{\infty} < 2NB_{\mathbf{t}}$$

holds for $0 \leq i < U$ and $\ell \in C$ in at most $f(\lambda_{\text{Snd}})/(\delta - \varepsilon)$ steps, where $f(\cdot)$ is a positive polynomial and $\varepsilon = 2^{-\lambda_{\text{Snd}}}$.

- **Simulatability:** *Let $\varepsilon = 2^{-\lambda_{\text{Sim}}}$ be negligible, and let $\sigma_1, \sigma_2, \mathfrak{s}_1$, and \mathfrak{s}_2 satisfy the following conditions.*

$$- \frac{\mathfrak{s}_2}{\sqrt{2}}, \mathfrak{s} \geq \frac{1}{b-1} \cdot \eta_{\varepsilon}(P\mathbb{Z}^N), \text{ where } \mathfrak{s} = (\mathfrak{s}_1^{-2} + V\mathfrak{s}_2^{-2})^{-1/2}$$

$$- \frac{\sigma_2}{\sqrt{2}}, \frac{\sigma}{\sqrt{2}} \geq \eta_{\varepsilon}(\mathbb{Z}^N), \text{ where } \sigma = \frac{1}{\sqrt{2}} \cdot (\sigma_1^{-2} + V\sigma_2^{-2})^{-1/2}$$

Then, there exists a PPT algorithm \mathcal{S} , which takes as input a challenge $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$ and a set $C \subset [n]$, and outputs $(\text{ct}_i^{(\ell)})_{0 \leq i < U}, (\text{comm}_j^{(\ell)}, \text{resp}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \notin C$, which is computationally indistinguishable from a valid execution of the protocol assuming the hardness of RLWE $_{R,q,\sigma,\sigma}$.

Proof. We prove the correctness, knowledge soundness, and simulatability of Π_{PoPK} as follows.

$\mathcal{S}_{\text{PoPK}}$

The inputs are a challenge $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$ and a set $C \subset [n]$.

- 1) Sample $\mu_{\mathbf{x}_i}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{0}, \mathfrak{s}_1)$, $\mathbf{r}_i^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_1}$ for $0 \leq i < U$, $\ell \notin C$.
- 2) Sample $\text{ct}_i^{(\ell)} \leftarrow \mathcal{U}(R_q^2)$ for $0 \leq i < U$, $\ell \notin C$.
- 3) Sample $\mathbf{y}_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$, $\mu_{\mathbf{y}_j}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{y}_j^{(\ell)}, \sqrt{U+1} \cdot \mathfrak{s}_2)$, $\mathfrak{s}_j^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sqrt{U+1} \cdot \sigma_2}$ for $0 \leq j < V$, $\ell \notin C$.
- 4) Set $\mu_{\mathbf{z}_j}^{(\ell)} = \mu_{\mathbf{y}_j}^{(\ell)} + \sum_{i=0}^{U-1} w_{i,j} \cdot \mu_{\mathbf{x}_i}^{(\ell)}$, $\mathbf{t}_j^{(\ell)} = \mathfrak{s}_j^{(\ell)} + \sum_{i=0}^{U-1} w_{i,j} \cdot \mathbf{r}_i^{(\ell)}$ for $0 \leq j < V$, $\ell \notin C$.
- 5) Set $\text{comm}_j^{(\ell)} = \text{Enc}(\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{t}_j^{(\ell)}) - \sum_{i=0}^{U-1} w_{i,j} \cdot \text{ct}_i^{(\ell)} \pmod{q}$ and $\text{resp}_j^{(\ell)} = (\mu_{\mathbf{z}_j}^{(\ell)}, \mathfrak{s}_j^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 6) Output $(\text{ct}_i^{(\ell)})_{0 \leq i < U}, (\text{comm}_j^{(\ell)})_{0 \leq j < V}$, and $(\text{resp}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \notin C$.

Figure 7: Simulator for Π_{PoPK}

- **Correctness:** By Lemma 2, the followings hold

$$\begin{aligned} \|\mu_{\mathbf{x}_i}^{(\ell)}\|_\infty &< 6(b+1)s_1, \quad \|\mu_{\mathbf{y}_j}^{(\ell)}\|_\infty < 6(b+1)\sqrt{U+1} \cdot s_2 \\ \|\mathbf{r}_i^{(\ell)}\|_\infty &< 6\sigma_1, \quad \|\mathbf{s}_j^{(\ell)}\|_\infty < 6\sqrt{U+1} \cdot \sigma_2 \end{aligned}$$

with high probability. Thus, it holds that $\|\mu_{\mathbf{z}_j}\|_\infty < B_z$ and $\|\mathbf{t}_j\|_\infty < B_t$ with high probability.

- **Knowledge Soundness:** We refer to the proof of Theorem 4 in [4].
- **Simulatability:** We can show that $\mathcal{S}_{\text{PoPK}}$ in Fig. 7 is a simulator for Π_{PoPK} in a similar manner to Theorem 4 in [23]. Below, we provide a brief sketch of the proof. We prove the simulatability of the protocol using the hybrid arguments $\mathcal{H}_0, \mathcal{H}_1, \mathcal{H}_2$ in Fig. 8. First, we show that the real transcript is statistically indistinguishable from \mathcal{H}_0 using the convolution lemma [35] for discrete Gaussian distributions. Second, we show that \mathcal{H}_0 is computationally indistinguishable from \mathcal{H}_1 due to the hardness of the Hint-RLWE problem. Next, \mathcal{H}_1 and \mathcal{H}_2 are statistically indistinguishable thanks to the randomized packing method. Finally, \mathcal{H}_2 and $\mathcal{S}_{\text{PoPK}}$ are statistically indistinguishable due to the convolution lemma again. Therefore, the real transcript is computationally indistinguishable from the simulated one. For a more detailed analysis of each hybrid argument, we refer to [23].

□

\mathcal{H}_0

The inputs are a challenge $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$, a set $C \subset [n]$, and messages $(\mathbf{x}_i^{(\ell)})_{0 \leq i < U}$ for $\ell \notin C$.

- 1) Sample $\mu_{\mathbf{x}_i}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{x}_i^{(\ell)}, s_1)$, $\mathbf{r}_i^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_1}$ for $0 \leq i < U$, $\ell \notin C$.
- 2) Set $\text{ct}_i^{(\ell)} = \text{Enc}(\mu_{\mathbf{x}_i}^{(\ell)}, \mathbf{r}_i^{(\ell)})$ for $0 \leq i < U$, $\ell \notin C$.
- 3) Sample $\mathbf{y}_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$, $\mu_{\mathbf{y}_j, i}^{(\ell)} \leftarrow \text{RandPack}(-\mathbf{W}_{i,j} \mathbf{x}_i^{(\ell)}, s_2)$, $\mu_{\mathbf{y}_j, U}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{y}_j^{(\ell)} + \sum_{i=0}^{U-1} \mathbf{W}_{i,j} \mathbf{x}_i^{(\ell)}, s_2)$, $\mathbf{s}_{j,i}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$ and $\mathbf{s}_{j,U}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$ for $0 \leq i < U$, $0 \leq j < V$, $\ell \notin C$, where $\mathbf{W}_{i,j}$ is the negacyclic matrix of $w_{i,j}$.
- 4) Set $\mu_{\mathbf{z}_j}^{(\ell)} = \mu_{\mathbf{y}_j, U}^{(\ell)} + \sum_{i=0}^{U-1} (\mu_{\mathbf{y}_j, i}^{(\ell)} + w_{i,j} \cdot \mu_{\mathbf{x}_i}^{(\ell)})$, $\mathbf{t}_j^{(\ell)} = \mathbf{s}_{j,U}^{(\ell)} + \sum_{i=0}^{U-1} (\mathbf{s}_{j,i}^{(\ell)} + w_{i,j} \cdot \mathbf{r}_i^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 5) Set $\text{comm}_j^{(\ell)} = \text{Enc}(\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{t}_j^{(\ell)}) - \sum_{i=0}^{U-1} w_{i,j} \cdot \text{ct}_i^{(\ell)} \pmod{q}$ and $\text{resp}_j^{(\ell)} = (\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{s}_j^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 6) Output $(\text{ct}_i^{(\ell)})_{0 \leq i < U}$, $(\text{comm}_j^{(\ell)})_{0 \leq j < V}$, and $(\text{resp}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \notin C$.

\mathcal{H}_1

The inputs are a challenge $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$, a set $C \subset [n]$, and messages $(\mathbf{x}_i^{(\ell)})_{0 \leq i < U}$ for $\ell \notin C$.

- 1) Sample $\mu_{\mathbf{x}_i}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{x}_i^{(\ell)}, s_1)$, $\mathbf{r}_i^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_1}$ for $0 \leq i < U$, $\ell \notin C$.
- 2) Sample $\text{ct}_i^{(\ell)} \leftarrow \mathcal{U}(R_q^2)$ for $0 \leq i < U$, $\ell \notin C$.
- 3) Sample $\mathbf{y}_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$, $\mu_{\mathbf{y}_j, i}^{(\ell)} \leftarrow \text{RandPack}(-\mathbf{W}_{i,j} \mathbf{x}_i^{(\ell)}, s_2)$, $\mu_{\mathbf{y}_j, U}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{y}_j^{(\ell)} + \sum_{i=0}^{U-1} \mathbf{W}_{i,j} \mathbf{x}_i^{(\ell)}, s_2)$, $\mathbf{s}_{j,i}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$, and $\mathbf{s}_{j,U}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$ for $0 \leq i < U$, $0 \leq j < V$, $\ell \notin C$, where $\mathbf{W}_{i,j}$ is the negacyclic matrix of $w_{i,j}$.

- 4) Set $\mu_{\mathbf{z}_j}^{(\ell)} = \mu_{\mathbf{y}_j, U}^{(\ell)} + \sum_{i=0}^{U-1} (\mu_{\mathbf{y}_j, i}^{(\ell)} + w_{i,j} \cdot \mu_{\mathbf{x}_i}^{(\ell)})$, $\mathbf{t}_j^{(\ell)} = \mathbf{s}_{j,U}^{(\ell)} + \sum_{i=0}^{U-1} (\mathbf{s}_{j,i}^{(\ell)} + w_{i,j} \cdot \mathbf{r}_i^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 5) Set $\text{comm}_j^{(\ell)} = \text{Enc}(\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{t}_j^{(\ell)}) - \sum_{i=0}^{U-1} w_{i,j} \cdot \text{ct}_i^{(\ell)} \pmod{q}$ and $\text{resp}_j^{(\ell)} = (\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{s}_j^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 6) Output $(\text{ct}_i^{(\ell)})_{0 \leq i < U}$, $(\text{comm}_j^{(\ell)})_{0 \leq j < V}$, and $(\text{resp}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \notin C$.

\mathcal{H}_2

The inputs are a challenge $(w_{i,j})_{0 \leq i < U, 0 \leq j < V}$, and a set $C \subset [n]$ for $\ell \notin C$.

- 1) Sample $\mu_{\mathbf{x}_i}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{0}, s_1)$, $\mathbf{r}_i^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_1}$ for $0 \leq i < U$, $\ell \notin C$.
- 2) Set $\text{ct}_i^{(\ell)} \leftarrow \mathcal{U}(R_q^2)$ for $0 \leq i < U$, $\ell \notin C$.
- 3) Sample $\mathbf{y}_j^{(\ell)} \leftarrow \mathcal{U}(\mathbb{Z}_p^D)$, $\mu_{\mathbf{y}_j, i}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{0}, s_2)$, $\mu_{\mathbf{y}_j, U}^{(\ell)} \leftarrow \text{RandPack}(\mathbf{y}_j^{(\ell)}, s_2)$, $\mathbf{s}_{j,i}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$, and $\mathbf{s}_{j,U}^{(\ell)} \leftarrow \mathcal{D}_{\mathbb{Z}^N, \sigma_2}$ for $0 \leq i < U$, $0 \leq j < V$, $\ell \notin C$.
- 4) Set $\mu_{\mathbf{z}_j}^{(\ell)} = \mu_{\mathbf{y}_j, U}^{(\ell)} + \sum_{i=0}^{U-1} (\mu_{\mathbf{y}_j, i}^{(\ell)} + w_{i,j} \cdot \mu_{\mathbf{x}_i}^{(\ell)})$, $\mathbf{t}_j^{(\ell)} = \mathbf{s}_{j,U}^{(\ell)} + \sum_{i=0}^{U-1} (\mathbf{s}_{j,i}^{(\ell)} + w_{i,j} \cdot \mathbf{r}_i^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 5) Set $\text{comm}_j^{(\ell)} = \text{Enc}(\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{t}_j^{(\ell)}) - \sum_{i=0}^{U-1} w_{i,j} \cdot \text{ct}_i^{(\ell)} \pmod{q}$ and $\text{resp}_j^{(\ell)} = (\mu_{\mathbf{z}_j}^{(\ell)}, \mathbf{s}_j^{(\ell)})$ for $0 \leq j < V$, $\ell \notin C$.
- 6) Output $(\text{ct}_i^{(\ell)})_{0 \leq i < U}$, $(\text{comm}_j^{(\ell)})_{0 \leq j < V}$, and $(\text{resp}_j^{(\ell)})_{0 \leq j < V}$ for $\ell \notin C$.

Figure 8: Hybrid arguments for the simulatability of Π_{PoPK}

B.7. Noise Bound of $\text{CT}_{\alpha \cdot \iota(\mathbf{C}_i)}$

We begin by analyzing the encryption noise. As analyzed in Lemma 5, the initial encryption noise is dominated by the term $\frac{b+1}{q} \cdot N\sqrt{3N} \cdot \|\mu\|_\infty$. Thus, we bound the noise for $\text{CT}_{\iota(\mathbf{A}_i)} = (\text{ct}_{\iota(\mathbf{A}_i)_j})_{0 \leq j < D}$ and $\text{CT}_{\tilde{\iota}(\mathbf{B}_i)} = (\text{ct}_{\tilde{\iota}(\mathbf{B}_i)_j})_{0 \leq j < D}$ as follows.

$$\|\text{Err}(\text{ct}_{\iota(\mathbf{A}_i)_j}, \iota(\mathbf{A}_i)_j)\|_\infty^{\text{can}}, \|\text{Err}(\text{ct}_{\tilde{\iota}(\mathbf{B}_i)_j}, \tilde{\iota}(\mathbf{B}_i)_j)\|_\infty^{\text{can}} \lesssim \frac{b+1}{q} \cdot 2N^2\sqrt{3N} \cdot B_z$$

This holds because Π_{PoPK} provides an upper bound for plaintexts as $2NB_z$. Next, we consider the noise growth from MatMul . As analyzed in Lemma 10, the dominant term is $d(b+1) \cdot \sqrt{2N^2 + 3N}(\mathbf{e}_A + \mathbf{e}_B)$, so we bound noise for $\text{CT}_{\iota(\mathbf{C}_i)} = (\text{ct}_{\iota(\mathbf{C}_i)_j})_{0 \leq j < D}$ as follows since $\text{CT}_{\iota(\mathbf{C}_i)} \leftarrow \text{MatMul}(\text{CT}_{\iota(\mathbf{A}_i)}, \text{CT}_{\tilde{\iota}(\mathbf{B}_i)})$.

$$\|\text{Err}(\text{ct}_{\iota(\mathbf{C}_i)_j}, \iota(\mathbf{C}_i)_j)\|_\infty^{\text{can}} \lesssim \frac{d(b+1)^2}{q} \cdot 4N^3\sqrt{6N} \cdot B_z$$

Lastly, we consider the noise growth from Mul when multiplying ct_α to $\text{CT}_{\iota(\mathbf{C}_i)}$. As analyzed in Lemma 8, the dominant term is $(b+1)\sqrt{2N^2 + 3N}(\|\text{Err}(\text{ct}, \mu)\|_\infty^{\text{can}} + \|\text{Err}(\text{ct}', \mu')\|_\infty^{\text{can}})$, so the noise for $\text{CT}_{\alpha \cdot \iota(\mathbf{C}_i)} = (\text{ct}_{\alpha \cdot \iota(\mathbf{C}_i)_j})_{0 \leq j < D}$ is bounded as follows, which becomes the upper bound B_{Dec} for ciphertext noise in Π_{Triple} .

$$\|\text{Err}(\text{ct}_{\alpha \cdot \iota(\mathbf{C}_i)_j}, \alpha \cdot \iota(\mathbf{C}_i)_j)\|_\infty^{\text{can}} \lesssim \frac{d(b+1)^3}{q} \cdot 8N^4\sqrt{3N} \cdot B_z = B_{\text{Dec}}$$