

Cloud-Assisted Multiparty Computation from Fully Homomorphic Encryption

Adriana López-Alt
New York University

Eran Tromer*
Tel-Aviv University

Vinod Vaikuntanathan†
University of Toronto

December 7, 2011

Abstract

We construct protocols for secure multiparty computation with the help of a computationally powerful party, namely the “cloud”. Our protocols are simultaneously efficient in a number of metrics:

- *Rounds*: our protocols run in 4 rounds in the semi-honest setting, and 5 rounds in the malicious setting.
- *Communication*: the number of bits exchanged in an execution of the protocol is *independent of the complexity of function f* being computed, and depends only on the length of the inputs and outputs.
- *Computation*: the computational complexity of all parties is *independent of the complexity of the function f* , whereas that of the cloud is linear in the size of the circuit computing f .

In the semi-honest case, our protocol relies on the “ring learning with errors” (RLWE) assumption, whereas in the malicious case, security is shown under the Ring LWE assumption as well as the existence of simulation-extractable NIZK proof systems and succinct non-interactive arguments. In the malicious setting, we also relax the communication and computation requirements above, and only require that they be “small” – polylogarithmic in the computation size and linear in the size of the joint size of the inputs.

Our constructions leverage the *key homomorphic* property of the recent fully homomorphic encryption scheme of Brakerski and Vaikuntanathan (CRYPTO 2011, FOCS 2011). Namely, these schemes allow combining encryptions of messages under different keys to produce an encryption (of the sum of the messages) under the sum of the keys. We also design an efficient, non-interactive *threshold decryption* protocol for these fully homomorphic encryption schemes.

*This work was partially supported by the Check Point Institute for Information Security and by the Israeli Centers of Research Excellence (I-CORE) program (center No. 4/11).

†This work was partially supported by an NSERC Discovery Grant, by DARPA under Agreement number FA8750-11-2-0225. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 1.1 | Our Results | 2 |
| 1.2 | Other Related Work | 3 |
| 2 | Preliminaries | 3 |
| 3 | Cloud-Assisted MPC | 4 |
| 4 | Cloud-Assisted MPC From Key-Homomorphic Threshold FHE | 6 |
| 4.1 | Construction | 9 |
| 4.2 | Security Under Semi-Honest Adversaries | 10 |
| 5 | Tolerating Malicious Adversaries | 13 |
| 6 | Instantiating Key-Homomorphic Threshold FHE | 16 |
| 6.1 | The Polynomial-LWE Assumption | 16 |
| 6.2 | A Somewhat-Homomorphic Scheme | 17 |
| 6.3 | Getting a (Leveled) Fully-Homomorphic Scheme via Re-linearization and Modulus Reduction | 18 |
| 6.3.1 | Re-linearization | 18 |
| 6.3.2 | Modulus Reduction | 20 |
| 6.4 | Key-Homomorphic Threshold FHE from PLWE | 21 |
| 7 | Implementing $\mathcal{F}_{\text{rand.bc}}^{\mathcal{E}}$ and $\mathcal{F}_{\text{keys}}^{\mathcal{E}}$ | 24 |
| 7.1 | Ciphertext Re-Randomization | 24 |
| 7.2 | Key Combination | 25 |
| 8 | Performance | 26 |
| A | Proofs | 31 |
| A.1 | Security Under Malicious Adversaries | 31 |
| A.2 | Ciphertext Re-Randomization | 38 |
| A.3 | Key Combination | 41 |

1 Introduction

Cloud computing is usually portrayed as a user outsourcing its computation to a provider. However, in reality each cloud provider has numerous users, who may wish to compute with each other. How can multiple cloud users conduct an expensive joint computation, when they wish to keep private all information except the final output? They could send all their inputs to the cloud and ask it to compute the output, but this would violate their confidentiality. Alternatively, they could conduct a secure multiparty computation protocol (MPC) [GMW87, BGW88, CCD88] among themselves, and bear the full burden of computation. In particular, standard MPC protocols incur a large blow-up in the computational complexity of all the players involved, as well as a commensurate increase in the communication complexity of the protocol, neither of which are acceptable in the client-server setting. Can the presence of the cloud make the multiparty computation more efficient, without compromising its security?

More precisely, what we would like is an MPC protocol where n parties contribute inputs x_1, \dots, x_n respectively, and they jointly compute $f(x_1, \dots, x_n)$ securely satisfying the following two properties:

- **Computational efficiency:** the running time of each party is polynomial in the length of its input and output. In particular, it should be independent of the complexity of the function f .
- **Communication efficiency:** the communication complexity of the protocol should be polynomial in the length of the inputs and outputs, but otherwise independent of the complexity of f .

Clearly, achieving computational efficiency alone requires the parties to seek the help of a (powerful) cloud capable of computing f . The protocol should be secure against a collusion of an arbitrary subset of parties together with the cloud.

The first possibility that comes to mind is to use a fully homomorphic encryption (FHE) scheme, which provides a mechanism to design communication-efficient protocols in the two-party setting. Starting from the groundbreaking work of Gentry [Gen09b], a number of progressively more efficient and simpler FHE constructions have surfaced [vDGHV10, SV10], most recently the works of Brakerski and Vaikuntanathan [BV11b, BV11a] and the work of Brakerski, Gentry and Vaikuntanathan [BGV11]. In the context of a single user outsourcing her computation to the cloud, the user can encrypt her input using her key under the fully homomorphic encryption scheme, and send the ciphertext to the server. The server homomorphically evaluates the desired function and returns the evaluated ciphertext back to the user who decrypts and gets the result of the computation.¹ Can we use fully homomorphic encryption in the *multiparty* setting, where the computation is performed jointly on the inputs of many players? A naïve way to approach this problem runs into the difficulty that we know of no way to operate on ciphertexts encrypted under different keys.

Gentry suggested an approach to do communication- and computation-efficient cloud-assisted MPC in the following way. The parties first run a multiparty coin-tossing protocol to generate a joint public key for (an arbitrary) FHE scheme, together with a secret sharing of the corresponding secret key (the i^{th} party gets the i^{th} share of the common secret key). Once this is done, the parties encrypt their inputs using the *common public key* and have the cloud compute an encryption of the result. They then run yet another MPC protocol to perform threshold decryption and recover the result. Clearly, since both the coin-tossing and the threshold decryption functionality are independent of the function f , this protocol is communication and computation-efficient (in the sense above). However, due to the use of generic MPC techniques, the protocol has a fairly large round-complexity and concrete computational complexity.

¹For this discussion, we assume that the adversaries are semi-honest.

Can this state of affairs be improved, by taking advantage of the algebraic structure that underlies many of the recent FHE schemes? We show this is indeed the case, by demonstrating concrete constructions.

1.1 Our Results

We show MPC protocols that are efficient in communication and computation as well as round complexity, using a special property of recent FHE constructions, namely *key homomorphism*. Informally, a key-homomorphic (public key) encryption scheme allows us to *deterministically* combine many public keys $\mathbf{pk}_1, \dots, \mathbf{pk}_N$ into a combined public key \mathbf{pk} , and simultaneously combine secret keys $\mathbf{sk}_1, \dots, \mathbf{sk}_N$ into a combined secret key \mathbf{sk} . Informally, this property allows the parties to come up with a joint public key under which they can all encrypt their respective inputs, without resorting to an expensive coin-tossing protocol as in Gentry’s solution [Gen09b, Gen09a]. Of course, for security, we require that semantic security under the joint public key \mathbf{pk} holds even given $t < N$ secret keys \mathbf{sk}_i . In a sense, \mathbf{sk}_i form the secret shares of a common secret key \mathbf{sk} whose corresponding public key is \mathbf{pk} . The novelty is that a key homomorphic encryption scheme allows us to come up with this structure in a rather natural and efficient way.

Key-homomorphic encryption schemes abound. For example, consider Regev’s public key encryption scheme [Reg09] where the public key $\mathbf{pk} = \mathbf{A}^\top \mathbf{s} + x$, where \mathbf{A} is a matrix that forms a global system parameter, and the secret key vector $\mathbf{sk} = \mathbf{s}$. Adding together two such public keys $\mathbf{pk}_1 = \mathbf{A}^\top \mathbf{s}_1 + x_1$ and $\mathbf{pk}_2 = \mathbf{A}^\top \mathbf{s}_2 + x_2$ gives us $\mathbf{pk} = \mathbf{pk}_1 + \mathbf{pk}_2 = \mathbf{A}^\top (\mathbf{s}_1 + \mathbf{s}_2) + (x_1 + x_2)$, a joint public key whose corresponding secret key \mathbf{sk} is the sum of \mathbf{sk}_1 and \mathbf{sk}_2 . This structure is by no means unique to lattice based schemes, but it appears in numerous number-theoretic systems as well, e.g., the El Gamal encryption scheme. In lattice schemes, they are especially useful since they give us fully homomorphic encryption schemes that are key-homomorphic as well.

As we argued before, key-homomorphic schemes provide us with a natural avenue to remove the coin-tossing protocol from Gentry’s schema for multiparty computation from FHE. In the context of FHE schemes, we also need a method of combining evaluation keys to produce a joint evaluation key which turns out to be trickier. In particular, we do not know of a nice non-interactive way of combining evaluation keys, but instead, we come up with a two-round protocol to do this.

Finally, the linear structure of these schemes give us a very efficient way to do *non-interactive threshold decryption*, similar to the work of Bendlin and D amgaard [BD10].

These ideas are fairly general and can be instantiated with a number of different combinations of the existing schemes. We choose to present the instantiation that combines the ring LWE based somewhat homomorphic encryption scheme from [BV11b] with the re-linearization technique of [BV11a] and the noise management technique of [BGV11] (which, in turn, relies on the modulus reduction technique of [BV11a]). We chose this instantiation since it is quite efficient and simplest to present. The reader is referred to Section 6 for details on the instantiations. The resulting protocol is secure against semi-honest adversaries that corrupt any $t < N$ parties, and runs in 4 rounds.

We then compile this protocol to be secure against malicious adversaries, using simulation-extractable NIZK proof systems [SP92, Sah99, CLOS02] as well as succinct non-interactive argument systems [Mic94, GKR08, BCCT11, GLR11]. The resulting protocol runs in 5 rounds.

Our main theorem is the following:

Theorem 1.1. *There are communication- and computation-efficient MPC protocols that are:*

- *secure against semi-honest corruptions of an arbitrary subset of players, under the Ring LWE assumption and runs in 4 rounds. The communication and computation of the parties is inde-*

pendent of the complexity of f , and the computation time of the server is linear in the size of the circuit computing f .

- *secure against malicious adversaries that corrupt an arbitrary number of players, under the Ring LWE assumption, the existence of simulation-extractable NIZK proof systems and succinct non-interactive argument systems. The protocol runs in 5 rounds. The communication of the parties is polylogarithmic in the size of the circuit computing f , their computation time is nearly linear in the total size of the inputs, and the computation time of the server is polynomial in the size of the circuit computing f .*

1.2 Other Related Work

The basic idea of using homomorphic encryption schemes in conjunction with threshold decryption to boost the efficiency of MPC protocols was first noticed by Cramer, Damgård and Nielsen [CDN01]. The idea of using a cloud to alleviate the computational efforts of parties was recently explored in the work on "server-aided MPC" by Kamara, Mohassel and Raykova [KMR11]. Their protocols, however, require some of the parties to do a large amount of computation, essentially proportional to the size of the function f being computed. Thus, their protocols are not computation-efficient in our sense. The works of Myers, Sergi and Shelat [MSas11], Bendlin et.al [BDOZ11], and Damgård et. al. [DPSZ11] construct MPC protocols (in the classical sense, without the help of a cloud) using FHE schemes (in the case of [MSas11]) and somewhat homomorphic schemes (in [BDOZ11, DPSZ11]). Their protocols are communication-efficient, but not computation-efficient in our sense, namely the computational complexity of the parties is proportional to the size of f . Halevi, Lindell and Pinkas [HLP11] recently considered the model of "secure computation on the web" wherein the goal is to minimize interaction between the parties. Unfortunately, as they show, their notion can only be achieved for a small class of functions. In contrast, we focus here on designing MPC protocols for arbitrary functions, of course at the cost of interaction.

2 Preliminaries

NOTATION. For an integer n , we use the notation $[n]$ to denote the set $[n] = \{1, \dots, n\}$. For a randomized function f , we write $f(x; r)$ to denote the unique output of f on input x with random coins r . We write $f(x)$ to denote a random variable for the output of $f(x; r)$ over the random coins r . For a distribution or random variable X , we write $x \leftarrow X$ to denote the operation of sampling a random x according to X . For a set S , we overload notation and use $s \leftarrow S$ to denote sampling s from the uniform distribution over S . For two distributions, X and Y , we use $X \stackrel{c}{\approx} Y$ to mean that X and Y are computationally indistinguishable, and $X \stackrel{s}{\approx} Y$ to mean that they are statistically close. Finally, we use $y := f(x)$ to denote the *deterministic* evaluation of f on input x with output y .

FULLY HOMOMORPHIC ENCRYPTION. We review the definition of (leveled) fully homomorphic encryption. We state the definition of [BGV11], which is a relaxation of the original definition in [Gen09b]. The main difference is that the definition in [Gen09b] requires that schemes $\mathcal{E}^{(D)}$ that evaluate circuits of depth at most D , all share the same decryption circuit. However, as in [BGV11], we relax this definition and let all algorithms (including decryption) depend on the depth D of the circuit that is evaluated.

Definition 2.1 (*C*-Homomorphic Encryption [Gen09b]). *For a class of circuits \mathcal{C} , we say that an encryption scheme $\mathcal{E} = (\text{Setup}, \text{Keygen}, \text{Enc}, \text{Dec}, \text{Eval})$ is \mathcal{C} -homomorphic if:*

- $\text{params} \leftarrow \text{Setup}(1^\kappa)$, for security parameter κ , outputs public parameters params . All other algorithms, Enc , Dec , Eval implicitly take params as input, even when not explicitly stated.
- $(\text{pk}, \text{sk}, \text{ek}) \leftarrow \text{Keygen}(\text{params})$, for a security parameter κ , outputs a public key pk , a secret key sk , and a (public) evaluation key ek .
- $m' \leftarrow \text{Dec}_{\text{sk}}(c)$, given a secret key sk and a ciphertext c , outputs a message m' .
- $\gamma := \text{Eval}(C, c_1, \dots, c_\ell)$, given a (description of) a circuit C along with an evaluation key pk and ℓ ciphertexts c_1, \dots, c_ℓ , outputs a ciphertext γ .

We require that for all $c \in \mathcal{C}$, all $(\text{pk}, \text{sk}, \text{ek}) \leftarrow \text{Keygen}(1^\kappa)$ and all plaintexts (m_1, \dots, m_ℓ) and ciphertexts (c_1, \dots, c_ℓ) such that c_i is in the support of $\text{Enc}_{\text{pk}}(m_i)$, if $\gamma := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_\ell)$, then $\text{Dec}_{\text{sk}}(\gamma) = C(m_1, \dots, m_\ell)$.

Definition 2.2 (Leveled Fully Homomorphic Encryption [Gen09b]). Let $\mathcal{C}^{(D)}$ be the class of all circuits of depth at most D (that use some specified complete set of gates). We say that a family of homomorphic encryption schemes $\{\mathcal{E}^{(D)} : D \in \mathbb{Z}^+\}$ is leveled fully homomorphic if, for all $D \in \mathbb{Z}^+$, it satisfies the following properties:

Correctness: $\mathcal{E}^{(D)}$ is $\mathcal{C}^{(D)}$ -homomorphic.

Compactness: The computational complexity of $\mathcal{E}^{(D)}$'s algorithms is polynomial (the same polynomial for all D) in the security parameter, D , and (in the case of the evaluation algorithm) the size of the circuit.

From now on, when we say *fully homomorphic*, we refer to *leveled fully homomorphic encryption*.

FUNCTIONALITIES. We review the following basic MPC functionalities, which we will use when describing our cloud-assisted MPC protocol.

- \mathcal{F}_f , the functionality computing a function f . See Figure 1.
- $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$, the CRS functionality that outputs a common reference string distributed according to a predetermined distribution \mathcal{D} . See Figure 2.
- \mathcal{F}_{BC} , the Broadcast functionality that upon receiving a message from a party P_i , forwards the message to *all* parties. See Figure 3.
- $\mathcal{F}_{\text{ZK.BC}}^R$, the “Prove-and-Broadcast” functionality $\mathcal{F}_{\text{ZK.BC}}^R$ that upon receiving a statement and witness pair (x, w) from a party P_i , verifies that $R(x, w) = 1$ (for some predetermined NP-relation R) and if so, sends x to all parties. See Figure 4.




Figure 1: Functionality \mathcal{F}_f

3 Cloud-Assisted MPC

We consider the problem of N parties who wish to securely compute a joint function of their inputs, but want to do so in a way such that the amount of data they interchange, as well as their computation

Functionality $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$

$\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ interacts with parties P_0, P_1, \dots, P_N and adversary \mathcal{S} , and is parametrized by an algorithm \mathcal{D} .
Upon receiving (sid, i) from P_i , compute $\text{crs} \leftarrow \mathcal{D}(1^\kappa)$, and send $(\text{sid}, i, \text{crs})$ to P_i and the adversary. Upon receiving subsequent messages (sid, j) from P_j , send $(\text{sid}, j, \text{crs})$ to P_j and the adversary \mathcal{S} , and halt.

Figure 2: CRS Functionality $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$

Functionality \mathcal{F}_{BC}

\mathcal{F}_{BC} interacts with parties P_0, P_1, \dots, P_N and adversary \mathcal{S} .
Upon receiving (sid, i, m) from P_i , send (sid, i, m) to all parties in P_1, \dots, P_N and the adversary \mathcal{S} , and halt.

Figure 3: Broadcast Functionality \mathcal{F}_{BC}

Functionality $\mathcal{F}_{\text{ZK.BC}}^R$

$\mathcal{F}_{\text{ZK.BC}}^R$ interacts with parties P_0, P_1, \dots, P_N and adversary \mathcal{S} , and is parametrized by an NP-relation R .
Upon receiving (sid, i, x, w) from P_i : If $R(x, w) = 1$, send (sid, i, x) to all parties and the adversary \mathcal{S} , and halt. Otherwise, halt.

Figure 4: “Prove-and-Broadcast” Functionality $\mathcal{F}_{\text{ZK.BC}}^R$

time, is independent of the complexity of the function. In order to do this, we rely on the computation power of a new party called the *server* or *cloud*, who will perform the computation of the function f . The computation should remain secure even though the server is untrusted. We formalize this idea below.

DEFINITION. For an N -input function f , we define a *cloud-assisted multiparty protocol* Π for f to be a protocol between N interactive Turing Machines P_1, \dots, P_N , called *clients*, and a *server* S , also an interactive Turing Machine, such that for all $\vec{x} = (x_1, \dots, x_N)$, the output of Π in an execution where P_i is given x_i as input (and the server S does not receive an input), is $y = f(\vec{x})$. We sometimes use P_0 to denote the server S . What differentiates a cloud-assisted MPC protocol from a standard MPC protocol is that we require the communication complexity of the protocol, as well as the computation time of the clients P_1, \dots, P_N to be *independent* of the complexity of the function f . On the other hand, the computation time of the server S must be *linear* in the circuit-size of f .²

SECURITY. We prove security of a cloud-assisted MPC protocol in the Ideal/Real paradigm (in the standalone setting). In the *ideal world*, the computation of f is performed through a trusted functionality \mathcal{F}_f that receives input x_i from each party P_i , computes $y = f(x_1, \dots, x_N)$ and gives y to all parties P_1, \dots, P_N , and the server S . It is clear that in the ideal world, the only information that any party learns is its own input and the output y . In the *real world*, however, this trusted functionality does not exist and so in order to compute $y = f(x_1, \dots, x_N)$, parties P_1, \dots, P_N and

²When considering security against semi-honest adversaries, we achieve precisely this. However, to tolerate malicious adversaries, we relax the definition and require only that the communication and computation time of the clients be “small”, much less than linear in the size of the circuit they are computing. Furthermore, we let the computation time of the server be polynomial in the size of the circuit.

server S run a protocol Π . Figure 5 gives a pictorial representation of the ideal world and the real world.

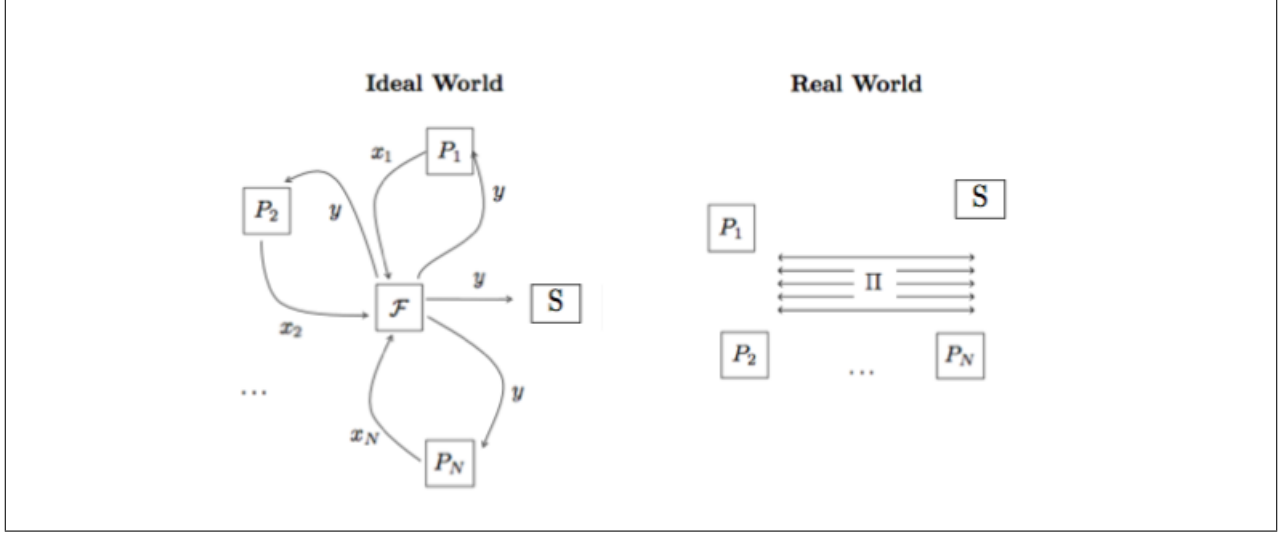


Figure 5: Ideal World vs. Real World

In either world, an adversary corrupting $t < N$ parties, and perhaps additionally the server S , receives all messages directed to the corrupted parties and controls the messages sent by them. We use $\text{IDEAL}_{\mathcal{F},\mathcal{S}}(\vec{x})$ to denote the joint output of the ideal-world adversary \mathcal{S} and of parties P_1, \dots, P_N and server P_0 in an ideal execution with functionality \mathcal{F} and inputs $\vec{x} = (x_1, \dots, x_N)$. Similarly, we use $\text{REAL}_{\Pi,\mathcal{A}}(\vec{x})$ to denote the joint output of the real-world adversary \mathcal{A} and the outputs of parties P_1, \dots, P_N and server P_0 in an execution of protocol Π with inputs $\vec{x} = (x_1, \dots, x_N)$. We say that a protocol Π *securely realizes* \mathcal{F} if for every real-world adversary \mathcal{A} , there exists an ideal-world adversary \mathcal{S} such that

$$\text{IDEAL}_{\mathcal{F},\mathcal{S}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi,\mathcal{A}}(\vec{x}),$$

4 Cloud-Assisted MPC From Key-Homomorphic Threshold FHE

In this section, we define a new primitive called *key-homomorphic threshold FHE* and show that it can be used to construct a cloud-assisted MPC protocol for any function f . Informally, a key-homomorphic threshold FHE scheme allows combining public keys $\text{pk}_1, \dots, \text{pk}_N$ into a combined public key pk , evaluation keys $\text{ek}_1, \dots, \text{ek}_N$ into a combined evaluation key ek , and secret keys $\text{sk}_1, \dots, \text{sk}_N$ into a combined secret key sk . The correctness property we require for this key combination is that if for all i , $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is in the support of Keygen , then $(\text{pk}, \text{sk}, \text{ek})$ is also in the support of Keygen . In other words, if for all i , pk_i and ek_i are *valid* public and evaluation keys for secret key sk_i , then pk and ek are valid public and evaluation keys for sk . For security, we require semantic security under the joint public key pk even given $t < N$ secret keys sk_i . We additionally require that the combined evaluation key ek does not reveal any information about the individual evaluation and secret keys ek_i, sk_i . We formalize this by requiring that there exist an algorithm SimCombKeys that given any number of public/evaluations key pairs $\{(\text{pk}_i, \text{ek}_i)\}_{i \in T}$, can simulate the remaining public/evaluations key pairs $\{(\text{pk}_j, \text{ek}_j)\}_{j \in [N] \setminus T}$ and the joint evaluation key ek .

A key-homomorphic threshold FHE scheme also has threshold decryption properties that allow for decryption of a ciphertext under the combined public key pk by independently using the secret keys sk_i . More formally, given a ciphertext c encrypting a plaintext m under pk , each party P_i can independently

compute a decryption share μ_i using sk_i . Combining the ciphertext c with the shares μ_1, \dots, μ_N will yield the plaintext m . For security, we require that the share μ_i does not reveal anything about the secret key sk_i . We formalize this by requiring that there exist an algorithm SimShareDec that given a ciphertext c , a plaintext m , and any number of shares $\{\mu_i\}_{i \in T}$, can simulate the remaining shares $\{\mu_j\}_{j \in [N] \setminus T}$, such that when all shares are combined, they decrypt c to m .

Finally, we require ciphertexts of a key-homomorphic threshold FHE scheme to be re-randomizable. In other words, we require that there exists an algorithm ReRand that given the public key pk and ciphertext $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_\ell)$, where c_1, \dots, c_ℓ are ciphertexts encrypting plaintexts m_1, \dots, m_ℓ under pk , outputs a randomization \hat{c} of c that is indistinguishable from a *fresh* encryption of $y = C(m_1, \dots, m_\ell)$, for all circuits C . We require this property because we can only guarantee security of threshold decryption (i.e., that the shares do not reveal information about the individual secret keys) if the ciphertext c to be decrypted is a fresh encryption of y . Unfortunately, since our instantiation of key-homomorphic threshold FHE is lattice-based, the only way we know how to randomize ciphertexts is by adding to c noise that is super-polynomially larger than the noise in the original ciphertexts c_i . Because of this, we need to define the encryption algorithm in key-homomorphic threshold FHE to take an extra bit $\text{mode} \in \{0, 1\}$ that decides the magnitude of the noise to be used in the encryption. There is no inherent reason why we need to work over two types of ciphertexts, this is just a technical artifact of our instantiation.

Our instantiation will also require the use of a common reference string, crs , that will be partitioned into N parts, $(\text{crs}_1, \dots, \text{crs}_N)$, so that crs_i corresponds to party P_i . The public and evaluation keys for P_i will depend on crs_i . We model this by having the Setup algorithm take an additional input N and output crs as part of the public parameters, and having Keygen take an additional input, $i \in [N]$. Again, there is no inherent reason why we need these additional inputs to Setup and Keygen , this is just a technical artifact of our instantiation.

Definition 4.1. A public-key key-homomorphic threshold FHE scheme \mathcal{E} with secret, public, and evaluation key-spaces SK, PK, EK and message-space \mathcal{M} is a tuple of algorithms $(\text{Setup}, \text{Keygen}, \text{Enc}, \text{Dec}, \text{Eval}, \text{CombinePK}, \text{CombineKeys}, \text{ShareDec}, \text{ShareCombine}, \text{ReRand})$ with the following syntax:

- $\text{Setup}, \text{Keygen}, \text{Dec}$ and Eval are as in a standard FHE. We only require that Setup takes an additional input N and Keygen takes an additional input i .

From here on, all algorithms implicitly take as input params , where $\text{params} \leftarrow \text{Setup}(1^\kappa, N)$, even if not explicitly stated. $\text{Keygen}, \text{Dec}$ and Eval also implicitly take params as input.

- $c \leftarrow \text{Enc}_{\text{pk}}(m, \text{mode})$. Takes as input a public key pk , a message $m \in \mathcal{M}$, and a bit $\text{mode} \in \{0, 1\}$, and outputs a ciphertext c . We will work with two types of ciphertexts; the bit mode will decide what type of ciphertext to output. Most of the time we will work with ciphertexts generated by $\text{Enc}(\cdot, 0)$, and will only use $\text{Enc}(\cdot, 1)$ in re-randomization (see below). When unspecified, we assume $\text{mode} = 0$.
- $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$. Takes as input N public keys $\{\text{pk}_i\}_{i \in [N]}$, and outputs a combined public key pk .
- $(\text{sk}, \text{ek}) \leftarrow \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]})$. Takes as input N key pairs $\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}$, and outputs combined secret key sk and evaluation key ek .
- $\mu_i \leftarrow \text{ShareDec}_{\text{sk}_i}(c)$. Takes as input secret key sk_i and a ciphertext c and outputs a decryption share μ_i .
- $m \leftarrow \text{ShareCombine}(c, \mu_1, \dots, \mu_N)$. Takes as input a ciphertext c and N decryption shares (μ_1, \dots, μ_N) and outputs a plaintext m .
- $\hat{c} \leftarrow \text{ReRand}_{\text{pk}}(c)$. Takes as input a public key pk and ciphertext c and outputs a rerandomization \hat{c} of c .

Correctness and security are given the following properties. Define the relation

$$R^{(keys)} = \left\{ (\text{params}, \text{pk}_i, \text{ek}_i, i), (\text{sk}_i, r_i^{(gen)}) \mid (\text{pk}_i, \text{sk}_i, \text{ek}_i) := \text{Keygen}(\text{params}, i; r_i^{(gen)}) \right\}$$

and say that the public/evaluation key pair $(\text{pk}_i, \text{ek}_i)$ is valid for secret key sk_i with respect to params , if there exists $r_i^{(gen)}$ such that $R^{(keys)}((\text{params}, \text{pk}_i, \text{ek}_i, i), (\text{sk}_i, r_i^{(gen)})) = 1$. We say that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple with respect to params if $(\text{pk}_i, \text{ek}_i)$ is a valid public/evaluation key pair for sk_i with respect to params . When params is clear from context, we simply say that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple.

Correct Key Combination: If $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in [N]}$ are all valid key tuples, $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$ and $(\text{sk}, \text{ek}) \leftarrow \text{CombineKeys}(\{\text{sk}_i\}_{i \in [N]}, \{\text{ek}_i\}_{i \in [N]})$ then $(\text{pk}, \text{sk}, \text{ek})$ is a valid key tuple.

Key-Simulation Indistinguishability: There exists a PPT simulator SimCombKeys such that for any set $T \subsetneq [N]$ with $t = |T|$ and $\bar{T} = [N] \setminus T$, $(\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}) \leftarrow \text{SimCombKeys}(\{\text{pk}_i, \text{ek}_i\}_{i \in T})$ takes as input t public/evaluation key pairs $\{\text{pk}_i, \text{ek}_i\}_{i \in T}$ and outputs $N - t$ public/evaluation key pairs $\{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}$ and an evaluation key ek , satisfying the following.

For all valid key tuples $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in [N]}$ and all subsets $T \subsetneq [N]$, the following two distributions are computationally indistinguishable:

$$(\text{params}, \text{SimCombKeys}(\{(\text{pk}_i, \text{ek}_i)\}_{i \in T})) \stackrel{c}{\approx} (\text{params}, \text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}})$$

where $\text{params} \leftarrow \text{Setup}(1^\kappa)$, $\{(\cdot, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa)\}_{j \in \bar{T}}$, and $(\cdot, \text{ek}) \leftarrow \text{CombineKeys}(\{\text{sk}_i, \text{ek}_i\}_{i \in [N]})$.

Semantic Security: For all $T \subsetneq N$, every $\text{mode} \in \{0, 1\}$, and all PPT adversaries \mathcal{A} , the advantage of \mathcal{A} in the following game is negligible (in κ).

Key Generation: The challenger runs $\text{params} \leftarrow \text{Setup}(1^\kappa)$. The adversary \mathcal{A} then sends $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$ to the challenger. The challenger runs $(\text{pk}_j, \text{sk}_j, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)$ for $j \in \bar{T} = [N] \setminus T$, computes $(\text{sk}, \text{ek}) \leftarrow \text{CombineKeys}(\{\text{sk}_i, \text{ek}_i\}_{i \in [N]})$ and $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$, and sends $(\text{params}, \text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}})$ to \mathcal{A} .

Challenge: \mathcal{A} sends plaintexts m_0, m_1 to the challenger, who chooses $b \leftarrow \{0, 1\}$, and sends $c^* = \text{Enc}_{\text{pk}}(m_b, \text{mode})$ to \mathcal{A} .

Output: \mathcal{A} outputs b' .

We define the advantage of \mathcal{A} to be $|\frac{1}{2} - \Pr[b' = b]|$.

Correct Share Decryption of Combined Key: For all valid key tuples $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in [N]}$, and all $c \in \mathcal{C}_{\text{pk}}$, where $\text{pk} = \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$ and \mathcal{C}_{pk} is the ciphertext-space defined by pk ,

$$\text{Dec}_{\text{sk}}(c) = \text{ShareCombine}(c, \text{ShareDec}_{\text{sk}_1}(c), \dots, \text{ShareDec}_{\text{sk}_N}(c))$$

Share-Simulation Indistinguishability: There exists a PPT simulator SimShareDec , where $\{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(c, m, \{\mu_i\}_{i \in T})$ takes as input a ciphertext c , a plaintext m , and t decryption shares $\{\mu_i\}_{i \in T}$ (where $T \subsetneq [N]$ is any subset of cardinality $t < N$) and outputs $N - t$ decryption shares $\{\mu_j\}_{j \in \bar{T}}$, satisfying the following.

For all valid key tuples $\{(\mathbf{pk}_i, \mathbf{sk}_i, \mathbf{ek}_i)\}_{i \in [N]}$, all subsets $T \subsetneq [N]$, all $m \in \mathcal{M}$, and all random strings $\{r_i^{(dec)}\}_{i \in T}$, the following two distributions are computationally indistinguishable:

$$\left(\text{params}, c, \text{SimShareDec}(c, m, \{\mu_i\}_{i \in T})\right) \stackrel{c}{\approx} \left(\text{params}, c, \{\mu_j\}_{j \in \bar{T}}\right),$$

where $\text{params} \leftarrow \text{Setup}(1^\kappa)$, $c \leftarrow \text{Enc}_{\mathbf{pk}}(m)$ and for $i \in T$, $\mu_i := \text{ShareDec}_{\mathbf{sk}_i}(c; r_i^{(dec)})$.

Ciphertext Re-Randomization: For all circuits C , all plaintexts (m_1, \dots, m_ℓ) , all random strings (r_1, \dots, r_ℓ) , the following distributions are statistically close:

$$\left(\text{params}, c_1, \dots, c_t, \text{ReRand}_{\mathbf{pk}}(c)\right) \stackrel{s}{\approx} \left(\text{params}, c_1, \dots, c_t, \text{Enc}_{\mathbf{pk}}(y, 1)\right),$$

where $\text{params} \leftarrow \text{Setup}(1^\kappa)$, $c_i := \text{Enc}_{\mathbf{pk}}(m_i, 0; r_i)$ for all $i \in [\ell]$, $c := \text{Eval}_{\mathbf{ek}}(c_1, \dots, c_\ell)$ and $y = C(m_1, \dots, m_\ell)$.

4.1 Construction

Before presenting our construction of cloud-assisted MPC from key-homomorphic threshold FHE, we describe two functionalities, $\mathcal{F}_{\text{KEYS}}^\mathcal{E}$ (Figure 6) and $\mathcal{F}_{\text{RAND.BC}}^\mathcal{E}$ (Figure 7) computing the **CombineKeys** and **ReRand** algorithms, respectively. For clarity, we will present our construction in the $(\mathcal{F}_{\text{KEYS}}^\mathcal{E}, \mathcal{F}_{\text{RAND.BC}}^\mathcal{E})$ -hybrid model, as this allows us to separate the messages that are sent for the purpose of computing f from the messages sent for the purpose of combining keys and from those sent for re-randomization. In Section 7, we show protocols implementing $\mathcal{F}_{\text{KEYS}}^\mathcal{E}$ and $\mathcal{F}_{\text{RAND.BC}}^\mathcal{E}$ for our specific instantiation of key-homomorphic threshold FHE.

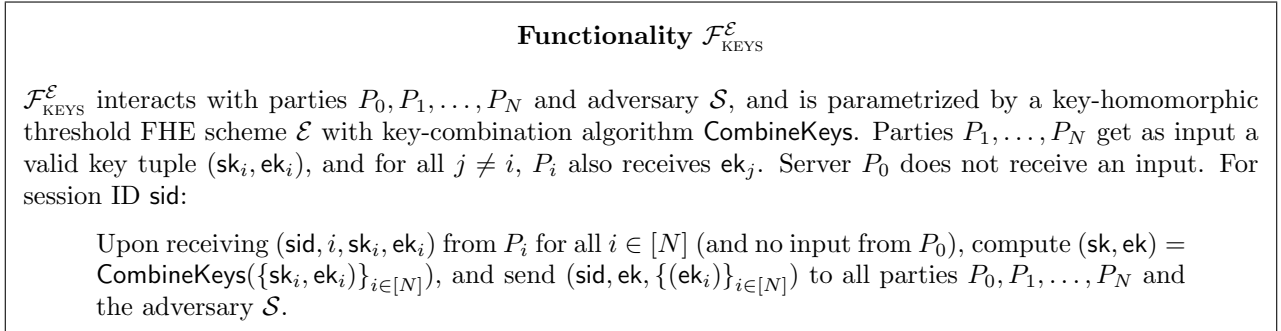


Figure 6: “Key-Combination” Functionality $\mathcal{F}_{\text{KEYS}}^\mathcal{E}$

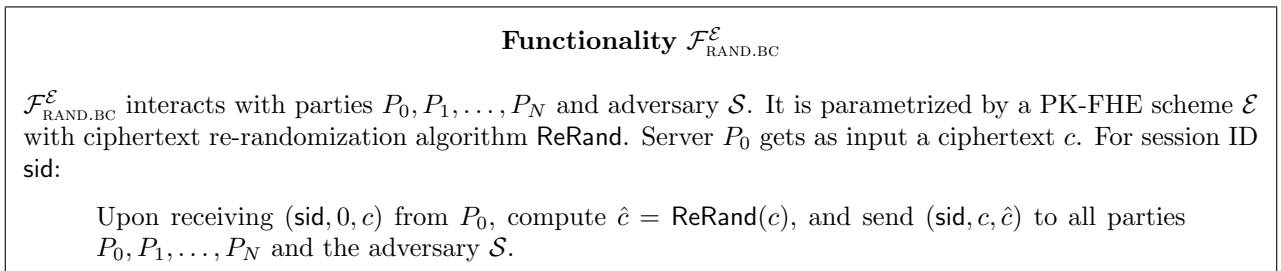


Figure 7: “Ciphertext Re-Randomization” Functionality $\mathcal{F}_{\text{RAND.BC}}^\mathcal{E}$

We describe our cloud-assisted MPC protocol Π^{SH} in detail in Figure 8. The basic idea behind our construction is essentially the same as in [Gen09b]: the parties run an MPC protocol to agree on

a joint public key pk and a joint evaluation key ek . The parties then encrypt their inputs x_i under pk , and the server homomorphically evaluates the function f on the encrypted inputs (x_1, \dots, x_N) using the joint evaluation key ek . Finally, the parties run an MPC protocol to threshold decrypt the output. What key-homomorphic threshold FHE gives us in comparison to the generic approach of [Gen09b] is *round efficiency*. Essentially, the joint public key can be computed deterministically from all public keys, the joint evaluation key ek , obtained from `CombineKeys`, can be computed in 2 rounds, and threshold decryption can be computed in 1 round: the parties simply compute decryption shares μ_i . It is not yet aparent from the description of the protocol why we can achieve only 4 rounds of interaction since we only given a description of our protocol in the $(\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}})$ -hybrid model. We explain this fully in Section 8, but the intuition behind it is that we can send many of the messages for computing f , the joint evaluation key and the re-randomization of c *in parallel*.

Protocol Π^{SH}

Let C be a circuit computing f . Let $\mathcal{E} = (\text{Setup}, \text{Keygen}, \text{Enc}, \text{Eval}, \text{CombinePK}, \text{CombineKeys}, \text{ShareDec}, \text{ShareCombine}, \text{ReRand})$ be a key-homomorphic threshold FHE scheme.

For session ID sid :

Phase 1: Each P_i calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with (sid, i) and receives params . It chooses an encryption key tuple, $(\text{pk}_i, \text{sk}_i, \text{ek}_i) \leftarrow \text{Keygen}(1^\kappa, i)$. P_i then sends $(\text{sid}, i, (\text{sk}_i, \text{ek}_i))$ to $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, and $(\text{sid}, i, \text{pk}_i)$ to \mathcal{F}_{BC} .

Phase 2: Upon receiving $(\text{sid}, \text{ek}, \{\text{ek}_i\}_{i \in [N]})$ from $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, and $(\text{sid}, i, \text{pk}_i)$ from \mathcal{F}_{BC} for all $i \in [N]$, each P_i computes $\text{pk} = \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$, and $c_i \leftarrow \text{Enc}_{\text{pk}}(x_i)$ and sends (sid, i, c_i) to \mathcal{F}_{BC} .

Phase 3: Upon receiving (sid, i, c_i) from \mathcal{F}_{BC} for all $i \in [N]$ and $(\text{sid}, \text{ek}, \{\text{ek}_i\}_{i \in [N]})$ from $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, server computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ and sends $(\text{sid}, 0, c)$ to $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$.

Phase 4: Upon receiving (sid, c, \hat{c}) from $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$, each P_i computes $\mu_i := \text{ShareDec}_{\text{sk}_i}(\hat{c})$ and sends (sid, i, μ_i) to \mathcal{F}_{BC} .

Local Computation: Upon receiving (sid, i, μ_i) from \mathcal{F}_{BC} for all $i \in [N]$, each P_i computes and outputs $y := \text{ShareCombine}(\hat{c}, \mu_1, \dots, \mu_N)$.

Figure 8: Protocol Π^{SH} . Implements \mathcal{F} in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}})$ -hybrid model, where \mathcal{D} is the output distribution of `Setup` $(1^\kappa, N)$. Secure against semi-honest adversaries.

Lemma 4.1. *Let f be an N -input function, and let \mathcal{F} be the ideal functionality computing f . Let \mathcal{E} be a key-homomorphic threshold FHE scheme. Then Π^{SH} described in Figure 8 securely implements \mathcal{F} against (static) semi-honest adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}})$ -hybrid model.*

Looking ahead, we will instantiate key-homomorphic threshold FHE based on the PLWE assumption, a special case of the “Ring LWE” assumption of [LPR10] (see Section 6). Combining Lemma 4.1 with protocols implementing $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ and $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ for this instantiation (shown in Section 7), we arrive at the following theorem.

Theorem 4.2. *Assuming the hardness of $\text{PLWE}_{f,q,\chi}$ for parameters f, q, χ as described in Section 6.1, there exists a protocol that securely implements \mathcal{F} against (static) semi-honest adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}})$ -hybrid model.*

4.2 Security Under Semi-Honest Adversaries

Proof of Lemma 4.1: Correctness follows from correctness of key combination, correctness of homomorphic evaluation, and correctness of share decryption of \mathcal{E} .

We now turn to security. Let \mathcal{A} be a real-world semi-honest adversary corrupting $t < N$ clients, and possibly the server. Let $T \subsetneq [N]$ be the set of corrupted clients. We describe the simulator \mathcal{S}^{SH} in Figure 9.

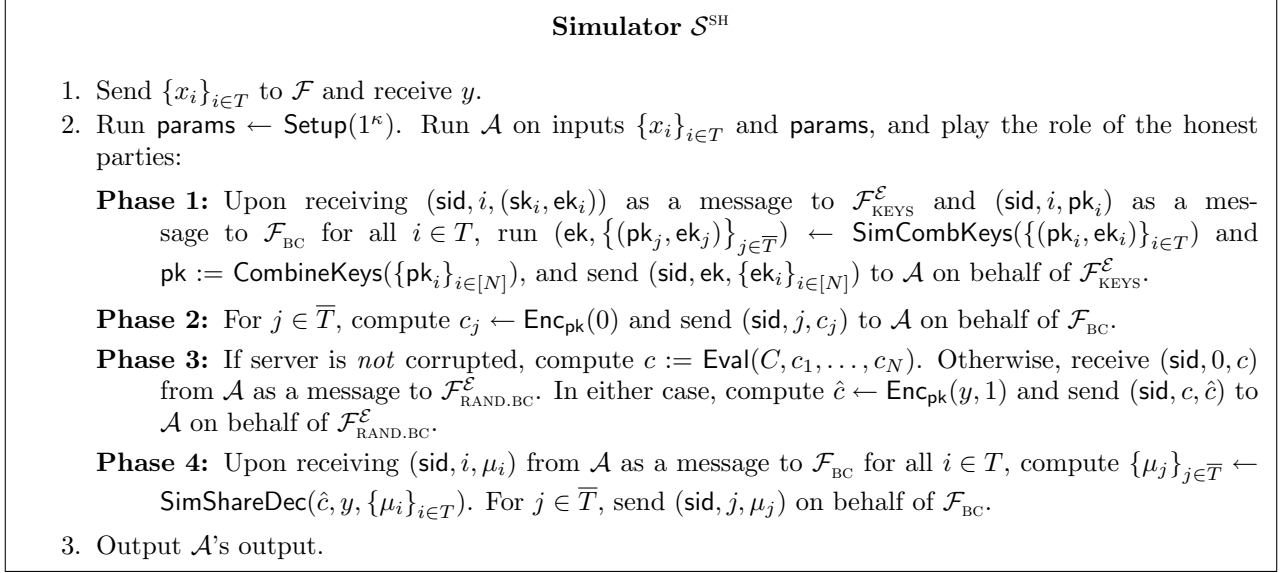


Figure 9: Simulator \mathcal{S}^{SH} for (static) semi-honest adversary \mathcal{A} corrupting clients $T \subsetneq [N]$ and possibly server S .

We prove that $\text{REAL}_{\Pi^{\text{SH}}, \mathcal{A}}(\vec{x}) \approx \text{IDEAL}_{\mathcal{F}, \mathcal{S}^{\text{SH}}}(\vec{x})$ via a series of hybrids. The view of the \mathcal{A} consists of

$$\text{View} = \left(\text{params}, \{x_i, (\text{pk}_i, \text{sk}_i, \text{ek}_i), c_i, \mu_i\}_{i \in T}, \text{pk}, \text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}, \{c_j\}_{j \in \bar{T}}, c, \hat{c}, \{\mu_j\}_{j \in \bar{T}} \right)$$

We simply prove that for all $\{x_i, (\text{pk}_i, \text{sk}_i, \text{ek}_i), c_i, \mu_i\}_{i \in T}$ output by \mathcal{A} , the joint distribution of these variables in the simulation is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol Π^{SH} . We let View_i be the joint distribution of these variables in hybrid i . In all hybrids, $\text{params} \leftarrow \text{Setup}(1^\kappa, N)$, $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ (this is true even when the server is corrupted since the adversary is semi-honest), and $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in N})$, but we gradually change how each of the other variables is computed.

Hybrid 0: This is the real-world execution of Π^{SH} . We have:

$$\begin{aligned} \{(\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}} \quad , \quad (\cdot, \text{ek}) &:= \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \\ \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{ReRand}(c) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \end{aligned}$$

Hybrid 1: We change how \hat{c} is created. Instead of using the ReRand algorithm, we simply sample a fresh encryption of $y = C(x_1, \dots, x_N)$ with $\text{mode} = 1$:

$$\begin{aligned} \{(\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}} \quad , \quad (\cdot, \text{ek}) &:= \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \\ \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \end{aligned}$$

We claim that $\text{View}_0 \stackrel{s}{\approx} \text{View}_1$ by *ciphertext re-randomization* of \mathcal{E} . The reduction gets $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$ from \mathcal{A} and computes $\{(\text{pk}_j, \text{sk}_j, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}}$ honestly. It computes (\cdot, ek) and pk using CombineKeys and CombinePK , respectively, and also computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}}$ and $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ if the server is honest (otherwise, it receives c from \mathcal{A}). Given \hat{c} which is either $\hat{c} = \text{ReRand}(c)$ or $\hat{c} = \text{Enc}_{\text{pk}}(y, 1)$, it computes $\{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}}$. Since \mathcal{A} is a semi-honest adversary, we know that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple for all $i \in T$. Likewise, we know that for all $i \in T$, there exist $r_i^{(x)}$ such that $c_i = \text{Enc}_{\text{pk}}(x_i, 0; r_i^{(x)})$. Therefore, $\text{View}_0 \stackrel{s}{\approx} \text{View}_1$ by *ciphertext re-randomization* of \mathcal{E} .

Hybrid 2: We change the way the decryption shares $\{\mu_j\}_{j \in \bar{T}}$ are generated. Instead of using the ShareDec algorithm with secret key sk_j , we now use the SimShareDec algorithm with $y = C(x_1, \dots, x_N)$ and $\{\mu_i\}_{i \in T}$:

$$\begin{aligned} & \{(\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}} \quad , \quad (\cdot, \text{ek}) := \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \\ & \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\mu_i\}_{i \in T}) \end{aligned}$$

We claim that $\text{View}_1 \stackrel{c}{\approx} \text{View}_2$ by *share-simulation indistinguishability* of \mathcal{E} . The reduction takes $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$ from \mathcal{A} and computes $\{(\text{pk}_j, \text{sk}_j, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}}$ honestly. It computes pk and ek using CombinePK and CombineKeys , respectively, and also computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}}$ and $\hat{c} = \text{Enc}_{\text{pk}}(y, 1)$. With $\{c_i\}_{i \in T}$ that \mathcal{A} sends, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$. Since \mathcal{A} is a semi-honest adversary, we know that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple for all $i \in T$. Likewise, we know that for all $i \in T$, there exists $r_i^{(dec)}$ such that $\mu_i := \text{ShareDec}(\hat{c}; r_i^{(dec)})$. Thus, $\text{View}_1 \stackrel{c}{\approx} \text{View}_2$ by *share-simulation indistinguishability* of \mathcal{E} .

Hybrid 3: This is actually a series of $N - t$ hybrids. We change how the ciphertexts $\{c_j\}_{j \in \bar{T}}$ are created, one at a time. Instead of having c_{j^*} be an encryption of input x_{j^*} , we have it be an encryption of 0.

$$\begin{aligned} & \{(\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j)\}_{j \in \bar{T}} \quad , \quad (\cdot, \text{ek}) := \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \\ & \{c_j = \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}, j \leq j^*} \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}, j > j^*} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1) \\ & \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\mu_i\}_{i \in T}) \end{aligned}$$

We claim two consecutive hybrids are computationally indistinguishable by *semantic security* of \mathcal{E} . The reduction forwards $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$ from \mathcal{A} to the challenger and receives ek , $\{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}$ from the challenger, which it forwards to \mathcal{A} . It uses the challenge ciphertext as c_{j^*} and computes the other ciphertexts c_j accordingly. With $\{c_i\}_{i \in T}$ that \mathcal{A} sends, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ (or receives it from \mathcal{A} if the server is corrupted). It also computes $\hat{c} = \text{Enc}_{\text{pk}}(y, 1)$ and $\{\mu_j \leftarrow \text{SimShareDec}(\hat{c}, y, \{\mu_i\}_{i \in T})\}_{j \in \bar{T}}$ from the values $\{\mu_i\}_{i \in T}$ sent by \mathcal{A} . We define View_3 to be the view having all c_j encrypting 0. Putting these $N - t$ hybrids together, we conclude that $\text{View}_2 \stackrel{c}{\approx} \text{View}_3$.

Hybrid 4: We change how the public/evaluation key pairs $\{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}$ and the combined evaluation key ek is generated. Instead of creating them using Keygen and CombineKeys , we use the SimCombKeys algorithm with $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$.

$$(\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}) \leftarrow \text{SimCombKeys}(\{(\text{pk}_i, \text{ek}_i)\}_{i \in T})$$

$$\{c_j = \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\mu_i\}_{i \in T})$$

We claim that $\text{View}_3 \stackrel{c}{\approx} \text{View}_4$ by *key-simulation indistinguishability* of \mathcal{E} . The reduction takes $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in T}$ from \mathcal{A} and computes $(\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}) \leftarrow \text{SimCombKeys}(\{(\text{pk}_i, \text{ek}_i)\}_{i \in T})$ and $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$. It also computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}}$ and $\hat{c} = \text{Enc}_{\text{pk}}(y, 1)$. With $\{c_i\}_{i \in T}$ that \mathcal{A} sends, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ (or receives it from \mathcal{A} if the server is corrupted). With $\{\mu_i\}_{i \in T}$ that \mathcal{A} sends, it computes $\{\mu_j \leftarrow \text{SimShareDec}(\hat{c}, y, \{\mu_i\}_{i \in T})\}_{j \in \bar{T}}$. Therefore, $\text{View}_2 \stackrel{c}{\approx} \text{View}_3$ by *key-simulation indistinguishability* of \mathcal{E} .

We have proved that $\text{View}_0 \stackrel{c}{\approx} \text{View}_4$. Since View_4 is the view that \mathcal{S}^{SH} produces, and View_0 is the view produced in an execution of Π^{SH} , we have that $\text{IDEAL}_{\mathcal{F}, \mathcal{S}^{\text{SH}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi^{\text{SH}}, \mathcal{A}}(\vec{x})$. \square

5 Tolerating Malicious Adversaries

The protocol described in Section 4 is not secure against malicious adversaries. We now describe the changes we need to make for our protocol to be able to tolerate this kind of attack. We assume a *rushing adversary*, that is, we assume that the adversary can choose his messages *adaptively*, depending on the messages from the honest players. In order to protect against malicious players, we need to guarantee that the public and evaluation keys chosen by the players are independent. To achieve this, we add a commitment key ck for a trapdoor commitment scheme (e.g. a semantically secure encryption scheme) to the crs . This guarantees that ck is generated honestly. We then have all parties commit, using ck , to the randomness $r_i^{(\text{gen})}$ that they will use in Keygen . The same is true for decryption shares; we must guarantee that they are independent, and therefore also have the parties commit to the the randomness $r_i^{(\text{dec})}$ that they will use in ShareDec . This is combined with having each player provide a zero-knowledge proof along with $(\text{pk}_i, \text{ek}_i)$ that proves that the keys were indeed created using the randomness $r_i^{(\text{gen})}$ he committed to previously. Similarly, we require each player to prove that he computed his decryption share μ_i with the randomness he committed to, and the secret-key that was output by Keygen on randomness $r_i^{(\text{gen})}$. Furthermore, we need to ensure that the ciphertexts c_i are independent, and so we have all players provide a proof of plaintext knowledge (POPK) for c_i .

There is a subtlety here. Because we have a rushing adversary, in the proof the simulator will have to provide simulated proofs for the honest parties and still require soundness from the proofs created by the adversary on behalf of corrupt players. We therefore need to use simulation-sound NIZKs (SS-NIZKs). For the same reason, we need simulation-extractability (SE) for the POPK of c_i : the simulator must still be able to extract from an adversarial proof, even if the adversary has seen many simulated proofs.

To protect against a malicious server, the parties must also verify the computation performed by the server. The ciphertext c is a deterministic function of the set of ciphertexts provided by the parties $\{c_i\}_{i \in N}$, so they could potentially verify the correctness of c by doing the homomorphic operations themselves. This is obviously unsatisfactory, as they would perform computation proportional to the size of the computation. This is precisely the problem we're trying to solve! Note, however, that precisely because the parties need to verify the computation, we cannot hope to have the communication complexity and computation time of the parties be independent of the complexity of f , as required by the definition of cloud-assisted MPC given in Section 3. We therefore relax the definition and require the communication complexity and computation time of parties P_1, \dots, P_N to be *polylogarithmic* in $|C|$ and the computation time of the server to be *polynomial* in $|C|$.

To prove that it did the evaluation correctly, we require the server to provide a short argument φ , which the parties can verify in little time. This can be done in one of four different ways, each with its own benefits and drawbacks. We comment that the circuit in the discussion below is the evaluation circuit, which can be viewed as the circuit resulting from replacing each gate in C by $\text{poly}(\kappa)$ many gates. Thus, both the size and depth of the evaluation circuit are (at most) a polynomial factor (in κ) larger than the size and depth of C .

1. Use the argument system of [Kil92, Kil95]. The benefit in doing this is that the communication complexity of the argument is polylogarithmic in the size of the computation, and the computation time of the verifier is linear in the input length (up to polylogarithmic factors). On the other hand, the system of [Kil92, Kil95] is an *interactive* system consisting of 3 rounds, and therefore this approach implies increasing the number of rounds in our protocol by 2 rounds.
2. Use computationally sound (CS) proofs [Mic94]. The benefit of using this approach is that the verifier runs in nearly linear time and the communication complexity is again polylogarithmic in the size of the computation. The drawback, however, is that CS proofs require the use of a random oracle, and thus, we could only hope to achieve security in the *random oracle model*.
3. Use the one-round argument system of [GKR08]. On the one hand, using this approach allows the server to provide an argument for the correct computation of c without PCPs, but on the other hand, both the communication complexity and the computation time of the verifier are polynomial in the *depth* of the circuit, thus restricting the type of functions we can compute to those in **NC**. The construction of [GKR08] also needs to assume the existence of a secure PIR scheme.
4. Use the one-round argument system of [BCCT11] or [GLR11]. With this approach, the communication complexity is again polylogarithmic in the size of the computation, the computation time of the verifier is nearly linear in the input size, and security holds in the standard model. However, the use of either of these argument systems requires us to make a non-falsifiable assumption [Nao03].

We give a full description of the modified protocol that is secure against malicious adversaries in Figure 10. We require the same properties from \mathcal{E} , namely, we assume \mathcal{E} is a key-homomorphic threshold FHE. We also need an argument system $\Phi = (\text{Setup}^\Phi, \text{Prove}, \text{Verify})$ for verifiable computation, and a trapdoor commitment scheme $(\text{Keygen}^{\text{Com}}, \text{Com})$. We describe the protocol in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{gen})}, \mathcal{F}_{\text{ZK.BC}}^{R(x)}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{dec})})$ -hybrid model, where $R^{(\text{gen})}$, $R^{(x)}$ and $R^{(\text{dec})}$ are NP relations:

$$\begin{aligned}
R^{(\text{gen})} &= \left\{ (\text{params}, \text{pk}_i, \text{ek}_i, \text{ck}, \rho_i^{(\text{gen})}, i), (r_i^{(\text{gen})}, s_i^{(\text{gen})}) \mid (\text{pk}_i, \cdot, \text{ek}_i) := \text{Keygen}(\text{params}, i; r_i^{(\text{gen})}) \wedge \right. \\
&\quad \left. \rho_i^{(\text{gen})} := \text{Com}_{\text{ck}}(r_i^{(\text{gen})}; s_i^{(\text{gen})}) \right\} \\
R^{(x)} &= \{(c, \text{pk}), (x, r) \mid c := \text{Enc}_{\text{pk}}(x, 0; r)\} \\
R^{(\text{dec})} &= \left\{ (\text{params}, c, \mu_i, \text{pk}_i, \rho_i^{(\text{dec})}, \rho_i^{(\text{gen})}, \text{ck}, i), (r_i^{(\text{dec})}, s_i^{(\text{dec})}, r_i^{(\text{gen})}, s_i^{(\text{gen})}) \mid \right. \\
&\quad (\text{pk}_i, \text{sk}_i, \cdot) \leftarrow \text{Keygen}(\text{params}, i; r_i^{(\text{gen})}) \wedge \rho_i^{(\text{gen})} := \text{Com}_{\text{ck}}(r_i^{(\text{gen})}; s_i^{(\text{gen})}) \wedge \\
&\quad \left. \mu_i := \text{ShareDec}_{\text{sk}_i}(c; r_i^{(\text{dec})}) \wedge \rho_i^{(\text{dec})} := \text{Com}_{\text{ck}}(r_i^{(\text{dec})}; s_i^{(\text{dec})}) \right\},
\end{aligned}$$

Lemma 5.1. *Let f be an N -input function, and let \mathcal{F} be the ideal functionality computing f . Let \mathcal{E} be a key-homomorphic threshold FHE scheme. Let Φ be an argument system for verifiable computation, and*

Protocol Π^{MAL}

Let C be a circuit computing f . Let $\mathcal{E} = (\text{Setup}, \text{Keygen}, \text{Enc}, \text{Eval}, \text{CombinePK}, \text{CombineKeys}, \text{ShareDec}, \text{ShareCombine}, \text{ReRand})$ be a key-homomorphic threshold FHE scheme. Let $\Phi = (\text{Setup}^\Phi, \text{Prove}, \text{Verify})$ be an argument system for verifiable computation, and let $(\text{Keygen}^{\text{Com}}, \text{Com})$ be a trapdoor commitment scheme.

For session ID sid :

Phase 1: Each P_i calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with (sid, i) and receives $(\text{params}, \text{ck})$. Samples randomness $r_i^{(\text{gen})}$ and $r_i^{(\text{dec})}$ and computes $\rho_i^{(\text{gen})} := \text{Com}_{\text{ck}}(r_i^{(\text{gen})}; s_i^{(\text{gen})})$ and $\rho_i^{(\text{dec})} \leftarrow \text{Com}_{\text{ck}}(r_i^{(\text{dec})}; s_i^{(\text{dec})})$. Sends $(\text{sid}, i, (\rho_i^{(\text{gen})}, \rho_i^{(\text{dec})}))$ to \mathcal{F}_{BC} .

Phase 2: Each P_i chooses an encryption key tuple: $(\text{pk}_i, \text{sk}_i, \text{ek}_i) := \text{Keygen}(1^\kappa, i; r_i^{(\text{gen})})$, and sends $(\text{sid}, i, (\text{sk}_i, \text{ek}_i))$ to $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, $(\text{sid}, i, \text{pk}_i)$ to \mathcal{F}_{BC} , and $(\text{sid}, i, (\text{pk}_i, \text{ek}_i, \text{ck}, \rho_i^{(\text{gen})}, i), (r_i^{(\text{gen})}, s_i^{(\text{gen})}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{gen})}}$.

Phase 3: Upon receiving $(\text{sid}, \text{ek}, \{(\text{ek}_i)\}_{i \in [N]})$ from $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, $(\text{sid}, i, \text{pk}_i)$ from \mathcal{F}_{BC} and $(\text{sid}, i, (\text{pk}_i, \text{ek}_i, \text{ck}, \rho_i^{(\text{gen})}, i))$ from $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{gen})}}$ for all $i \in [N]$, each P_i computes $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$ and $c_i := \text{Enc}_{\text{pk}}(x_i; r_i^{(x)})$ and sends $(\text{sid}, i, (c_i, \text{pk}), (x_i, r_i^{(x)}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R^{(x)}}$.

Phase 4: Upon receiving $(\text{sid}, i, (c_i, \text{pk}))$ from $\mathcal{F}_{\text{ZK.BC}}^{R^{(x)}}$ for all $i \in [N]$ and $(\text{sid}, \text{ek}, \{\text{ek}_i\}_{i \in [N]})$ from $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, server S computes $c := \text{Eval}(c_1, \dots, c_N)$, and sends $(\text{sid}, 0, c)$ to $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$. Also computes $\varphi \leftarrow \text{Prove}(c, c_1, \dots, c_N)$ and sends $(\text{sid}, 0, \varphi)$ to \mathcal{F}_{BC} .

Phase 5: Upon receiving (sid, c, \hat{c}) from $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$, and φ from \mathcal{F}_{BC} , each client P_i verifies φ , computes $\mu_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(\text{dec})})$ and sends $(\text{sid}, i, (\hat{c}, \mu_i, \text{pk}_i, \rho_i^{(\text{dec})}, \text{ck}, i), (\text{sk}_i, r_i^{(\text{dec})}, s_i^{(\text{dec})}, r_i^{(\text{gen})}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{dec})}}$.

Local Computation: Upon receiving $(\text{sid}, i, (\hat{c}, \mu_i, \text{pk}_i, \rho_i^{(\text{dec})}, \text{ck}, i))$ from $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{dec})}}$ for all $i \in [N]$, each client P_i computes and outputs $y := \text{ShareCombine}(c, \mu_1, \dots, \mu_N)$.

Figure 10: Protocol Π^{MAL} . Implements \mathcal{F} in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{gen})}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(x)}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{dec})}})$ -hybrid model, where \mathcal{D} is the distribution of $(\text{params}, \text{ck})$, where $\text{params} \leftarrow \text{Setup}(1^\kappa, N)$ and $(\text{ck}, \cdot) \leftarrow \text{Keygen}^{\text{Com}}(1^\kappa)$. Secure against malicious adversaries.

let $R^{(x)}, R^{(\text{dec})}$ be NP-relations as described above. Then the protocol Π^{MAL} described in Figure 10 securely implements \mathcal{F} against (static) malicious adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(x)}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{dec})}})$ -hybrid model.

The proof of Lemma 5.1 is given in Appendix A.1. Instantiating key-homomorphic threshold FHE based on the PLWE assumption, a special case of the ‘‘Ring LWE’’ assumption of [LPR10], as in Section 6, and combining Lemma 5.1 with protocols implementing $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ and $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ for this instantiation (shown in Section 7), SS-NIZKs for $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{gen})}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{dec})}}$, and SE-NIZKs for $\mathcal{F}_{\text{ZK.BC}}^{R^{(x)}}$, we arrive at the following theorem:

Theorem 5.2. *Assuming the hardness of $\text{PLWE}_{f,q,\chi}$ for parameters f, q, χ as described in Section 6.1, there exists a protocol that securely implements \mathcal{F} against (static) malicious adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}})$ -hybrid model.*

6 Instantiating Key-Homomorphic Threshold FHE

We now show how to instantiate public-key key-homomorphic threshold FHE (KHT-FHE, Definition 4.1). Our construction will be based on PLWE assumption [BV11b], which is a special case of the Ring LWE assumption of [LPR10]. We therefore first review the PLWE assumption (Section 6.1) and then present the scheme. To do the latter, we first review the somewhat homomorphic scheme of [BV11b] (Section 6.2) and show how to use the re-linearization and modulus reduction techniques of [BV11a, BGV11] to convert it into a (leveled) fully homomorphic scheme (Section 6.3). We then show that the scheme satisfies key homomorphism and also show how to instantiate the algorithms CombinePK, CombineKeys, ShareDec, ShareCombine, ReRand and simulators SimCombKeys, SimShareDec (Section 6.4).

6.1 The Polynomial-LWE Assumption

In this section, we describe the *polynomial LWE* (or, PLWE) assumption from [BV11b], which is a special case of the “ring learning with errors” (RLWE) assumption of Lyubashevsky, Peikert and Regev [LPR10]. The PLWE assumption is analogous to the by now standard “learning with errors” (LWE) assumption, first defined by Regev [Reg05, Reg09] (generalizing the learning parity with noise assumption of Blum et al. [BFKL93]). In the PLWE assumption, we consider rings $R \doteq \mathbb{Z}[x]/\langle f(x) \rangle$ and $R_q \doteq R/qR$ for some degree n integer polynomial $f(x) \in \mathbb{Z}[x]$ and a prime integer $q \in \mathbb{Z}$. Note that $R_q \equiv \mathbb{Z}_q[x]/\langle f(x) \rangle$, i.e. the ring of degree n polynomials modulo $f(x)$ with coefficients in \mathbb{Z}_q . Addition in these rings is done component-wise in their coefficients (thus, their additive group is isomorphic to \mathbb{Z}^n and \mathbb{Z}_q^n respectively). Multiplication is simply polynomial multiplication modulo $f(x)$ (and also q , in the case of the ring R_q). Thus an element in R (or R_q) can be viewed as a degree n polynomial over \mathbb{Z} (or \mathbb{Z}_q). We represent such an element using the vector of its coefficients. For an element $a(x) = a_0 + a_1x + \dots + a_{n-1}x^{n-1} \in R$, we let $\|a\| = \max |a_i|$ denote its ℓ_∞ norm.

The $\text{PLWE}_{f,q,\chi}$ assumption is that for a random ring element $s \leftarrow R_q$, given any polynomial number of samples of the form $(a_i, b_i = a_i \cdot s + e_i) \in (R_q)^2$, where a_i is uniformly random in R_q and e_i is drawn from the error distribution χ , the b_i 's are computationally indistinguishable from uniform in R_q . We use the *hermite normal form* of the assumption, as in [BV11b], where the secret s is sampled from the noise distribution χ rather than being uniform in R_q . This presentation is more useful for the purposes of this paper and it turns out that to be equivalent to the original one up to obtaining one additional sample [AiCPS09, LPR10].

Definition 6.1 (The PLWE Assumption - Hermite Normal Form [BV11b, LPR10]). *For all $\kappa \in \mathbb{N}$, let $f(x) = f_\kappa(x) \in \mathbb{Z}[x]$ be a polynomial of degree $n = n(\kappa)$, let $q = q(\kappa) \in \mathbb{Z}$ be a prime integer, let the ring $R \doteq \mathbb{Z}[x]/\langle f(x) \rangle$ and $R_q \doteq R/qR$, and let χ denote a distribution over the ring R .*

The polynomial LWE assumption $\text{PLWE}_{f,q,\chi}$ states that for any $\ell = \text{poly}(\kappa)$ it holds that

$$\{(a_i, a_i \cdot s + e_i)\}_{i \in [\ell]} \stackrel{c}{\approx} \{(a_i, u_i)\}_{i \in [\ell]},$$

where s is sampled from the noise distribution χ , a_i are uniform in R_q , the “error polynomials” e_i are sampled from the error distribution χ , and finally, the ring elements u_i are uniformly random over R_q .

Fact 6.2 ([BV11b]). *The $\text{PLWE}_{f,q,\chi}^\ell$ assumption implies that,*

$$\{(a_i, a_i \cdot s + 2 \cdot e_i)\}_{i \in [\ell]} \stackrel{c}{\approx} \{(a_i, u_i)\}_{i \in [\ell]}.$$

where a_i, s, e_i and u_i are as in Definition 6.1.

Choice of Parameters The correctness, security, and homomorphic properties of the scheme of [BV11b] rely on specific choices of the polynomial $f(x)$, the modulus q , and the error distribution χ . For security parameter κ :

- We set $f(x)$ to be the $2^{\lceil \log \kappa \rceil}$ th cyclotomic polynomial. This implies that $f(x) = \Phi_{2^{\lceil \log \kappa \rceil}}(x) = x^n + 1$ for $n = 2^{\lceil \log \kappa \rceil - 1}$. Since $n \in (\kappa/4, \kappa]$, all asymptotics can be stated in terms of n .
- The error distribution χ is the (truncated) discrete Gaussian distribution $D_{\mathbb{Z}^n, r, B(r)}$ for standard deviation $r > 0$ and truncation bound $B(r)$. A sample from this distribution defines a polynomial $e(x) \in R$ with each coefficient bounded by $B(r)$. If $B(r) = \omega(\log \kappa \cdot r)$, then the truncated distribution $D_{\mathbb{Z}^n, r, B(r)}$ is statistically close from the (non-truncated) discrete Gaussian $D_{\mathbb{Z}^n, r}$.

The Worst-case to Average-case Connection. We state a worst-case to average-case reduction from the shortest vector problem on ideal lattices to the PLWE problem for our setting of parameters. The reduction stated below is a special case the results of [LPR10].

Theorem 6.1 (A special case of [LPR10]). *Let κ be the security parameter. Let $k \in \mathbb{N}$ and let $m = 2^{\lceil \log \kappa \rceil}$ be a power of two. Let $\Phi_m(x) = x^n + 1$ be the m^{th} cyclotomic polynomial of degree $n = \varphi(m) = m/2$. Let $r \geq \omega(\sqrt{\log n})$ be a real number, and let $q \equiv 1 \pmod{m}$ be a prime integer. Let $R = \mathbb{Z}[x]/\langle \Phi_{2^{\lceil \log \kappa \rceil}}(x) \rangle$. Then:*

- *There is a randomized reduction from $2^{\omega(\log n)} \cdot (q/r)$ -approximate R-SVP to $\text{PLWE}_{\Phi_{2^{\lceil \log \kappa \rceil}}, q, \chi}$ where $\chi = D_{\mathbb{Z}^n, r}$ is the discrete Gaussian distribution. The reduction runs in time $\text{poly}(n, q)$.*
- *There is a randomized reduction from $(n^2 q/r) \cdot (n(\ell + 1)/\log(n(\ell + 1)))^{1/4}$ -approximate R-SVP to $\text{PLWE}_{\Phi_{2^{\lceil \log \kappa \rceil}}, q, \chi}^{(\ell)}$ where $\chi = D_{\mathbb{Z}^n, r}$ is the discrete Gaussian distribution. The reduction runs in time $\text{poly}(n, q, \ell)$.*

6.2 A Somewhat-Homomorphic Scheme

We describe the scheme of Brakerski and Vaikuntanathan [BV11b] based on PLWE.³ For security parameter κ , the scheme is parameterized by a prime number q , a degree n polynomial $f(x) \in \mathbb{Z}[x]$, and an error distribution χ over the ring $R_q \doteq \mathbb{Z}_q[x]/\langle f(x) \rangle$. The parameters n, f, q and χ are public and we assume that given κ , there are polynomial-time algorithms that output f and q , and sample from the error distribution χ .

- $\text{Setup}(1^\kappa)$: Sample a ring element vector $\mathbf{a} \leftarrow R_q^n$, and set $\text{params} := \mathbf{a}$.
- $\text{Keygen}(\text{params})$: Sample a ring element $s \leftarrow \chi$ and a ring element vector $\mathbf{x} \leftarrow \chi^n$, and set

$$\text{sk} := s \quad \text{and} \quad \text{pk} := \mathbf{p} = \mathbf{a}s + 2\mathbf{x} \in R_q^n$$

- $\text{Enc}(\text{sk}, m)$: To encrypt a bit $m \in \{0, 1\}$, sample $\mathbf{e} \leftarrow \chi^n$, and compute

$$c_0 := \langle \mathbf{p}, \mathbf{e} \rangle + m \in R_q \quad \text{and} \quad c_1 := \langle \mathbf{a}, \mathbf{e} \rangle \in R_q$$

and output the ciphertext $\mathbf{c} := (c_0, c_1) \in R_q^2$.

³The scheme presented in [BV11b] is a private-key FHE. We make the natural modifications and present it as a public-key scheme.

- $\text{Dec}(\text{sk}, \mathbf{c})$: To decrypt, we first compute $c_0 - c_1 s \in R_q$, and then output $m = c_0 - c_1 s \pmod{2}$ as the message. Note that the condition for correct decryption is that the ℓ_∞ norm of the polynomial $c_0 - c_1 s$ is smaller than $q/2$.

Theorem 6.2 ([BV11b]). *Let $n, q, f(x)$ be as in the scheme, let $r = \text{poly}(n)$ and $q = 2^{n^\epsilon}$ for some $0 < \epsilon < 1$. Then, the scheme allows evaluation of degree- $O(n^\epsilon/\log n)$ polynomials with at most $2^{O(n^\epsilon/\log n)}$ terms, and is secure under the worst-case hardness of approximating shortest vectors on ideal lattices to within a factor of $O(2^{n^\epsilon})$.*

We will use the fact that ciphertexts of this scheme are *pseudorandom*. This follows from the fact that \mathbf{p} is a PLWE sample and therefore pseudorandom, and the fact that for \mathbf{e} chosen according to the discrete gaussian, $\langle \mathbf{a}, \mathbf{e} \rangle$ is statistically close to uniform.

6.3 Getting a (Leveled) Fully-Homomorphic Scheme via Re-linearization and Modulus Reduction

We now describe how to transform the somewhat homomorphic scheme from Section 6.2 into a (leveled) fully homomorphic scheme. We use the re-linearization and modulus reduction techniques of [BV11a, BGV11], as opposed to the “squashing” and “bootstrapping” blueprint of [Gen09b] used in [BV11b]. We do this in order to keep the `CombineKeys` algorithm and thus the N -party protocol implementing the “combine keys” functionality $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$, simple.

6.3.1 Re-linearization

Let us first look at the symbolic linear function $f_{\mathbf{c}}(x) : R_q \rightarrow \{0, 1\}$ defined as: $f_{\mathbf{c}}(x) = c_1 - c_0 x \pmod{2}$. We can describe homomorphic addition and multiplication in terms of f . First note that decrypting a ciphertext \mathbf{c} corresponds to simply evaluating $f_{\mathbf{c}}(s)$. For ciphertexts \mathbf{c} and \mathbf{c}' encrypting m and m' respectively, we have:

$$m + m' = f_{\mathbf{c}}(s) + f_{\mathbf{c}'}(s) = (c_0 + c'_0) - (c_1 + c'_1)s$$

Therefore, it is natural to define $\mathbf{c}_{\text{add}} = \mathbf{c} + \mathbf{c}' = (c_0 + c'_0, c_1 + c'_1)$. Multiplication is a bit more tricky. We have:

$$m \cdot m' = f_{\mathbf{c}}(s) \cdot f_{\mathbf{c}'}(s) = (c_0 - c_1 s)(c'_0 - c'_1 s) = c_0 c'_0 - (c_0 c'_1 + c_1 c'_0)s + c_1 c'_1 s^2 = h_0 + h_1 s + h_2 s^2$$

In [BV11b], the authors let $\mathbf{c}_{\text{mult}} = (h_0, h_1, h_2)$, so that the size of the ciphertext grows linearly with the number of multiplications performed. This requires the secret key to be $\mathbf{s} = (1, s, s^2, \dots, s^D)$, where D is the maximal degree of homomorphism tolerated by the scheme. In our setting of multiparty computation, this would require a very complex protocol Π_{KEYS} to compute `CombineKeys`, and therefore we opt for another strategy. Instead of letting the ciphertext grow with each multiplication, we use the re-linearization technique of Brakerski and Vaikuntanathan [BV11a] to reduce the size of the resulting ciphertext c_{mult} after each multiplication.

To be able to homomorphically evaluate a polynomial of degree D while keeping the ciphertext size constant, we have `Keygen` sample D different polynomials $s_0, \dots, s_D \leftarrow \chi$, one for each level, and set $\text{sk} = (s_0, \dots, s_D)$, where s_0 is used to create the public key $\text{pk} := \mathbf{p} = \mathbf{a}s_0 + \mathbf{x}$ and s_D is used for decryption.⁴ The evaluation key is computed as follows. For all $d \in [D], \tau \in \{0, \dots, \lfloor \log q \rfloor\}$, sample

⁴The secret key sk could potentially be composed of only s_D . However, we let it contain all s_i 's because we will need them when we implement the `CombineKeys` algorithm.

$(a_{d,\tau}, b_{d,\tau} = a_{d,\tau}s_d + 2e_{d,\tau}) \in R_q^2$ and $(a'_{d,\tau}, b'_{d,\tau} = a'_{d,\tau}s_d + 2e'_{d,\tau}) \in R_q^2$, where $a_{d,\tau}, a'_{d,\tau} \leftarrow R_q$ and $e_{d,\tau}, e'_{d,\tau} \leftarrow \chi$. Compute

$$\begin{aligned}\tilde{c}_{0,d,\tau} &:= b_{d,\tau} + 2^\tau s_{d-1} \in R_q & \text{and} & & \tilde{c}_{1,d,\tau} &:= a_{d,\tau} \\ \tilde{z}_{0,d,\tau} &:= b'_{d,\tau} + 2^\tau s_{d-1}^2 \in R_q & \text{and} & & \tilde{z}_{1,d,\tau} &:= a'_{d,\tau} \text{ }^5\end{aligned}$$

We let

$$\begin{aligned}\tilde{\mathbf{c}}_{0,d} &= (\tilde{c}_{0,d,0}, \dots, \tilde{c}_{0,d, \lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{c}}_1 &= (\tilde{c}_{1,d,0}, \dots, \tilde{c}_{1,d, \lfloor \log q \rfloor}) \\ \tilde{\mathbf{z}}_{0,d} &= (\tilde{z}_{0,d,0}, \dots, \tilde{z}_{0,d, \lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{z}}_1 &= (\tilde{z}_{1,d,0}, \dots, \tilde{z}_{1,d, \lfloor \log q \rfloor})\end{aligned}$$

and set the evaluation key

$$\text{ek} := \{\tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d}\}_{d \in [D]}$$

SH.Eval $(p, (\mathbf{c}_1, \dots, \mathbf{c}_N))$: We show how to evaluate an N -variate polynomial $p : \{0, 1\}^N \rightarrow \{0, 1\}$ of degree D . To this end, we show how to homomorphically add and multiply two elements in $\{0, 1\}$.

- Given two ciphertexts $\mathbf{c} = (c_0, c_1)$ and $\mathbf{c}' = (c'_0, c'_1)$ under the same secret key s_{d-1} , output the ciphertext $\mathbf{c}_{\text{add}} = \mathbf{c} + \mathbf{c}' = (c_0 + c'_0, c_1 + c'_1) \in R_q^2$, as an encryption of the *sum* of the underlying messages.
- Given two ciphertexts $\mathbf{c} = (c_0, c_1)$ and $\mathbf{c}' = (c'_0, c'_1)$ under the same secret key s_{d-1} , an encryption of their *product* is computed as follows.

Define $h_0, h_1, h_2 \in R_q$ as before: $h_0 + h_1x + h_2x^2 = (c_0 - c_1x)(c'_0 - c'_1x)$. Then in particular, $h_0 + h_1s_{d-1} + h_2s_{d-1}^2 \pmod{2} = m \cdot m'$. Let $\mathbf{u} = (u_0, \dots, u_{\lfloor \log q \rfloor}) \in R_2^{\lfloor \log q \rfloor}$ for $u_i \in R_2$ be the binary representation of h_1 , and let $\mathbf{v} = (v_0, \dots, v_{\lfloor \log q \rfloor}) \in R_2^{\lfloor \log q \rfloor}$ for $v_i \in R_2$ be the binary representation of h_2 , so that:

$$h_1 = \sum_{\tau=0}^{\lfloor \log q \rfloor} 2^\tau u_\tau \quad \text{and} \quad h_2 = \sum_{\tau=0}^{\lfloor \log q \rfloor} 2^\tau v_\tau$$

Output

$$\mathbf{c}_{\text{mult}} = (h_0 + \langle \mathbf{u}, \tilde{\mathbf{c}}_{0,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{0,d} \rangle, \langle \mathbf{u}, \tilde{\mathbf{c}}_{1,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{1,d} \rangle),$$

which is a valid encryption of $m \cdot m'$ under secret key s_d . To see why this is the case, simply notice that:

$$\langle \mathbf{u}, \tilde{\mathbf{c}}_{0,d} \rangle = \sum_{\tau=0}^{\lfloor \log q \rfloor} u_\tau c_{0,d,\tau} = \sum_{\tau=0}^{\lfloor \log q \rfloor} u_\tau a_{d,\tau} s_d + \sum_{\tau=0}^{\lfloor \log q \rfloor} u_\tau 2^\tau s_{d-1} + 2e = \langle \mathbf{u}, \tilde{\mathbf{c}}_{1,d} \rangle s_d + h_1 + s_{d-1} + 2e$$

and

$$\langle \mathbf{v}, \tilde{\mathbf{z}}_{0,d} \rangle = \sum_{\tau=0}^{\lfloor \log q \rfloor} v_\tau z_{0,d,\tau} = \sum_{\tau=0}^{\lfloor \log q \rfloor} v_\tau a'_{d,\tau} s_d + \sum_{\tau=0}^{\lfloor \log q \rfloor} v_\tau 2^\tau s_{d-1}^2 + 2e' = \langle \mathbf{v}, \tilde{\mathbf{z}}_{1,d} \rangle s_d + h_2 + s_{d-1}^2 + 2e',$$

where e and e' are “small” error terms. Thus,

$$h_0 + \langle \mathbf{u}, \tilde{\mathbf{c}}_{0,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{0,d} \rangle - (\langle \mathbf{u}, \tilde{\mathbf{c}}_{1,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{1,d} \rangle) s_d \pmod{2} = h_0 + h_1 s_{d-1} + h_2 s_{d-1}^2 \pmod{2} = m \cdot m'$$

⁵The reader may notice that $(\tilde{c}_{0,d,\tau}, \tilde{c}_{1,d,\tau})$ and $(\tilde{z}_{0,d,\tau}, \tilde{z}_{1,d,\tau})$ are private-key *pseudo-encryptions* of $2^\tau s_{d-1}$ and $2^\tau s_{d-1}^2$, respectively, under secret key s_d . We say they are pseudo-encryptions because even though they were created as normal ciphertexts, they cannot be decrypted in the usual way since $2^\tau s_{d-1}$ and $2^\tau s_{d-1}^2$ are not in the message space $\{0, 1\}$.

6.3.2 Modulus Reduction

We have shown how to keep the ciphertext size constant when performing evaluation. However, the magnitude of the noise in the ciphertext still grows with each operation. In order to handle this, we use the modulus reduction technique of [BV11a, BGV11] as a noise management technique. We use the notation of [BGV11]. For a secret key s , we let $\mathbf{s} = (1, s)$ and re-write the decryption function as: $[\langle \mathbf{c}, \mathbf{s} \rangle]_q \pmod{2}$, where $[\cdot]_q$ denotes reducing modulo q .

Modulus reduction allows us to transform a ciphertext $\mathbf{c} \in R_q^2$ into a different ciphertext $\mathbf{c}' \in R_p^2$ while preserving correctness:

$$[\langle \mathbf{c}', \mathbf{s} \rangle]_p = [\langle \mathbf{c}, \mathbf{s} \rangle]_q \pmod{2}$$

The transformation from \mathbf{c} to \mathbf{c}' involves simply scaling by (p/q) and rounding appropriately. Moreover, if s is short and p is sufficiently smaller than q , the “noise” in the ciphertext actually decreases.

Lemma 6.3 ([BGV11]). *Let p and q be two odd moduli, and let \mathbf{c} be an integer vector. Define \mathbf{c}' to be the integer vector closest to $(p/q) \cdot \mathbf{c}$ such that $\mathbf{c}' = \mathbf{c} \pmod{2}$. Then, for any \mathbf{s} with $\|[\langle \mathbf{c}, \mathbf{s} \rangle]_q\| < q/2 - (q/p) \cdot \|\mathbf{s}\|_1$, we have*

$$[\langle \mathbf{c}', \mathbf{s} \rangle]_p = [\langle \mathbf{c}, \mathbf{s} \rangle]_q \pmod{2} \quad \text{and} \quad \|[\langle \mathbf{c}', \mathbf{s} \rangle]_p\| < (p/q) \cdot \|[\langle \mathbf{c}, \mathbf{s} \rangle]_q\| + \|\mathbf{s}\|_1$$

where $\|\mathbf{s}\|$ and $\|\mathbf{s}\|_1$ are the ℓ_∞ and ℓ_1 norms of \mathbf{s} , respectively.

To see why the lemma statement is true, consider a coordinate $(\langle \mathbf{c}, \mathbf{s} \rangle)_i$ of $\langle \mathbf{c}, \mathbf{s} \rangle$. We know that there exists $k \in \mathbb{Z}$ such that:

$$[(\langle \mathbf{c}, \mathbf{s} \rangle)_i]_q = (\langle \mathbf{c}, \mathbf{s} \rangle)_i - kq \in \left[-\frac{q}{2} + \frac{q}{p} \|\mathbf{s}\|_1, \frac{q}{2} - \frac{q}{p} \|\mathbf{s}\|_1 \right],$$

so that

$$((p/q) \cdot \mathbf{c}, \mathbf{s})_i - kp \in \left[-\frac{p}{2} + \|\mathbf{s}\|_1, \frac{p}{2} - \|\mathbf{s}\|_1 \right]$$

Substituting $(p/q) \cdot \mathbf{c}$ by \mathbf{c}' incurs an error of at most $\|\mathbf{s}\|_1$. Therefore,

$$(\langle \mathbf{c}', \mathbf{s} \rangle)_i - kp \in \left[-\frac{p}{2}, \frac{p}{2} \right] \quad \text{and} \quad [(\langle \mathbf{c}', \mathbf{s} \rangle)_i]_p = (\langle \mathbf{c}', \mathbf{s} \rangle)_i - kp$$

This proves the second part of the lemma. To prove the first part, notice that since p and q are both odd, we know $kp \equiv kq \pmod{2}$. Moreover, we chose \mathbf{c}' such that $\mathbf{c}' \equiv \mathbf{c} \pmod{2}$. We thus have

$$\begin{aligned} (\langle \mathbf{c}', \mathbf{s} \rangle)_i - kp &\equiv (\langle \mathbf{c}, \mathbf{s} \rangle)_i - kq \pmod{2} \\ [(\langle \mathbf{c}, \mathbf{s} \rangle)_i]_p &\equiv [(\langle \mathbf{c}, \mathbf{s} \rangle)_i]_q \pmod{2} \end{aligned}$$

We now show the changes we need to make to the evaluation algorithm. Let D be the maximum level of homomorphism we wish to achieve, and let q_0, \dots, q_D be a ladder of decreasing moduli. As before, **Keygen** samples D different polynomials $s_0, \dots, s_D \leftarrow \chi$, one for each level, and sets $\mathbf{sk} = (s_0, \dots, s_D)$, where $\mathbf{pk} := \mathbf{p} = \mathbf{a}\mathbf{s}_0 + \mathbf{x}$ and s_D is used for decryption. The evaluation key is computed as before, except that $\tilde{c}_{0,d,\tau}, \tilde{c}_{1,d,\tau}, \tilde{z}_{0,d,\tau}, \tilde{z}_{1,d,\tau}$ are reduced modulo q_{d-1} . In more detail: for all $d \in [D], \tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$, sample $(a_{d,\tau}, b_{d,\tau} = a_{d,\tau}s_d + 2e_{d,\tau}) \in R_{q_{d-1}}^2$ and $(a'_{d,\tau}, b'_{d,\tau} = a'_{d,\tau}s_d + 2e'_{d,\tau}) \in R_{q_{d-1}}^2$, where $a_{d,\tau}, a'_{d,\tau} \leftarrow R_{q_{d-1}}$ and $e_{d,\tau}, e'_{d,\tau} \leftarrow \chi$. Compute

$$\begin{aligned} \tilde{c}_{0,d,\tau} &:= b_{d,\tau} + 2^\tau s_{d-1} \in R_{q_{d-1}} & \text{and} & & \tilde{c}_{1,d,\tau} &:= a_{d,\tau} \\ \tilde{z}_{0,d,\tau} &:= b'_{d,\tau} + 2^\tau s_{d-1}^2 \in R_{q_{d-1}} & \text{and} & & \tilde{z}_{1,d,\tau} &:= a'_{d,\tau} \end{aligned}$$

Let

$$\begin{aligned}\tilde{\mathbf{c}}_{0,d} &= (\tilde{c}_{0,d,0}, \dots, \tilde{c}_{0,d, \lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{c}}_1 &= (\tilde{c}_{1,d,0}, \dots, \tilde{c}_{1,d, \lfloor \log q \rfloor}) \\ \tilde{\mathbf{z}}_{0,d} &= (\tilde{z}_{0,d,0}, \dots, \tilde{z}_{0,d, \lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{z}}_1 &= (\tilde{z}_{1,d,0}, \dots, \tilde{z}_{1,d, \lfloor \log q \rfloor})\end{aligned}$$

and set the evaluation key

$$\text{ek} := \{\tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d}\}_{d \in [D]}$$

$\text{SH.Eval}(p, (\mathbf{c}_1, \dots, \mathbf{c}_N))$: We show how to evaluate an N -variate polynomial $p : \{0, 1\}^N \rightarrow \{0, 1\}$ of degree D . To this end, we show how to homomorphically add and multiply two elements in $\{0, 1\}$.

- Given two ciphertexts $\mathbf{c} = (c_0, c_1)$ and $\mathbf{c}' = (c'_0, c'_1)$ under the same secret key s_{d-1} and modulus q_{d-1} , output the ciphertext $\mathbf{c}_{\text{add}} = \mathbf{c} + \mathbf{c}' = (c_0 + c'_0, c_1 + c'_1) \in R_{q_{d-1}}^2$, as an encryption of the sum of the underlying messages.
- Given two ciphertexts $\mathbf{c} = (c_0, c_1)$ and $\mathbf{c}' = (c'_0, c'_1)$ under the same secret key s_{d-1} and modulus q_{d-1} , an encryption of their product is computed as follows. First, re-linearize: define $h_0, h_1, h_2 \in R_{q_{d-1}}$ as before, $h_0 + h_1x + h_2x^2 = (c_0 - c_1x)(c'_0 - c'_1x)$. Let $\mathbf{u} = (u_0, \dots, u_{\lfloor \log q \rfloor}) \in R_2^{\lfloor \log q \rfloor}$ for $u_i \in R_2$ be the binary representation of h_1 , and let $\mathbf{v} = (v_0, \dots, v_{\lfloor \log q \rfloor}) \in R_2^{\lfloor \log q \rfloor}$ for $v_i \in R_2$ be the binary representation of h_2 . Let

$$\mathbf{c}_{\text{mult}} = (h_0 + \langle \mathbf{u}, \tilde{\mathbf{c}}_{0,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{0,d} \rangle, \langle \mathbf{u}, \tilde{\mathbf{c}}_{1,d} \rangle + \langle \mathbf{v}, \tilde{\mathbf{z}}_{1,d} \rangle)$$

Finally, reduce the modulus: let $\mathbf{c}'_{\text{mult}}$ be the integer vector closest to $(q_d/q_{d-1}) \cdot \mathbf{c}_{\text{mult}}$ such that $\mathbf{c}'_{\text{mult}} = \mathbf{c}_{\text{mult}} \pmod{2}$. Output $\mathbf{c}'_{\text{mult}}$.

6.4 Key-Homomorphic Threshold FHE from PLWE

We describe an instantiation of key-homomorphic threshold FHE based on the FHE described in Section 6.2 and Section 6.3. Our scheme is parametrized by two noise distributions, χ_0 and χ_1 , with $\chi_b = D_{\mathbb{Z}^n, r_b, B(r_b)}$, as described in Section 6.1.

Algorithms. We instantiate algorithms Dec and Eval as described in Section 6.2 and Section 6.3. Below we show how to instantiate the algorithms Keygen, Enc, CombinePK, CombineKeys, ShareDec, ShareCombine, ReRand and simulators SimCombKeys, SimShareDec. For threshold decryption we use a scheme analogous to the one of [BD10] for standard LWE. We note, however, that threshold decryption in [BD10] is only secure (ie. shares do not reveal information about the secret key) if the ciphertext that is being decrypted has a random (or statistically close to random) c_1 . This is the reason why in our cloud-assisted MPC protocol, we need the parties to jointly *re-randomize* the ciphertext $c = \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ before they can threshold decrypt it.

- $\text{Setup}(1^\kappa, N)$: Let $\mathbf{a} \leftarrow R_q^n$. For $i \in [N], d \in \{0, \dots, D\}, \tau \in \{0, \dots, \lfloor \log q_d \rfloor\}$, sample $a_{d,\tau}^{(i)}, a'_{d,\tau}^{(i)} \leftarrow R_{q_d}$. Output

$$\text{params} = \left(\mathbf{a}, \left\{ a_{d,\tau}^{(i)}, a'_{d,\tau}^{(i)} \right\}_{d,\tau,i} \right)$$

- $\text{Keygen}(1^\kappa, i, \text{params})$: Sample D different polynomials $s_0^{(i)}, \dots, s_D^{(i)} \leftarrow \chi_0$ and a vector $\mathbf{x} \leftarrow \chi^n$. Set $\text{sk}_i := (s_0^{(i)}, \dots, s_D^{(i)})$ and $\text{pk}_i := \mathbf{a}s_0 + \mathbf{x}$. The evaluation key is computed as before, using

the $a_{d,\tau}^{(i)}, a'_{d,\tau}^{(i)}$ in **params** instead of sampling them randomly from $R_{q_{d-1}}$. In more detail: for all $d \in [D], \tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$, sample $e_{d,\tau}^{(i)}, e'_{d,\tau}^{(i)} \leftarrow \chi_0$ and compute

$$\begin{aligned}\tilde{c}_{0,d,\tau}^{(i)} &:= a_{d,\tau}^{(i)} + 2e_{d,\tau}^{(i)} + 2^\tau s_{d-1}^{(i)} \in R_{q_{d-1}} & \text{and} & & \tilde{c}_{1,d,\tau}^{(i)} &:= a_{d,\tau}^{(i)} \\ \tilde{z}_{0,d,\tau}^{(i)} &:= a'_{d,\tau}^{(i)} + 2e'_{d,\tau}^{(i)} + 2^\tau (s_{d-1}^{(i)})^2 \in R_{q_{d-1}} & \text{and} & & \tilde{z}_{1,d,\tau}^{(i)} &:= a'_{d,\tau}^{(i)}\end{aligned}$$

Let

$$\begin{aligned}\tilde{\mathbf{c}}_{0,d}^{(i)} &= (\tilde{c}_{0,d,0}^{(i)}, \dots, \tilde{c}_{0,d,\lfloor \log q \rfloor}^{(i)}) & \text{and} & & \tilde{\mathbf{c}}_1^{(i)} &= (\tilde{c}_{1,d,0}^{(i)}, \dots, \tilde{c}_{1,d,\lfloor \log q \rfloor}^{(i)}) \\ \tilde{\mathbf{z}}_{0,d}^{(i)} &= (\tilde{z}_{0,d,0}^{(i)}, \dots, \tilde{z}_{0,d,\lfloor \log q \rfloor}^{(i)}) & \text{and} & & \tilde{\mathbf{z}}_1^{(i)} &= (\tilde{z}_{1,d,0}^{(i)}, \dots, \tilde{z}_{1,d,\lfloor \log q \rfloor}^{(i)})\end{aligned}$$

and set the evaluation key

$$\mathbf{ek}_i := \left\{ \tilde{\mathbf{c}}_{0,d}^{(i)}, \tilde{\mathbf{c}}_{1,d}^{(i)}, \tilde{\mathbf{z}}_{0,d}^{(i)}, \tilde{\mathbf{z}}_{1,d}^{(i)} \right\}_{d \in [D]}$$

- **CombinePK**($\{\mathbf{pk}_i\}_{i \in [N]}\})$: Set $\mathbf{pk} := \mathbf{pk}_1 + \dots + \mathbf{pk}_N$.
- **CombineKeys**($\{\mathbf{sk}_i\}_{i \in [N]}, \{\mathbf{ek}_i\}_{i \in [N]}\})$: Parse $\mathbf{sk}_i = (s_0^{(i)}, \dots, s_D^{(i)})$ and define $s_j = \sum_{i=1}^N s_j^{(i)}$. Set

$$\mathbf{sk} := (s_0, \dots, s_D)$$

Parse $\mathbf{ek}_i = \left\{ \tilde{\mathbf{c}}_{0,d}^{(i)}, \tilde{\mathbf{c}}_{1,d}^{(i)}, \tilde{\mathbf{z}}_{0,d}^{(i)}, \tilde{\mathbf{z}}_{1,d}^{(i)} \right\}_{d \in [D]}$. The evaluation key is computed as in Section 6.3 for secret key $\mathbf{sk} = (s_0, \dots, s_D)$. For all $d \in [D], \tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$, sample $(a_{d,\tau}, b_{d,\tau} = a_{d,\tau}s_d + 2e_{d,\tau}) \in R_{q_{d-1}}^2$ and $(a'_{d,\tau}, b'_{d,\tau} = a'_{d,\tau}s_d + 2e'_{d,\tau}) \in R_{q_{d-1}}^2$, where $a_{d,\tau}, a'_{d,\tau} \leftarrow R_{q_{d-1}}$ and $e_{d,\tau}, e'_{d,\tau} \leftarrow \chi$. Compute

$$\begin{aligned}\tilde{c}_{0,d,\tau} &:= b_{d,\tau} + 2^\tau s_{d-1} \in R_{q_{d-1}} & \text{and} & & \tilde{c}_{1,d,\tau} &:= a_{d,\tau} \\ \tilde{z}_{0,d,\tau} &:= b'_{d,\tau} + 2^\tau s_{d-1}^2 \in R_{q_{d-1}} & \text{and} & & \tilde{z}_{1,d,\tau} &:= a'_{d,\tau}\end{aligned}$$

Let

$$\begin{aligned}\tilde{\mathbf{c}}_{0,d} &= (\tilde{c}_{0,d,0}, \dots, \tilde{c}_{0,d,\lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{c}}_1 &= (\tilde{c}_{1,d,0}, \dots, \tilde{c}_{1,d,\lfloor \log q \rfloor}) \\ \tilde{\mathbf{z}}_{0,d} &= (\tilde{z}_{0,d,0}, \dots, \tilde{z}_{0,d,\lfloor \log q \rfloor}) & \text{and} & & \tilde{\mathbf{z}}_1 &= (\tilde{z}_{1,d,0}, \dots, \tilde{z}_{1,d,\lfloor \log q \rfloor})\end{aligned}$$

and set

$$\mathbf{ek} := \left\{ \tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d} \right\}_{d \in [D]}$$

- **Enc_{sk}**($m, \text{mode}; \mathbf{e}$) : If $\|\mathbf{e}\| \geq B_{\text{mode}}$, output \perp . Otherwise compute

$$c_0 := \langle \mathbf{p}, \mathbf{e} \rangle + m \in R_{q_0} \quad \text{and} \quad c_1 := \langle \mathbf{a}, \mathbf{e} \rangle$$

and output $\mathbf{c} = (c_0, c_1)$. Also recall that when unspecified, we assume $\text{mode} = 0$.

Without loss of generality, assume below that ciphertext \mathbf{c} is a ciphertext under secret key s_D and modulus q_D .

- **ShareDec_{sk_i}**(\mathbf{c}) : Parse $\mathbf{c} = (c_0, c_1)$ and $\mathbf{sk}_i = (s_0^{(i)}, \dots, s_D^{(i)})$. Sample $e^{(i)} \leftarrow \chi_0$ and output

$$\mu_i := c_1 s_D^{(i)} + 2e^{(i)} \in R_{q_D}$$

- $\text{ShareCombine}(\mathbf{c}, \mu_1, \dots, \mu_N)$: Parse $\mathbf{c} = (c_0, c_1)$. Output $c_0 - \sum_{i=1}^N \mu_i \pmod{2}$.
- $\text{ReRand}_{\text{pk}}(\mathbf{c}, 1)$: Sample $\mathbf{e} \leftarrow \chi_1^n$ and output $\hat{\mathbf{c}} \leftarrow \mathbf{c} + \text{Enc}_{\text{pk}}(0; \mathbf{e})$.

Theorem 6.4. *The scheme described above is a key-homomorphic threshold FHE encryption scheme.*

Proof. (Sketch)

Correct Key Combination: It is easily seen that

$$\text{pk} = \text{pk}_1 + \dots + \text{pk}_N = \sum_{i=1}^N (\mathbf{a}s_0^{(i)} + \mathbf{x}^{(i)}) = \mathbf{a} \left(\sum_{i=1}^N s_0^{(i)} \right) + \left(\sum_{i=0}^N \mathbf{x}^{(i)} \right) = \mathbf{a}s_0 + \mathbf{x}'$$

Furthermore, ek is correct by construction.

Key-Simulation Indistinguishability: We define:

$\text{SimCombKeys}(\{(\text{pk}_i, \text{ek}_i)\}_{i \in T})$: Let $\bar{T} = [N] \setminus T$. For $j \in \bar{T}$, lets $\text{pk}_j \leftarrow R_{q_0}$. Sets $\tilde{\mathbf{c}}_{1,d}^{(j)} = (a_{d,1}^{(j)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{(j)})$ and $\tilde{\mathbf{z}}_{1,d}^{(j)} = (a'_{d,1}^{(j)}, \dots, a'_{d, \lfloor \log q_{d-1} \rfloor}^{(j)})$. Chooses $\tilde{\mathbf{c}}_{0,d}^{(j)}, \tilde{\mathbf{c}}_{1,d}^{(j)} \leftarrow R_{q_d}^{\lfloor \log q_{d-1} \rfloor}$, and sets $\text{ek}_j = \left\{ \tilde{\mathbf{c}}_{0,d}^{(j)}, \tilde{\mathbf{c}}_{1,d}^{(j)}, \tilde{\mathbf{z}}_{0,d}^{(j)}, \tilde{\mathbf{z}}_{1,d}^{(j)} \right\}_{d \in [D]}$. Chooses $\text{ek} = (\tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d}) \leftarrow R_{q_{d-1}}^{4 \lfloor \log q_{d-1} \rfloor}$.

We use a hybrid argument, one per each $d \in \{0, \dots, D\}$. First, by PLWE with secret $s_0^{(j)}$, we know that pk_j and $(\tilde{\mathbf{c}}_{0,d}^{(j)}, \tilde{\mathbf{c}}_{1,d}^{(j)}, \tilde{\mathbf{z}}_{0,d}^{(j)}, \tilde{\mathbf{z}}_{1,d}^{(j)})$ as output by SimCombKeys is computationally indistinguishable from an honestly created public key and honestly created tuples. The remaining hybrids are similar. For $d \neq 0$, by PLWE with secret $s_d^{(j)}$, we know that $(\tilde{\mathbf{c}}_{0,d}^{(j)}, \tilde{\mathbf{c}}_{1,d}^{(j)}, \tilde{\mathbf{z}}_{0,d}^{(j)}, \tilde{\mathbf{z}}_{1,d}^{(j)})$ as output by SimCombKeys is computationally indistinguishable from honestly created tuples, and therefore so is ek_j . By PLWE with secret $s_d = \sum_{i=1}^N s_d^{(i)}$, which is computationally hidden from \mathcal{A} (since \mathcal{A} only knows $\{s_d^{(i)}\}_{i \in T}$), we know that ek as output by SimCombKeys is computationally indistinguishable from an honestly created ek .

Correct Share Decryption of Combined Key: Let $\mathbf{c} = (c_0, c_1)$ be an encryption of m under secret key s_D . Then $c_0 = c_1 s_D + 2e + m$ and $\text{Dec}_{\text{sk}}(\mathbf{c}) = m$ as long as $\|e\| \ll q_D/2$. Let $\mu_i = c_1 s_D^{(i)} + 2e^{(i)}$. Then

$$\begin{aligned} \text{ShareCombine}(\mathbf{c}, \mu_1, \dots, \mu_N) &= c_0 - \sum_{i=1}^N \mu_i = (c_1 s_D + 2e + m) - \left(\sum_{i=1}^N c_1 s_D^{(i)} + 2e^{(i)} \right) \pmod{2} \\ &= c_1 s_D - c_1 \sum_{i=1}^N s_D^{(i)} + 2 \left(e - \sum_{i=1}^N e^{(i)} \right) + m \pmod{2} = m, \end{aligned}$$

as long as $\left\| e - \sum_{i=1}^N e^{(i)} \right\| \ll q_D/2$

Share-Simulation Indistinguishability: We define:

$\text{SimShareDec}(\mathbf{c}, m, \{\mu_i\}_{i \in T})$: Parse $\mathbf{c} = (c_0, c_1)$. Fix $j^* \in \bar{T}$. For all $j \neq j^*$, let $\mu_j \leftarrow R_{q_D}$. Let $\mu_{j^*} := c_0 - m - \sum_{i \in [N], i \neq j^*} \mu_i$ and output $\{\mu_j\}_{j \in \bar{T}}$

By correct share decryption, we know that regardless of how $\{\mu_j\}_{j \in \bar{T}}$ were created, $\mu_{j^*} := c_0 - m - \sum_{i \in [N], i \neq j^*} \mu_i$. Since this is a deterministic function of the rest of the variables, we simply need to prove that

$$\left\{ \mu_j := c_1 s_D^{(j)} + 2e^{(j)} \right\}_{j \in \bar{T}, j \neq j^*} \stackrel{c}{\approx} \{u_j \leftarrow R_{q_D}\}_{j \in \bar{T}, j \neq j^*} \quad (1)$$

Because $\mathbf{c} \leftarrow \text{Enc}_{\text{sk}}(m)$, we know that the distribution of c_1 is statistically close to uniform in R_{q_D} and thus, (1) simply follows from a series of hybrid arguments, one for each j , in which we use the pseudo-randomness of a PLWE sample with secret $s_D^{(j)}$.

Semantic Security: Semantic security follows from the pseudorandomness of PLWE. Given any $\{\text{pk}_i\}_{i \in T}$, if $\{\text{pk}_j\}_{j \in \bar{T}}$ are honestly generated then in particular, they are pseudorandom (by PLWE), and thus $\text{pk} = \text{pk}_1 + \dots + \text{pk}_N$ is also pseudorandom. Thus, a ciphertext encrypted under this public key is pseudorandom and semantic security follows.

Ciphertext Re-Randomization: The output of $c = \text{Eval}_{\text{ek}}(c_1, \dots, c_\ell)$ is not indistinguishable from a fresh encryption $c' \leftarrow \text{Enc}_{\text{pk}}(y, 0)$ since the noise in c might be much larger than in c' . We remedy this by adding noise that is super-polynomially bigger than the noise in c (in the form of an encryption of 0, $\text{Enc}_{\text{pk}}(0, 1)$). We choose χ_0 and χ_1 appropriately. Recall that χ_0 and χ_1 are (truncated) discrete Gaussians $D_{\mathbb{Z}^n, r_0, B(r_0)}$ and $D_{\mathbb{Z}^n, r_1, B(r_1)}$, which are statistically close to the discrete Gaussian distributions $D_{\mathbb{Z}^n, r_0}$ and $D_{\mathbb{Z}^n, r_1}$. We choose r_1 to be a super-polynomial factor larger than r_0 : $r_1 \geq 2^{\omega(\log \kappa)} r_0$. Recall that

$$\text{ReRand}(c) = \text{Eval}_{\text{ek}}(c_1, \dots, c_\ell) + \text{Enc}_{\text{pk}}(0, 1)$$

where the first term, $\text{Eval}_{\text{ek}}(c_1, \dots, c_\ell)$ has noise distributed according to $D_{\mathbb{Z}^n, \text{poly}(\kappa) \cdot r_0}$, while the second has noise distributed according to $D_{\mathbb{Z}^n, r_1}$. Because r_1 is super-polynomially bigger than r_0 , we have that the noise distributions of $\text{ReRand}(c)$ and a fresh encryption $\text{Enc}_{\text{pk}}(y, 1)$ are statistically close. □

7 Implementing $\mathcal{F}_{\text{rand.bc}}^{\mathcal{E}}$ and $\mathcal{F}_{\text{keys}}^{\mathcal{E}}$

We show how to implement $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ and $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ for our key-homomorphic threshold FHE scheme.

7.1 Ciphertext Re-Randomization

Protocol $\Pi_{\text{RAND}}^{\text{SH}}$ shown in Figure 11 securely implements $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ in the \mathcal{F}_{BC} -hybrid model, against (static) semi-honest adversaries. In $\Pi_{\text{RAND}}^{\text{SH}}$, we simply ask each party to send a random encryption of 0. Since corrupted parties are semi-honest, this re-randomizes the ciphertext.

Theorem 7.1. *Protocol $\Pi_{\text{RAND}}^{\text{SH}}$ described in Figure 11 securely implements $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ against (static) semi-honest adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}})$ -hybrid model.*

The proof of Theorem 7.1 is given in Appendix A.2. Protocol $\Pi_{\text{RAND}}^{\text{MAL}}$ shown in Figure 12 securely implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{zero})})$ -hybrid model. To tolerate malicious adversaries, we need to guarantee that the encryption z_i is a “fresh” encryption; in particular, we need to guarantee that z_i is not correlated to any other z_j . In order to guarantee this, we place a commitment key for a trapdoor

Protocol $\Pi_{\text{RAND}}^{\text{SH}}$

Inputs: All parties P_1, \dots, P_N and server S get public key pk as input. Server S additionally gets a ciphertext $c \in \mathcal{C}_{\text{pk}}$, where \mathcal{C}_{pk} is the ciphertext space defined by pk .

For session ID sid :

Round 1: Each party P_i samples $z_i \leftarrow \text{Enc}_{\text{pk}}(0, 1)$ and sends (sid, i, z_i) to \mathcal{F}_{BC} .

Round 2: Upon receiving (sid, i, z_i) for all $i \in [N]$ from \mathcal{F}_{BC} , server computes $\hat{c} := c + \sum_{i=1}^N z_i$ and sends $(\text{sid}, 0, c, \hat{c})$ to \mathcal{F}_{BC} .

Figure 11: Protocol $\Pi_{\text{RAND}}^{\text{SH}}$. Implements $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ in the \mathcal{F}_{BC} -hybrid model. Secure against semi-honest adversaries.

commitment scheme $(\text{Keygen}^{\text{Com}}, \text{Com})$ in the crs . This way we know that it is correctly distributed. We then have each party commit in advance, using ck , to the randomness $r_i^{(0)}$ they will use to create z_i . Along with z_i parties also provide a SS-NIZK proof that z_i was indeed computed with randomness $r_i^{(0)}$. Formally, we require each party to compute a SS-NIZK for the NP-relation:

$$R^{(\text{zero})} = \{ (\rho, \text{ck}, z, \text{pk}), (r, s) \mid \rho \leftarrow \text{Com}_{\text{ck}}(r; s) \wedge z := \text{Enc}_{\text{pk}}(0, 1; r) \}$$

Protocol $\Pi_{\text{RAND}}^{\text{MAL}}$

Inputs: All parties P_1, \dots, P_N and server S get public key pk as input. Server S additionally gets a ciphertext $c \in \mathcal{C}_{\text{pk}}$, where \mathcal{C}_{pk} is the ciphertext space defined by pk .

For session ID sid :

Round 1 Each party P_i :

- Calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with sid and receives (sid, ck) . Chooses encryption and commitment randomness $r_i^{(0)}$ and $s_i^{(0)}$, and computes a commitment $\rho_i^{(0)} \leftarrow \text{Com}(r_i^{(0)}; s_i^{(0)})$.
- Sends $(\text{sid}, i, \rho_i^{(0)})$ to \mathcal{F}_{BC} .

Round 2: Each party P_i , samples $z_i \leftarrow \text{Enc}_{\text{pk}}(0, 1; r_i^{(0)})$ and sends $(\text{sid}, i, (\rho_i^{(0)}, z_i, \text{pk}), (r_i^{(0)}, s_i^{(0)}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{zero})}}$.

Round 3: The server S :

- Upon receiving $(\text{sid}, i, (\rho_i^{(0)}, z_i, \text{pk}))$ from $\mathcal{F}_{\text{ZK.BC}}^{R^{(\text{zero})}}$ for every $i \in [N]$, compute $\hat{c} := c + \sum_{i=1}^N z_i$ and send $(\text{sid}, S, c, \hat{c})$ to \mathcal{F}_{BC} .

Figure 12: Protocol $\Pi_{\text{RAND}}^{\text{MAL}}$. Implements $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{zero})}})$ -hybrid model, where \mathcal{D} is the distribution defined by $\text{Keygen}^{\text{Com}}$. Secure against malicious adversaries.

Theorem 7.2. *Protocol $\Pi_{\text{RAND}}^{\text{MAL}}$ described in Figure 12 securely implements $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$ against (static) malicious adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{F}_{\text{ZK.BC}}^{R^{(\text{zero})}})$ -hybrid model.*

The proof of Theorem 7.2 is given in Appendix A.2.

7.2 Key Combination

Protocol $\Pi_{\text{KEYS}}^{\text{SH}}$ shown in Figure 13 securely implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ in the $\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ -hybrid model, against (static) semi-honest adversaries. We wish to obtain pseudoencryptions of s_{d-1} and $(s_{d-1})^2$, where s_{d-1} is the combined secret key $s_d = \sum_{i=1}^N s_{d-1}^{(i)}$. In $\Pi_{\text{KEYS}}^{\text{SH}}$, parties begin by sending encryptions of

0, $w_{d,\tau}^{(k)}, w'_{d,\tau}^{(k)}$ using each of the $a_{d,\tau}^{(k)}, a'_{d,\tau}^{(k)}$ in **params**. Using these and the pseudoencryptions in the evaluation keys ek_k , by key-homomorphism they can then compute pseudoencryptions $\tilde{x}_{d,\tau}^{(k)}$ of $2^\tau s_{d-1}^{(k)}$ under the *combined* secret key s_d . Adding these yields a pseudoencryption of $2^\tau s_{d-1}$. Furthermore, multiplying these by $s_{d-1}^{(i)}$ and re-randomizing, yields a pseudoencryption of $2^\tau s_{d-1}^{(i)} s_{d-1}^{(k)}$, under the combined secret key s_d . Adding these pseudoencryptions for all k, i yields a pseudoencryption of $2^\tau (s_d)^2$:

$$2^\tau \sum_{i,k=1}^N s_{d-1}^{(i)} s_{d-1}^{(k)} = 2^\tau \left(\sum_i s_d^{(i)} \right)^2 = 2^\tau (s_d)^2$$

Theorem 7.3. *Protocol $\Pi_{\text{KEYS}}^{\text{SH}}$ described in Figure 13 securely implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ against (static) semi-honest adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}})$ -hybrid model.*

The proof of Theorem 7.3 is given in Appendix A.3. Figure 14 describes protocol $\Pi_{\text{KEYS}}^{\text{MAL}}$ that securely implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{enc})}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{rand})})$ -hybrid model, against (static) malicious adversaries. To tolerate malicious adversaries, all we have to change is to have the parties provide SS-NIZK proofs that they computed the $w_{d,\tau}^{(i)}$'s correctly, as encryptions of 0, using secret $s_d^{(i)}$, and also that multiplied in $s_{d-1}^{(i)}$ and re-randomized correctly. More formally, they provide SS-NIZK proofs for the following NP relations:

$$\begin{aligned} R^{(\text{enc})} &= \{ (a, c), (s, e) \mid c = as + 2e \} \\ R^{(\text{rand})} &= \{ (\mathbf{c}, \mathbf{c}', \text{pk}), (s, r) \mid \mathbf{c}' = \mathbf{s}\mathbf{c} + \text{Enc}_{\text{pk}}(0; r) \} \end{aligned}$$

Theorem 7.4. *Protocol $\Pi_{\text{KEYS}}^{\text{MAL}}$ described in Figure 14 securely implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ against (static) malicious adversaries, in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{enc})}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{rand})})$ -hybrid model.*

The proof of Theorem 7.4 is given in Appendix A.3.

8 Performance

In this section we argue that the protocol from Section 4 can be executed in 4 rounds in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}})$ -hybrid model. We also analyze the communication complexity and the computation time of the parties P_i and server S . We first look at the semi-honest case and then turn our attention to the changes needed to the protocol secure against malicious adversaries.

In Section 7, we described protocols $\Pi_{\text{RAND}}^{\text{SH}}$ and $\Pi_{\text{KEYS}}^{\text{SH}}$ that securely implement the functionality $\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}$, each in 2 rounds. Since the circuit C is never used in these protocols, the communication complexity and the computation time of all parties and the server in $\Pi_{\text{RAND}}^{\text{SH}}$ and $\Pi_{\text{KEYS}}^{\text{SH}}$ is independent of $|C|$. We want to run these protocols inside Π^{SH} in the least number of rounds, subject to the following restrictions:

- Parties must know pk before Round 1 of $\Pi_{\text{RAND}}^{\text{SH}}$.
- Round 2 of $\Pi_{\text{RAND}}^{\text{SH}}$ must be placed in Phase 3 of Π^{SH} .
- Server must know ek in (and therefore $\Pi_{\text{KEYS}}^{\text{SH}}$ must be completed by) Phase 3 of Π^{SH} , so that the server can perform homomorphic operations.

We do this by sending some of the messages in parallel:

Protocol $\Pi_{\text{KEYS}}^{\text{SH}}$

Inputs: All parties P_i get as input a valid key tuple $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$, and for $j \neq i$, P_i and S also receive $(\text{pk}_j, \text{ek}_j)$.

For session ID sid , each party P_i :

Round 1: Calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with sid and receives $(\text{sid}, \text{params})$, where $\text{params} = \left(\dots, \left\{ a_{d,\tau}^{(k)}, a_{d,\tau}^{\prime(k)} \right\}_{d,\tau,k}, \dots \right)$.

For $k \in [N]$:

- Parses $\text{ek}_k := \left\{ \tilde{\mathbf{c}}_{0,d}^{(k)}, \tilde{\mathbf{c}}_{1,d}^{(k)}, \tilde{\mathbf{z}}_{0,d}^{(k)}, \tilde{\mathbf{z}}_{1,d}^{(k)} \right\}_{d \in [D]}$.
- Verifies that $\tilde{\mathbf{c}}_{1,d}^{(k)} = (a_{d,1}^{(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{(k)})$ and $\tilde{\mathbf{z}}_{1,d}^{(k)} = (a_{d,1}^{\prime(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{\prime(k)})$
- Samples $e_{d,\tau}^{(i,k)} \leftarrow \chi_0$ and computes

$$w_{d,\tau}^{(i,k)} := a_{d,\tau}^{(k)} s_d^{(i)} + 2e_{d,\tau}^{(i,k)}$$

- Sends $(\text{sid}, i, w_{d,\tau}^{(i,k)})$ to \mathcal{F}_{BC} .

Round 2: Upon receiving $(\text{sid}, j, w_{d,\tau}^{(j,k)})$ from \mathcal{F}_{BC} for all $j, k \in [N]$, $d \in [D]$, $\tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$:

- Computes

$$\tilde{x}_{0,d,\tau}^{(k)} := \sum_{j \neq k} w_{d,\tau}^{(j,k)} + c_{0,d,\tau}^{(k)} = a_{d,\tau}^{(k)} s_d + 2 \left(\sum_{j \neq k} e_{d,\tau}^{(j,k)} + e_{d,\tau}^{(k)} \right) + 2^\tau s_{d-1}^{(k)} \quad \text{and} \quad \tilde{x}_{1,d,\tau}^{(k)} = a_{d,\tau}^{(k)}$$

$$\tilde{y}_{0,d,\tau}^{(i,k)} = s_{d-1}^{(i)} \tilde{x}_{d,\tau}^{(k)} \quad \text{and} \quad \tilde{y}_{1,d,\tau}^{(i,k)} = s_{d-1}^{(i)} a_{d,\tau}^{(k)}$$

- Let $\tilde{\mathbf{x}}_{d,\tau}^{(k)} = (\tilde{x}_{0,d,\tau}^{(k)}, \tilde{x}_{1,d,\tau}^{(k)})$ and $\tilde{\mathbf{y}}_{d,\tau}^{(i,k)} = (\tilde{y}_{0,d,\tau}^{(i,k)}, \tilde{y}_{1,d,\tau}^{(i,k)})$. Computes

$$\tilde{\mathbf{c}}_{d,\tau}^{(k)} = (\tilde{c}_{0,d,\tau}^{(k)}, \tilde{c}_{1,d,\tau}^{(k)}) \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0)$$

$$\tilde{\mathbf{z}}_{d,\tau}^{(i,k)} = (\tilde{z}_{0,d,\tau}^{(i,k)}, \tilde{z}_{1,d,\tau}^{(i,k)}) \leftarrow \tilde{\mathbf{y}}_{d,\tau}^{(i,k)} + \text{Enc}_{\text{pk}}(0)$$

- Sends $(\text{sid}, i, \tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(i,k)})$ to \mathcal{F}_{BC} .

(Local Computation): Upon receiving $(\text{sid}, j, \tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)})$ from \mathcal{F}_{BC} for all $j, k \in [N]$, $d \in [D]$, $\tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$:

- Computes

$$\tilde{c}_{0,d,\tau} := \sum_k \tilde{c}_{0,d,\tau}^{(k)} \quad \text{and} \quad \tilde{c}_{1,d,\tau} := \sum_k \tilde{c}_{1,d,\tau}^{(k)}$$

$$\tilde{z}_{0,d,\tau} := \sum_{j,k} \tilde{z}_{0,d,\tau}^{(j,k)} \quad \text{and} \quad \tilde{z}_{1,d,\tau} := \sum_{j,k} \tilde{z}_{1,d,\tau}^{(j,k)}$$

- For $b \in \{0, 1\}$, sets $\tilde{\mathbf{c}}_{b,d} := (\tilde{c}_{b,d,1}, \dots, \tilde{c}_{b,d, \lfloor \log q_{d-1} \rfloor})$ and $\tilde{\mathbf{z}}_{b,d} := (\tilde{z}_{b,d,1}, \dots, \tilde{z}_{b,d, \lfloor \log q_{d-1} \rfloor})$, and $\text{ek} := (\tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d})$.

Figure 13: Protocol $\Pi_{\text{KEYS}}^{\text{SH}}$. Implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}})$ -hybrid model, where \mathcal{D} is the distribution defined by Setup. Secure against semi-honest adversaries.

Round 1: Each party computes its key tuple, as in Π^{SH} , and also runs Round 1 of $\Pi_{\text{KEYS}}^{\text{SH}}$

Round 2: Each party computes an encryption of its input, as in Π^{SH} . Since all parties know pk , also runs Round 1 of $\Pi_{\text{RAND}}^{\text{SH}}$ and Round 2 of $\Pi_{\text{KEYS}}^{\text{SH}}$.

Protocol $\Pi_{\text{KEYS}}^{\text{MAL}}$

Inputs: All parties P_i get as input a valid key tuple $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$, and for $j \neq i$, P_i and S also receive $(\text{pk}_j, \text{ek}_j)$.

For session ID sid , each party P_i :

Round 1: Calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with sid and receives $(\text{sid}, \text{params})$, where $\text{params} = \left(\dots, \left\{ a_{d,\tau}^{(k)}, a_{d,\tau}^{\prime(k)} \right\}_{d,\tau,k}, \dots \right)$.

For $k \in [N]$,

- Parses $\text{ek}_k := \left\{ \tilde{\mathbf{c}}_{0,d}^{(k)}, \tilde{\mathbf{c}}_{1,d}^{(k)}, \tilde{\mathbf{z}}_{0,d}^{(k)}, \tilde{\mathbf{z}}_{1,d}^{(k)} \right\}_{d \in [D]}$.
- Verifies that $\tilde{\mathbf{c}}_{1,d}^{(k)} = (a_{d,1}^{(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{(k)})$ and $\tilde{\mathbf{z}}_{1,d}^{(k)} = (a_{d,1}^{\prime(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{\prime(k)})$
- Samples $e_{d,\tau}^{(i,k)} \leftarrow \chi_0$ and computes

$$w_{d,\tau}^{(i,k)} := a_{d,\tau}^{(k)} s_d^{(i)} + 2e_{d,\tau}^{(i,k)}$$

- Sends $(\text{sid}, i, (w_{d,\tau}^{(i,k)}, a_{d,\tau}^{(i,k)}), (s_d^{(i)}, e_{d,\tau}^{(i,\tau)}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R(\text{enc})}$.

Round 2: Upon receiving $(\text{sid}, j, (w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}))$ from $\mathcal{F}_{\text{ZK.BC}}^{R(\text{enc})}$ for all $j, k \in [N]$, $d \in [D]$, $\tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$:

- Computes

$$\tilde{x}_{0,d,\tau}^{(k)} := \sum_{j \neq k} w_{d,\tau}^{(j,k)} + c_{0,d,\tau}^{(k)} = a_{d,\tau}^{(k)} s_d + 2 \left(\sum_{j \neq k} e_{d,\tau}^{(j,k)} + e_{d,\tau}^{(k)} \right) + 2^\tau s_{d-1}^{(k)} \quad \text{and} \quad \tilde{x}_{1,d,\tau}^{(k)} = a_{d,\tau}^{(k)}$$

$$\tilde{y}_{0,d,\tau}^{(i,k)} = s_{d-1}^{(i)} \tilde{x}_{d,\tau}^{(k)} \quad \text{and} \quad \tilde{y}_{1,d,\tau}^{(i,k)} = s_{d-1}^{(i)} a_{d,\tau}^{(k)}$$

- Let $\tilde{\mathbf{x}}_{d,\tau}^{(k)} = (\tilde{x}_{0,d,\tau}^{(k)}, \tilde{x}_{1,d,\tau}^{(k)})$ and $\tilde{\mathbf{y}}_{d,\tau}^{(i,k)} = (\tilde{y}_{0,d,\tau}^{(i,k)}, \tilde{y}_{1,d,\tau}^{(i,k)})$. Chooses randomness $r_{d,\tau}^{(k)}, r_{d,\tau}^{(i,k)}$ and computes

$$\tilde{\mathbf{c}}_{d,\tau}^{(k)} = (\tilde{c}_{0,d,\tau}^{(k)}, \tilde{c}_{1,d,\tau}^{(k)}) \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(k)})$$

$$\tilde{\mathbf{z}}_{d,\tau}^{(i,k)} = (\tilde{z}_{0,d,\tau}^{(i,k)}, \tilde{z}_{1,d,\tau}^{(i,k)}) \leftarrow \tilde{\mathbf{y}}_{d,\tau}^{(i,k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(i,k)})$$

- Sends $(\text{sid}, i, (\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}), (1, r_{d,\tau}^{(k)}))$ and $(\text{sid}, i, (\tilde{\mathbf{z}}_{d,\tau}^{(i,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}), (s_d^{(i)}, r_{d,\tau}^{(i,k)}))$ to $\mathcal{F}_{\text{ZK.BC}}^{R(\text{rand})}$.

(Local Computation): Upon receiving $(\text{sid}, j, (\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}))$ and $(\text{sid}, j, (\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}))$ from $\mathcal{F}_{\text{ZK.BC}}^{R(\text{rand})}$ for all $j, k \in [N]$, $d \in [D]$, $\tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$:

- Computes

$$\tilde{c}_{0,d,\tau} := \sum_k \tilde{c}_{0,d,\tau}^{(k)} \quad \text{and} \quad \tilde{c}_{1,d,\tau} := \sum_k \tilde{c}_{1,d,\tau}^{(k)}$$

$$\tilde{z}_{0,d,\tau} := \sum_{j,k} \tilde{z}_{0,d,\tau}^{(j,k)} \quad \text{and} \quad \tilde{z}_{1,d,\tau} := \sum_{j,k} \tilde{z}_{1,d,\tau}^{(j,k)}$$

- For $b \in \{0, 1\}$, sets $\tilde{\mathbf{c}}_{b,d} := (\tilde{c}_{b,d,1}, \dots, \tilde{c}_{b,d, \lfloor \log q_{d-1} \rfloor})$ and $\tilde{\mathbf{z}}_{b,d} := (\tilde{z}_{b,d,1}, \dots, \tilde{z}_{b,d, \lfloor \log q_{d-1} \rfloor})$, and $\text{ek} := (\tilde{\mathbf{c}}_{0,d}, \tilde{\mathbf{c}}_{1,d}, \tilde{\mathbf{z}}_{0,d}, \tilde{\mathbf{z}}_{1,d})$.

Figure 14: Protocol $\Pi_{\text{KEYS}}^{\text{MAL}}$. Implements $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\mathcal{D}}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{enc})}, \mathcal{F}_{\text{ZK.BC}}^{R(\text{rand})})$ -hybrid model, where \mathcal{D} is the distribution defined by **Setup**. Secure against malicious adversaries.

Round 3: The server knows ek because $\Pi_{\text{KEYS}}^{\text{SH}}$ was completed in the previous round. The server performs the homomorphic evaluation, as in Π^{SH} , and runs Round 2 or $\Pi_{\text{RAND}}^{\text{SH}}$.

Round 4: The parties share decryption shares as in $\Pi_{\text{KEYS}}^{\text{SH}}$.

(Local Computation): The parties combine the ciphertext and decryption shares to obtain the output.

It is clear that the communication complexity and the computation time of the parties in the resulting protocol is *independent* of $|C|$, and polynomial in $(\sum |x_i|, |y|, N, \kappa)$. The computation time of the server, on the other hand, is linear in $|C|$ and polynomial in $(\sum |x_i|, |y|, N, \kappa)$.

In the malicious case, $\Pi_{\text{KEYS}}^{\text{MAL}}$ has exactly the same number of rounds and same restrictions as $\Pi_{\text{KEYS}}^{\text{SH}}$, but the protocols Π^{MAL} and $\Pi_{\text{RAND}}^{\text{MAL}}$ require one more round than their semi-honest counterparts, where the parties commit to the randomness they will use throughout the protocol. Therefore, we can achieve 5 rounds, in exactly the same way as before, except that we have one more round in the beginning, where the parties send commitments to randomness. Furthermore, the server must provide and the parties verify a short argument for the validity of c , that is, for the correct homomorphic evaluation of $|C|$. As we saw in Section 5, this can be done in one of four ways. Table 1 summarizes the round and communication complexity of the resulting protocol for each of the 4 options, together the computation time of the parties. The computation time of the server is always polynomial in $|C|$, as well as in $(\sum |x_i|, |y|, N, \kappa)$.

| VC Proof | RC | CC(poly(\cdot)) | CT(poly(\cdot)) | Model / Assumptions |
|-----------------|----|---------------------|---------------------|---------------------|
| [Kil92, Kil95] | 7 | $\log(C)$ | $\log(C)$ | CRHF |
| [Mic94] | 5 | $\log(C)$ | $\log(C)$ | Random Oracle |
| [BCCT11, GLR11] | 5 | $\log(C)$ | $\log(C)$ | Non-Falsifiable |
| [GKR08] | 5 | $d, \log(C)$ | $d, \log(C)$ | Secure PIR Scheme |

Table 1: Performance of Π^{MAL} in the $(\mathcal{F}_{\text{BC}}, \mathcal{F}_{\text{CRS}}^{\text{D}})$ -hybrid model, for N clients and circuit C of depth d . Communication complexity (CC) and computation time (CT) displayed are in addition to $\text{poly}(\sum |x_i|, |y|, N, \kappa)$.

References

- [AiCPS09] Benny Applebaum, David Cash inand Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618. Springer, 2009.
- [BCCT11] Nir Bitansky, Ran Canetti, Alessandro Chiesa, and Eran Tromer. From extractable collision resistance to succinct non-interactive arguments of knowledge, and back again. *Cryptology ePrint Archive: Report 2011/443*, 2011.
- [BD10] Rikke Bendlin and Ivan Damgård. Threshold decryption and zero-knowledge proofs for lattice-based cryptosystems. In Daniele Micciancio, editor, *TCC*, volume 5978 of *Lecture Notes in Computer Science*, pages 201–218. Springer, 2010.
- [BDOZ11] Rikke Bendlin, Ivan Damgård, Claudio Orlandi, and Sarah Zakarias. Semi-homomorphic encryption and multiparty computation. In Kenneth G. Paterson, editor, *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 169–188. Springer, 2011.

- [BFKL93] Avrim Blum, Merrick L. Furst, Michael J. Kearns, and Richard J. Lipton. Cryptographic primitives based on hard learning problems. In Douglas R. Stinson, editor, *CRYPTO*, volume 773 of *Lecture Notes in Computer Science*, pages 278–291. Springer, 1993.
- [BGV11] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully homomorphic encryption without bootstrapping. *Manuscript (in submission)*, 2011.
- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *STOC*, pages 1–10, 1988.
- [BV11a] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) lwe. *To Appear in FOCS*, 2011.
- [BV11b] Zvika Brakerski and Vinod Vaikuntanathan. Fully homomorphic encryption from ring-lwe and security for key dependent messages. In Rogaway [Rog11], pages 505–524.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *STOC*, pages 11–19, 1988.
- [CDN01] Ronald Cramer, Ivan Damgård, and Jesper Buus Nielsen. Multiparty computation from threshold homomorphic encryption. In *EUROCRYPT*, pages 280–299, 2001.
- [CLOS02] Ran Canetti, Yehuda Lindell, Rafail Ostrovsky, and Amit Sahai. Universally composable two-party and multi-party secure computation. In *STOC*, pages 494–503, 2002.
- [DPSZ11] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. *Cryptology ePrint Archive*, Report 2011/535, 2011. <http://eprint.iacr.org/>.
- [Gen09a] Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. crypto.stanford.edu/craig.
- [Gen09b] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *STOC*, pages 169–178. ACM, 2009.
- [Gil10] Henri Gilbert, editor. *Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, French Riviera, May 30 - June 3, 2010. Proceedings*, volume 6110 of *Lecture Notes in Computer Science*. Springer, 2010.
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Cynthia Dwork, editor, *STOC*, pages 113–122. ACM, 2008.
- [GLR11] Shafi Goldwasser, Huijia Lin, and Aviad Rubinfeld. Delegation of computation without rejection problem from designated verifier cs-proofs. *Cryptology ePrint Archive: Report 2011/456*, 2011.
- [GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In *STOC*, pages 218–229, 1987.
- [HLP11] Shai Halevi, Yehuda Lindell, and Benny Pinkas. Secure computation on the web: Computing without simultaneous interaction. In Rogaway [Rog11], pages 132–150.

- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *STOC*, pages 723–732. ACM, 1992.
- [Kil95] Joe Kilian. Improved efficient arguments (preliminary version). In Don Coppersmith, editor, *CRYPTO*, volume 963 of *Lecture Notes in Computer Science*, pages 311–324. Springer, 1995.
- [KMR11] Seny Kamara, Payman Mohassel, and Mariana Raykova. Outsourcing multi-party computation. Cryptology ePrint Archive, Report 2011/272, 2011. <http://eprint.iacr.org/>.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Gilbert [Gil10], pages 1–23.
- [Mic94] Silvio Micali. Cs proofs (extended abstracts). In *FOCS*, pages 436–453. IEEE, 1994.
- [MSas11] Steven Myers, Mona Sergi, and abhi shelat. Threshold fully homomorphic encryption and secure computation. Cryptology ePrint Archive, Report 2011/454, 2011. <http://eprint.iacr.org/>.
- [Nao03] Moni Naor. On cryptographic assumptions and challenges. In Dan Boneh, editor, *CRYPTO*, volume 2729 of *Lecture Notes in Computer Science*, pages 96–109. Springer, 2003.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *STOC*, pages 84–93. ACM, 2005.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
- [Rog11] Phillip Rogaway, editor. *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*. Springer, 2011.
- [Sah99] Amit Sahai. Non-malleable non-interactive zero knowledge and adaptive chosen-ciphertext security. In *FOCS*, pages 543–553, 1999.
- [SP92] Alfredo De Santis and Giuseppe Persiano. Zero-knowledge proofs of knowledge without interaction (extended abstract). In *FOCS*, pages 427–436, 1992.
- [SV10] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. In *Public Key Cryptography*, pages 420–443, 2010.
- [vDGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Gilbert [Gil10], pages 24–43.

A Proofs

A.1 Security Under Malicious Adversaries

Proof of Lemma 5.1: Correctness follows from correctness of key combination, correctness of homomorphic evaluation, correctness of share decryption of \mathcal{E} , and correctness of the NIZK proof systems $\Pi^{(x)}$, $\Pi^{(gen)}$, $\Pi^{(dec)}$ and the short argument system Φ .

We now turn to security. To show more clearly why we need simulation-sound and simulation-extractable NIZKs (as opposed to standard NIZKs), we instantiate in Π^{MAL} from Figure 10, $\mathcal{F}_{\text{ZK.BC}}^{R(\text{gen})}$, $\mathcal{F}_{\text{ZK.BC}}^{R(\text{dec})}$ with SS-NIZKs, and $\mathcal{F}_{\text{ZK.BC}}^{R(x)}$ with a SE-NIZK and present the proof of Lemma 5.1 after this instantiation. Let \mathcal{A} be a real-world malicious adversary corrupting $t < N$ clients, and possibly the server. Let $T \subsetneq [N]$ be the set of corrupted clients. We describe the simulator \mathcal{S}^{MAL} in Figure 15. We prove that $\text{REAL}_{\Pi^{\text{MAL}}, \mathcal{A}}(\vec{x}) \stackrel{c}{\approx} \text{IDEAL}_{\mathcal{F}, \mathcal{S}^{\text{MAL}}}(\vec{x})$ via a series of hybrids. The view of the \mathcal{A} consists of

$$\text{View} = \left(\text{params}, \left\{ x_i, \rho_i^{(\text{gen})}, \rho_i^{(\text{dec})}, (\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(\text{gen})}, c_i, \pi_i^{(x)}, \mu_i, \pi_i^{(\text{dec})} \right\}_{i \in T}, \right. \\ \left. \left\{ (\rho_j^{(\text{gen})}, \rho_j^{(\text{dec})}) \right\}_{j \in \bar{T}}, \text{pk}, \text{ek}, \left\{ (\text{pk}_j, \text{ek}_j), \pi_j^{(\text{gen})} \right\}_{j \in \bar{T}}, \left\{ c_j, \pi_j^{(x)} \right\}_{j \in \bar{T}}, c, \varphi, \hat{c}, \left\{ \mu_j, \pi_j^{(\text{dec})} \right\}_{j \in \bar{T}} \right)$$

We simply prove that for all $\{x_i, \rho_i^{(\text{gen})}, \rho_i^{(\text{dec})}, (\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(\text{gen})}, c_i, \pi_i^{(x)}, \mu_i, \pi_i^{(\text{dec})}\}_{i \in T}$ output by \mathcal{A} , the joint distribution of these variables in the simulation is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol Π^{MAL} . We let View_i be the joint distribution of these variables in hybrid i . In all hybrids, $\text{params} \leftarrow \text{Setup}(1^\kappa, N)$, $(\text{crs}^{(x)}, \text{tk}^{(x)}, \text{xtk}) \leftarrow \text{Setup}^{\text{X}}(1^\kappa)$, $(\text{crs}^{(\text{dec})}, \text{tk}^{(\text{dec})}) \leftarrow \text{Setup}^{\text{Gen}}(1^\kappa)$, $(\text{crs}^{(\text{gen})}, \text{tk}^{(\text{gen})}) \leftarrow \text{Setup}^{\text{Dec}}(1^\kappa)$, $(\text{ck}, \text{xtk}^{(\text{com})}) \leftarrow \text{Keygen}^{\text{Com}}(1^\kappa)$, $\text{pk} \leftarrow \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$. Also, we can assume $c = \text{Eval}(C, c_1, \dots, c_N)$, even when the server is corrupted, as this is guaranteed by soundness of the short argument system Φ . We gradually change how each of the other variables is computed.

Hybrid 0: This is the real-world execution of Π^{MAL} . We have:

$$\rho_j^{(\text{gen})} \leftarrow \text{Com}_{\text{ck}}(r_j^{(\text{gen})}; s_j^{(\text{gen})}) \quad , \quad \rho_j^{(\text{dec})} \leftarrow \text{Com}_{\text{ck}}(r_j^{(\text{dec})}; s_j^{(\text{dec})}) \\ \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j; r_j^{(\text{gen})}) \right\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(\text{gen})} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(\text{gen})}, j), (r_j^{(\text{gen})}, s_j^{(\text{dec})})) \right\}_{j \in \bar{T}} \\ (\cdot, \text{ek}) := \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(x)} \leftarrow \text{Prove}^{\text{X}}(c_j, \text{pk}, (x_j, r_j^{(x)})) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{ReRand}(c) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(x)} \leftarrow \text{Prove}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(\text{gen})}, \rho_j^{(\text{dec})}, \text{ck}, j), (r_j^{(\text{dec})}, s_j^{(\text{dec})}, r_j^{(\text{gen})}, s_j^{(\text{gen})}) \right\}_{j \in \bar{T}}$$

Hybrid 1: We now change the proofs $\pi_j^{(x)}$ so that they are simulated proofs, created with the Sim^{X} algorithm, instead of created honestly using Prove^{X} .

$$\rho_j^{(\text{gen})} \leftarrow \text{Com}_{\text{ck}}(r_j^{(\text{gen})}; s_j^{(\text{gen})}) \quad , \quad \rho_j^{(\text{dec})} \leftarrow \text{Com}_{\text{ck}}(r_j^{(\text{dec})}; s_j^{(\text{dec})}) \\ \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j; r_j^{(\text{gen})}) \right\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(\text{gen})} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(\text{gen})}, j), (r_j^{(\text{gen})}, s_j^{(\text{dec})})) \right\}_{j \in \bar{T}} \\ (\cdot, \text{ek}) := \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{ReRand}(c) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \\ \left\{ \pi_j^{(x)} \leftarrow \text{Prove}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(\text{gen})}, \rho_j^{(\text{dec})}, \text{ck}, j), (r_j^{(\text{dec})}, s_j^{(\text{dec})}, r_j^{(\text{gen})}, s_j^{(\text{gen})}) \right\}_{j \in \bar{T}}$$

We have that $\text{View}_0 \stackrel{c}{\approx} \text{View}_1$ by the *zero-knowledge* property of $\Pi^{(x)}$.

Simulator \mathcal{S}^{MAL}

1. Run $\text{params} \leftarrow \text{Setup}(1^\kappa, N)$, $(\text{crs}^{(x)}, \text{tk}^{(x)}, \text{xtk}) \leftarrow \text{Setup}^X(1^\kappa)$, $(\text{crs}^{(dec)}, \text{tk}^{(dec)}) \leftarrow \text{Setup}^{\text{Gen}}(1^\kappa)$, $(\text{crs}^{(gen)}, \text{tk}^{(gen)}) \leftarrow \text{Setup}^{\text{Dec}}(1^\kappa)$, and $(\text{ck}, \text{xtk}^{(com)}) \leftarrow \text{Keygen}^{\text{Com}}(1^\kappa)$. Let $\text{params}' := (\text{params}, \text{crs}^{(x)}, \text{crs}^{(dec)}, \text{crs}^{(gen)}, \text{ck})$ and run \mathcal{A} on inputs $\{x_i\}_{i \in T}$ and params' while playing the role of the honest parties:

Phase 1: For all $j \in \bar{T}$, compute $\rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(0)$ and $\rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0)$. Send $(\text{sid}, j, \rho_j^{(gen)}, \rho_j^{(dec)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Phase 2:

- Upon receiving $(\text{sid}, i, \rho_i^{(gen)}, \rho_i^{(dec)})$ as a message to $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ for all $i \in T$, use the commitment trapdoor to extract randomness $r_i^{(gen)}, r_i^{(dec)}$:

$$r_i^{(gen)} := \text{Ext}_{\text{xtk}^{(com)}}^{\text{Com}}(\rho_i^{(gen)}) \quad , \quad r_i^{(dec)} := \text{Ext}_{\text{xtk}^{(com)}}^{\text{Com}}(\rho_i^{(dec)})$$

- For all $i \in T$, let $(\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_i, \widetilde{\text{ek}}_i) \leftarrow \text{Keygen}(1^\kappa, i; r_i^{(gen)})$.
- For all $j \in \bar{T}$, compute $\pi_j^{(gen)} \leftarrow \text{Sim}_{\text{tk}^{(dec)}}^{\text{Gen}}(\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j)$ and send $(\text{sid}, j, (\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j), \pi_j^{(gen)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .
- Upon receiving $(\text{sid}, i, (\text{pk}_i, \text{ek}_i, \text{ck}, \rho_i^{(gen)}, i), \pi_i^{(gen)})$ as a message to \mathcal{F}_{BC} for all $i \in T$, verify that $\pi_i^{(gen)}$ is valid. Upon receiving $(\text{sid}, i, (\text{pk}_i, \text{sk}_i, \text{ek}_i))$ as a message to $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$ for all $i \in T$, check that $(\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_i, \widetilde{\text{ek}}_i) = (\text{pk}_i, \text{sk}_i, \text{ek}_i)$ and halt if the key tuples are not the same.
- Let $\text{pk} := \text{CombinePK}(\{\text{pk}_j\}_{j \in [N]})$, and run $(\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}) \leftarrow \text{SimCombKeys} \left(\left\{ (\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_i, \widetilde{\text{ek}}_i) \right\}_{i \in T} \right)$. Send $(\text{sid}, \text{ek}, \{(\text{pk}_i, \text{ek}_i)\}_{i \in [N]})$ to \mathcal{A} on behalf of $\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}$.

Phase 3:

- For $j \in \bar{T}$, compute $c_j \leftarrow \text{Enc}_{\text{pk}}(0)$ and $\pi_j^{(x)} \leftarrow \text{Sim}_{\text{tk}^{(x)}}^X(c_j, \text{pk})$. Send $(\text{sid}, j, (c_j, \text{pk}), \pi_j^{(x)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .
- Upon receiving $(\text{sid}, i, (c_i, \text{pk}), \pi_i^{(x)})$ as a message to \mathcal{F}_{BC} for all $i \in T$, verify that $\pi_i^{(x)}$ is valid, and if so, use the extraction trapdoor to extract the witness \tilde{x}_i :

$$\tilde{x}_i := \text{Ext}_{\text{xtk}^{(x)}}^X((c_i, \text{pk}), \pi_i^{(x)})$$

Send $\{\tilde{x}_i\}_{i \in T}$ to \mathcal{F} and receive \tilde{y} .

Phase 4: If server is *not* corrupted, compute $c := \text{Eval}(C, c_1, \dots, c_N)$ and $\varphi \leftarrow \text{Prove}(c, c_1, \dots, c_N)$ honestly, and send $(\text{sid}, 0, \varphi)$ to \mathcal{F}_{BC} . Otherwise, receive $(\text{sid}, 0, c)$ from \mathcal{A} as a message to $\mathcal{F}_{\text{RAND, BC}}^{\mathcal{E}}$, and $(\text{sid}, 0, \varphi)$ as a message to \mathcal{F}_{BC} . In the latter case, verify that φ is valid. In either case, compute $\hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1)$ and send $(\text{sid}, 0, c, \hat{c})$ on behalf of $\mathcal{F}_{\text{RAND, BC}}^{\mathcal{E}}$.

Phase 5:

- For $i \in T$, compute $\tilde{\mu}_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(dec)})$.
- Compute $\{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}_{\text{sk}_i}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T})$, and for $j \in \bar{T}$, send (sid, j, μ_j) on behalf of \mathcal{F}_{BC} .
- Upon receiving (sid, i, μ_i) from \mathcal{A} as a message to \mathcal{F}_{BC} for all $i \in T$, verify that $\tilde{\mu}_i = \mu_i$, and halt if they are not the same.

2. Output \mathcal{A} 's output.

Figure 15: Simulator \mathcal{S}^{MAL} for (static) malicious adversary \mathcal{A} corrupting clients $T \subsetneq [N]$ and possibly server S .

Hybrid 2: We change how \hat{c} is created. We let

$$\tilde{x}_i \leftarrow \text{Dec}_{\text{sk}}(c_i) \quad \text{for } i \in T \quad \text{and} \quad \tilde{x}_j := x_j \quad \text{for } j \in \bar{T}$$

Instead of using the ReRand algorithm, we simply sample a fresh encryption of $\tilde{y} = C(\tilde{x}_1, \dots, \tilde{x}_N)$ with $\text{mode} = 1$:

$$\begin{aligned}
\rho_j^{(gen)} &\leftarrow \text{Com}_{\text{ck}}(r_j^{(gen)} ; s_j^{(gen)}) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(r_j^{(dec)} ; s_j^{(dec)}) \\
&\left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j ; r_j^{(gen)}) \right\}_{j \in \bar{T}} \\
&\left\{ \pi_j^{(gen)} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j), (r_j^{(gen)}, s_j^{(dec)})) \right\}_{j \in \bar{T}} \\
(\text{sk}, \text{ek}) &:= \text{CombineKeys} \left(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\
&\left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \\
&\left\{ \pi_j^{(x)} \leftarrow \text{Prove}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j), (r_j^{(dec)}, s_j^{(dec)}, r_j^{(gen)}, s_j^{(gen)}) \right\}_{j \in \bar{T}}
\end{aligned}$$

We claim that $\text{View}_1 \stackrel{s}{\approx} \text{View}_2$ by *ciphertext re-randomization* of \mathcal{E} . The reduction computes $\{\rho_j^{(gen)}, \rho_j^{(dec)}\}_{j \in \bar{T}}$ honestly and receives $\{\rho_i^{(gen)}, \rho_i^{(dec)}\}_{i \in T}$. It also computes $\{(\text{pk}_j, \text{sk}_j, \text{ek}_j), \pi_j^{(gen)}\}_{j \in \bar{T}}$ honestly, receives $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(gen)}\}_{i \in T}$ from \mathcal{A} , and verifies that all $\pi_i^{(gen)}$ are valid. It computes $(\text{pk}, \text{sk}, \text{ek})$ using CombinePK and CombineKeys and computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(x_j), \pi_j^{(x)}\}_{j \in \bar{T}}$ honestly. It receives $\{c_i, \pi_i^{(x)}\}_{i \in T}$ from \mathcal{A} and verifies that all $\pi_i^{(x)}$ are valid. If the server is honest, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ and φ honestly; otherwise it receives c, φ from \mathcal{A} , in which case it verifies φ . It computes $\{\tilde{x}_i \leftarrow \text{Dec}_{\text{sk}}(c_i)\}_{i \in T}$ and $\tilde{y} = C(\tilde{x}_1, \dots, \tilde{x}_N)$. Given \hat{c} which is either $\hat{c} = \text{ReRand}(c)$ or $\hat{c} = \text{Enc}_{\text{pk}}(y, 1)$, it computes $\{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c}), \pi_j^{(dec)}\}_{j \in \bar{T}}$ honestly.

If all proofs $\pi_i^{(gen)}$ verify, by soundness of $\Pi^{(gen)}$ we know that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple for all $i \in T$. Likewise, by soundness of $\Pi^{(x)}$ we know that for all $i \in T$ there exists $x_i, r_i^{(x)}$ such that $c_i = \text{Enc}_{\text{pk}}(x_i, 0; r_i^{(x)})$, and by construction we have $\tilde{x}_i = x_i$ (except with negligible probability). Therefore, $\text{View}_1 \stackrel{s}{\approx} \text{View}_2$ by *ciphertext re-randomization* of \mathcal{E} .

Hybrid 3: We now change how the \tilde{x}_i are created. Instead of decrypting the ciphertext c_i , we obtain \tilde{x}_i by using the proof of knowledge extractor of $\Pi^{(x)}$.

$$\tilde{x}_i \leftarrow \text{Ext}^{\text{X}}(\pi_i^{(x)}) \quad \text{for } i \in T \quad \text{and} \quad \tilde{x}_j := x_j \quad \text{for } j \in \bar{T}$$

We claim that $\text{View}_2 \stackrel{c}{\approx} \text{View}_3$ by the *simulation-extractability* property of $\Pi^{(x)}$. Note that we must rely on the stronger property of simulation-extractability instead of standard proof of knowledge since the adversary sees the simulated proofs $\pi_j^{(x)}$ *before* producing the proofs $\pi_i^{(x)}$. We therefore need to be able to extract the witness \tilde{x}_i from the proof $\pi_i^{(x)}$ even if the adversary sees simulated proofs before creating $\pi_i^{(x)}$. This is precisely what simulation-extractability guarantees.

Hybrid 4: We now change the proofs $\pi_j^{(dec)}$ so that they are simulated proofs, created with the Sim^{Dec} algorithm, instead of created honestly using $\text{Prove}^{\text{Dec}}$.

$$\begin{aligned}
\rho_j^{(gen)} &\leftarrow \text{Com}_{\text{ck}}(r_j^{(gen)} ; s_j^{(gen)}) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(r_j^{(dec)} ; s_j^{(dec)}) \\
&\left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j ; r_j^{(gen)}) \right\}_{j \in \bar{T}}
\end{aligned}$$

$$\begin{aligned}
& \left\{ \pi_j^{(gen)} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j), (r_j^{(gen)}, s_j^{(dec)})) \right\}_{j \in \bar{T}} \\
& (\cdot, \text{ek}) := \text{CombineKeys} \left(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(dec)} \leftarrow \text{Sim}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j) \right\}_{j \in \bar{T}}
\end{aligned}$$

We have that $\text{View}_3 \stackrel{c}{\approx} \text{View}_4$ by the *zero-knowledge* property of $\Pi^{(dec)}$.

Hybrid 5: We now change the commitment $\rho_j^{(dec)}$ so that it is a commitment of 0, instead of the real randomness used in decryption.

$$\begin{aligned}
& \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(r_j^{(gen)} ; s_j^{(gen)}) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\
& \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j ; r_j^{(gen)}) \right\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(gen)} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j), (r_j^{(gen)}, s_j^{(dec)})) \right\}_{j \in \bar{T}} \\
& (\cdot, \text{ek}) := \text{CombineKeys} \left(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j \leftarrow \text{ShareDec}_{\text{sk}_j}(\hat{c})\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(dec)} \leftarrow \text{Sim}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j) \right\}_{j \in \bar{T}}
\end{aligned}$$

We have that $\text{View}_4 \stackrel{c}{\approx} \text{View}_5$ by the *hiding* property of the trapdoor commitment scheme Com .

Hybrid 6: In this hybrid, for $i \in T$ we let

$$r_i^{(dec)} := \text{Ext}^{\text{Com}}(\rho_i^{(dec)}) \quad \text{and} \quad \tilde{\mu}_i := \text{ShareDec}_{\text{sk}_i}(\hat{c} ; r_i^{(dec)})$$

and halt if $\tilde{\mu}_i \neq \mu_i$, where $\{\mu_i\}_{i \in T}$ are the shares sent by \mathcal{A} . We claim that $\text{View}_5 \stackrel{c}{\approx} \text{View}_6$. This is guaranteed by the *soundness* of $\Pi^{(dec)}$ and the *extraction* property of the trapdoor commitment.

Hybrid 7: We change the way the decryption shares $\{\mu_j\}_{j \in \bar{T}}$ are generated. Instead of using the ShareDec algorithm with secret key sk_j , we now use the SimShareDec algorithm with $\tilde{y} = C(\tilde{x}_1, \dots, \tilde{x}_N)$ and $\{\tilde{\mu}_i\}_{i \in T}$:

$$\begin{aligned}
& \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(r_j^{(gen)} ; s_j^{(gen)}) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\
& \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j ; r_j^{(gen)}) \right\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(gen)} \leftarrow \text{Prove}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j), (r_j^{(gen)}, s_j^{(dec)})) \right\}_{j \in \bar{T}} \\
& (\cdot, \text{ek}) := \text{CombineKeys} \left(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\
& \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T})
\end{aligned}$$

$$\left\{ \pi_j^{(dec)} \leftarrow \text{Sim}^{\text{Dec}}((\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j)) \right\}_{j \in \bar{T}}$$

We claim that $\text{View}_6 \stackrel{c}{\approx} \text{View}_7$ by *share-simulation indistinguishability* of \mathcal{E} . The reduction computes $\{\rho_j^{(gen)}, \rho_j^{(dec)}\}_{j \in \bar{T}}$ and receives $\{\rho_i^{(gen)}, \rho_i^{(dec)}\}_{i \in T}$. It also computes $\{(\text{pk}_j, \text{sk}_j, \text{ek}_j), \pi_j^{(gen)}\}_{j \in \bar{T}}$ honestly, receives $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(gen)}\}_{i \in T}$ from \mathcal{A} , and verifies that all $\pi_i^{(gen)}$ are valid. It computes $(\text{pk}, \text{sk}, \text{ek})$ using `CombinePK` and `CombineKeys` and computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(x_j), \pi_j^{(x)}\}_{j \in \bar{T}}$ with a simulated proof. It receives $\{c_i, \pi_i^{(x)}\}_{i \in T}$ from \mathcal{A} and verifies that all $\pi_i^{(x)}$ are valid. If the server is honest, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ and φ honestly; otherwise it receives c, φ from \mathcal{A} , in which case it verifies φ . It computes $\{\tilde{x}_i \leftarrow \text{Ext}^X(\pi_i^{(x)})\}_{i \in T}$ and $\tilde{y} = C(\tilde{x}_1, \dots, \tilde{x}_N)$ and $\hat{c} \leftarrow \text{Enc}(\tilde{y}, 1)$. It then computes $r_i^{(dec)} \leftarrow \text{Ext}^{\text{Com}}(\rho_i^{(dec)})$ and $\tilde{\mu}_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(dec)})$, $\{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T})$ and simulated proofs $\{\pi_j^{(dec)}\}_{j \in \bar{T}}$. It receives $\{\mu_i, \pi_i^{(dec)}\}_{i \in T}$ from \mathcal{A} , verifies all proofs $\pi_i^{(dec)}$, and halts if $\tilde{\mu}_i \neq \mu_i$ for any $i \in T$.

If all proofs $\pi_i^{(gen)}$ verify, by soundness of $\Pi^{(gen)}$ we know that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple for all $i \in T$. Likewise, by *simulation-soundness* of $\Pi^{(dec)}$ we know that for all $i \in T$ there exists $r_i^{(dec)}$ such that $\mu_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(dec)})$. By construction, $\hat{c} \leftarrow \text{Enc}(\tilde{y})$. Thus, $\text{View}_5 \stackrel{c}{\approx} \text{View}_6$ by *share-simulation indistinguishability* of \mathcal{E} .

Hybrid 8: We now change the proofs $\pi_j^{(gen)}$ so that they are simulated proofs, created with the Sim^{Gen} algorithm, instead of created honestly using $\text{Prove}^{\text{Gen}}$.

$$\begin{aligned} & \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(r_j^{(gen)}; s_j^{(gen)}) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\ & \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j; r_j^{(gen)}) \right\}_{j \in \bar{T}} \quad , \quad \left\{ \pi_j^{(gen)} \leftarrow \text{Sim}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j)) \right\}_{j \in \bar{T}} \\ & (\text{pk}, \cdot, \text{ek}) := \text{CombineKeys} \left(\{(\text{pk}_i, \text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^X((c_j, \text{pk})) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\tilde{\mu}_i\}_{i \in T}) \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{Dec}}((\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j)) \right\}_{j \in \bar{T}} \end{aligned}$$

We have that $\text{View}_7 \stackrel{c}{\approx} \text{View}_8$ by the *zero-knowledge* property of $\Pi^{(gen)}$.

Hybrid 9: We change the commitment $\rho_j^{(gen)}$ so that it is a commitment of 0, instead of the real randomness used in key generation.

$$\begin{aligned} & \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(0) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\ & \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j; r_j^{(gen)}) \right\}_{j \in \bar{T}} \quad , \quad \left\{ \pi_j^{(gen)} \leftarrow \text{Sim}^{\text{Gen}}((\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j)) \right\}_{j \in \bar{T}} \\ & (\cdot, \text{ek}) := \text{CombineKeys} \left(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]} \right) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}} \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^X((c_j, \text{pk})) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\tilde{\mu}_i\}_{i \in T}) \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{Dec}}((\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j)) \right\}_{j \in \bar{T}} \end{aligned}$$

We have that $\text{View}_8 \stackrel{c}{\approx} \text{View}_9$ by the *hiding* property of the trapdoor commitment scheme Com .

Hybrid 10: In this hybrid, for $i \in T$ we let

$$r_i^{(gen)} := \text{Ext}^{\text{Com}}(\rho_i^{(gen)}) \quad \text{and} \quad (\widetilde{\text{pk}}_i, \cdot, \cdot) := \text{Keygen}(1^\kappa, i; r_i^{(gen)})$$

and halt if $\widetilde{\text{pk}}_i \neq \text{pk}_i$, where $\{(\text{pk}_i, \cdot, \cdot)\}_{i \in T}$ are the public keys sent by \mathcal{A} . We claim that $\text{View}_9 \stackrel{c}{\approx} \text{View}_{10}$. This is guaranteed by the *simulation-soundness* of $\Pi^{(gen)}$ and the *extraction* property of the trapdoor commitment.

Hybrid 11: This is actually a series of $N - t$ hybrids. We change how the ciphertexts $\{c_j\}_{j \in \bar{T}}$ are created, one at a time. Instead of having c_{j^*} be an encryption of input x_{j^*} , we have it be an encryption of 0.

$$\begin{aligned} & \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(0) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\ & \left\{ (\text{pk}_j, \cdot, \text{ek}_j) \leftarrow \text{Keygen}(1^\kappa, j; r_j^{(gen)}) \right\}_{j \in \bar{T}} \quad , \quad \left\{ \pi_j^{(gen)} \leftarrow \text{Sim}^{\text{Gen}}(\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j) \right\}_{j \in \bar{T}} \\ & (\cdot, \text{ek}) := \text{CombineKeys}(\{(\text{sk}_i, \text{ek}_i)\}_{i \in [N]}) \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}, j \leq j^*} \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(x_j)\}_{j \in \bar{T}, j > j^*} \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T}) \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(gen)}, \rho_j^{(dec)}, \text{ck}, j) \right\}_{j \in \bar{T}} \end{aligned}$$

We claim two consecutive hybrids are computationally indistinguishable by *semantic security* of \mathcal{E} . The reduction computes $\{\rho_j^{(gen)}, \rho_j^{(dec)}\}_{j \in \bar{T}}$ and receives $\{\rho_i^{(gen)}, \rho_i^{(dec)}\}_{i \in T}$. It computes $r_i^{(gen)} := \text{Ext}^{\text{Com}}(\rho_i^{(gen)})$ and sends $\{(\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_i, \widetilde{\text{ek}}_i) := \text{Keygen}(1^\kappa, i; r_i^{(gen)})\}_{i \in T}$ to the challenger. It forwards $(\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}})$ from the challenger to \mathcal{A} and sends to \mathcal{A} simulated proofs $\pi_j^{(gen)}$. Upon receiving $(\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(gen)}$ from \mathcal{A} , it verifies $\pi_i^{(gen)}$ and checks that $\widetilde{\text{pk}}_i = \text{pk}_i$ and halts otherwise. It uses the challenge ciphertext as c_{j^*} and computes the other ciphertexts c_j accordingly, with simulated proofs $\pi_i^{(x)}$ for all $j \in \bar{T}$. When it receives $c_i, \pi_i^{(x)}$, it verifies $\pi_i^{(x)}$. If the server is honest, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ and φ honestly; otherwise it receives c, φ from \mathcal{A} , in which case it verifies φ . It computes $r_i^{(dec)} \leftarrow \text{Ext}^{\text{Com}}(\rho_i^{(dec)})$ and $\tilde{\mu}_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(dec)})$, $\{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T})$ and simulated proofs $\{\pi_j^{(dec)}\}_{j \in \bar{T}}$. It receives $\{\mu_i, \pi_i^{(dec)}\}_{i \in T}$ from \mathcal{A} , verifies all proofs $\pi_i^{(dec)}$, and halts if $\tilde{\mu}_i \neq \mu_i$ for any $i \in T$. We define View_{11} to be the view having all c_j encrypting 0. Putting these $N - t$ hybrids together, we conclude that $\text{View}_{10} \stackrel{c}{\approx} \text{View}_{11}$.

Hybrid 12: We change how the public/evaluation key pairs $\{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}$ and the combined evaluation key ek are generated. Instead of creating them using Keygen and CombineKeys , we use the SimCombKeys algorithm.

$$\begin{aligned} & \rho_j^{(gen)} \leftarrow \text{Com}_{\text{ck}}(0) \quad , \quad \rho_j^{(dec)} \leftarrow \text{Com}_{\text{ck}}(0) \\ & (\text{ek}, \{(\text{pk}_j, \text{ek}_j)\}_{j \in \bar{T}}) \leftarrow \text{SimCombKeys}(\{\text{pk}_i, \text{ek}_i\}_{i \in T}) \\ & \left\{ \pi_j^{(gen)} \leftarrow \text{Sim}^{\text{Gen}}(\text{pk}_j, \text{ek}_j, \text{ck}, \rho_j^{(gen)}, j) \right\}_{j \in \bar{T}} \quad , \quad \{c_j = \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}} \\ & \left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{X}}(c_j, \text{pk}) \right\}_{j \in \bar{T}} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(\tilde{y}, 1) \quad , \quad \{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, y, \{\tilde{\mu}_i\}_{i \in T}) \end{aligned}$$

$$\left\{ \pi_j^{(x)} \leftarrow \text{Sim}^{\text{Dec}}(\hat{c}, \mu_j, \text{pk}_j, \rho_j^{(\text{gen})}, \rho_j^{(\text{dec})}, \text{ck}, j) \right\}_{j \in \bar{T}}$$

We claim that $\text{View}_{11} \stackrel{c}{\approx} \text{View}_{12}$ by *key-simulation indistinguishability* of \mathcal{E} . The reduction computes $\{\rho_j^{(\text{gen})}, \rho_j^{(\text{dec})}\}_{j \in \bar{T}}$ honestly and receives $\{\rho_i^{(\text{gen})}, \rho_i^{(\text{dec})}\}_{i \in T}$. It computes $r_i^{(\text{gen})} := \text{Ext}^{\text{Com}}(\rho_i^{(\text{gen})})$ and $\{(\widetilde{\text{pk}}_i, \widetilde{\text{sk}}_i, \widetilde{\text{ek}}_i) := \text{Keygen}(1^\kappa, i; r_i^{(\text{gen})})\}_{i \in T}$. Given keys $(\text{ek}, \{(\widetilde{\text{pk}}_j, \widetilde{\text{ek}}_j)\}_{j \in \bar{T}})$ which are either created honestly using `Keygen` and `CombineKeys`, or simulated using `SimCombKeys`, it forwards these to \mathcal{A} . It then receives $\{(\text{pk}_i, \text{sk}_i, \text{ek}_i), \pi_i^{(\text{gen})}\}_{i \in T}$ from \mathcal{A} , and verifies $\pi_i^{(\text{gen})}$. It halts if $\widetilde{\text{pk}}_i \neq \text{pk}_i$ for any $i \in T$; otherwise it computes $\text{pk} := \text{CombinePK}(\{\text{pk}_i\}_{i \in [N]})$. It also computes $\{c_j \leftarrow \text{Enc}_{\text{pk}}(0)\}_{j \in \bar{T}}$ with simulated proofs $\pi_j^{(x)}$, receives $\{(c_i, \pi_i^{(x)})\}_{i \in T}$ from \mathcal{A} and verifies $\pi_i^{(x)}$. If the server is honest, it computes $c := \text{Eval}_{\text{ek}}(C, c_1, \dots, c_N)$ and φ honestly; otherwise it receives c, φ from \mathcal{A} , in which case it verifies φ . It computes $\{\tilde{x}_i \leftarrow \text{Ext}^X(\pi_i^{(x)})\}_{i \in T}$ and $\tilde{y} = C(\tilde{x}_1, \dots, \tilde{x}_N)$ and $\hat{c} \leftarrow \text{Enc}(\tilde{y}, 1)$. It then computes $r_i^{(\text{dec})} \leftarrow \text{Ext}^{\text{Com}}(\rho_i^{(\text{dec})})$ and $\tilde{\mu}_i := \text{ShareDec}_{\text{sk}_i}(\hat{c}; r_i^{(\text{dec})})$, $\{\mu_j\}_{j \in \bar{T}} \leftarrow \text{SimShareDec}(\hat{c}, \tilde{y}, \{\tilde{\mu}_i\}_{i \in T})$ and simulated proofs $\{\pi_j^{(\text{dec})}\}_{j \in \bar{T}}$. It receives $\{\mu_i, \pi_i^{(\text{dec})}\}_{i \in T}$ from \mathcal{A} , verifies all proofs $\pi_i^{(\text{dec})}$, and halts if $\tilde{\mu}_i \neq \mu_i$ for any $i \in T$.

If all proofs $\pi_i^{(\text{gen})}$ verify, by *simulation-soundness* of $\Pi^{(\text{gen})}$ we know that $(\text{pk}_i, \text{sk}_i, \text{ek}_i)$ is a valid key tuple for all $i \in T$. Thus, $\text{View}_{11} \stackrel{c}{\approx} \text{View}_{12}$ by *key-simulation indistinguishability* of \mathcal{E} .

We have proved that $\text{View}_0 \stackrel{c}{\approx} \text{View}_{12}$. Since View_{12} is the view that \mathcal{S}^{MAL} produces, and View_0 is the view produced in an execution of Π^{MAL} , we have that $\text{IDEAL}_{\mathcal{F}, \mathcal{S}^{\text{MAL}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi^{\text{MAL}}, \mathcal{A}}(\vec{x})$. \square

A.2 Ciphertext Re-Randomization

For security, we are not concerned with *privacy* as the parties do not have any inputs that they wish to keep private. However, we are concerned with *correctness*; we need to ensure that the output \hat{c} of the protocol, is indistinguishable from a fresh encryption of the plaintext y . This is fairly straight-forward in the case of a semi-honest adversary but requires more care when dealing with a malicious adversary. In this case, we also rely on simulation-sound NIZK proofs. As before, we show the proof of security for protocol $\Pi_{\text{RAND}}^{\text{MAL}}$ after the functionality $\mathcal{F}_{\text{ZK.BC}}^{R(\text{gen})}$ has been instantiated with simulation-sound NIZK proofs.

Proof of Theorem 7.1: We describe the simulator $\mathcal{S}_{\text{RAND}}^{\text{SH}}$ for a (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ (and possibly server S) in Figure 17. We prove that $\text{IDEAL}_{\Pi_{\text{RAND}}^{\text{SH}}, \mathcal{S}_{\text{RAND}}^{\text{SH}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}, \mathcal{A}}(\vec{x})$ using a series of hybrids. The view of \mathcal{A} consists of

$$\text{View} = (\text{pk}, c, z_1, \dots, z_N, \hat{c})$$

We simply prove that the joint distribution of these variables in the simulation (where the ciphertext output is a fresh encryption of y) is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol $\Pi_{\text{RAND}}^{\text{SH}}$. This is easy to see because z_{j^*} is a deterministic function of \hat{c} and $\{z_j\}_{i \in N, i \neq j^*}$, all z_i 's for $i \neq j^*$ are distributed the same in both cases (since the adversary is semi-honest), and the distribution of \hat{c} is statistically close to $\text{Enc}_{\text{pk}}(y, 1)$ since the noise in each z_i is a super-polynomial factor larger than the noise in c . \square

Proof of Theorem 7.2: We describe the simulator $\mathcal{S}_{\text{RAND}}^{\text{MAL}}$ for a (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ (and possibly server S) in Figure 17. We prove that $\text{IDEAL}_{\Pi_{\text{RAND}}^{\text{MAL}}, \mathcal{S}_{\text{RAND}}^{\text{MAL}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi_{\text{RAND}}^{\text{MAL}}, \mathcal{A}}(\vec{x})$.

Simulator $\mathcal{S}_{\text{RAND}}^{\text{SH}}$

Inputs: All parties P_1, \dots, P_N and server S get public key pk as input. Server S additionally gets a ciphertext $c \in \mathcal{C}_{\text{pk}}$. The simulator also receives input y , which is the plaintext of c .

- Upon receiving (sid, i, z_i) as a message to \mathcal{F}_{BC} for all $i \in T$:
 - Let $\hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1)$, and let $j^* \in \bar{T}$. For all $j \in \bar{T}, j \neq j^*$, set $z_j := \text{Enc}_{\text{pk}}(0, 1)$. Set $z_{j^*} := \hat{c} - \sum_{i \in [N], i \neq j^*} z_i$
 - For all $j \in \bar{T}$, send (sid, j, z_j) to \mathcal{A} on behalf of \mathcal{F}_{BC} .
- Output whatever \mathcal{A} outputs.

Figure 16: Simulator $\mathcal{S}_{\text{RAND}}^{\text{SH}}$ for (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ in protocol $\Pi_{\text{RAND}}^{\text{SH}}$

$\text{REAL}_{\mathcal{F}_{\text{RAND.BC}}^{\mathcal{E}}, \mathcal{A}}(\vec{x})$ using a series of hybrids. The view of the \mathcal{A} consists of

$$\text{View} = \left(\text{pk}, \text{ck}, c, \left\{ \rho_i^{(0)}, z_i, \pi_i^{(0)} \right\}_{i \in [N]}, \hat{c} \right)$$

Simulator $\mathcal{S}_{\text{RAND}}^{\text{MAL}}$

Inputs: All parties P_1, \dots, P_N and server S get public key pk , and a commitment key ck as input. Server S additionally gets a ciphertext $c \in \mathcal{C}_{\text{pk}}$. The simulator also receives input y , which is the plaintext of c .

Round 1: For $j \in \bar{T}$, compute commitments $\rho_j^{(0)} \leftarrow \text{Com}_{\text{ck}}(0)$ and send $(\text{sid}, j, \rho_j^{(0)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Round 2: Upon receiving $(\text{sid}, i, \rho_i^{(0)})$ as a message to \mathcal{F}_{BC} for all $i \in T$

- Let

$$r_i^{(0)} := \text{Ext}^{\text{Com}}(\rho_i^{(0)}) \quad \text{and} \quad \tilde{z}_i := \text{Enc}_{\text{pk}}(0, 1; r_i^{(0)})$$
- Compute $\hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1)$, and let $j^* \in \bar{T}$. For all $j \in \bar{T}, j \neq j^*$, set $\tilde{z}_j := \text{Enc}_{\text{pk}}(0, 1)$. Set $\tilde{z}_{j^*} := \hat{c} - \sum_{i \in [N], i \neq j^*} \tilde{z}_i$.
- For $j \in \bar{T}$, compute $\pi_j^{(0)} \leftarrow \text{Sim}^{\text{Zero}}(\rho_j^{(0)}, z_j, \text{pk})$.
- For $j \in \bar{T}$, send $(\text{sid}, j, \tilde{z}_j, \pi_j^{(0)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC}

Round 3: Upon receiving $(\text{sid}, i, z_i, \pi_i^{(0)})$ as a message to \mathcal{F}_{BC} for all $i \in T$, verify $\pi_i^{(0)}$ and check that $\tilde{z}_i = z_i$; halt if they are not equal. If the server S is honest, send $(\text{sid}, 0, c, \hat{c})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Output whatever \mathcal{A} outputs.

Figure 17: Simulator $\mathcal{S}_{\text{RAND}}^{\text{MAL}}$ for (static) malicious adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ in protocol $\Pi_{\text{RAND}}^{\text{MAL}}$

We simply prove that the joint distribution of these variables in the simulation is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol $\Pi_{\text{RAND}}^{\text{MAL}}$. We let View_i be the joint distribution of these variables in hybrid i .

Hybrid 0: This is the real-world execution of $\Pi_{\text{RAND}}^{\text{MAL}}$. We have

$$\left\{ \rho_j^{(0)} := \text{Com}_{\text{ck}}(r_j^{(0)}; s_j^{(0)}) \right\}_{j \in \bar{T}}, \quad \left\{ z_j \leftarrow \text{Enc}(0, 1; r_j^{(0)}) \right\}_{j \in \bar{T}}$$

$$\left\{ \pi_j^{(0)} \leftarrow \text{Prove}^{\text{Zero}}((\rho_j^{(0)}, z_j, \text{pk}), (r_j^{(0)}, s_j^{(0)})) \right\} \quad , \quad \hat{c} := c + \sum_{i=1}^N z_i$$

Hybrid 1: We now change the proofs $\pi_j^{(0)}$ so that they are simulated proofs, created with the Sim^{Zero} algorithm, instead of created honestly using $\text{Prove}^{\text{Zero}}$.

$$\left\{ \rho_j^{(0)} := \text{Com}_{\text{ck}}(r_j^{(0)} ; s_j^{(0)}) \right\}_{j \in \bar{T}} \quad , \quad \left\{ z_j \leftarrow \text{Enc}(0, 1 ; r_j^{(0)}) \right\}_{j \in \bar{T}}$$

$$\left\{ \pi_j^{(0)} \leftarrow \text{Sim}^{\text{Zero}}(\rho_j^{(0)}, z_j, \text{pk}) \right\} \quad , \quad \hat{c} := c + \sum_{i=1}^N z_i$$

$\text{View}_0 \stackrel{c}{\approx} \text{View}_1$ by the *zero-knowledge* property of $\Pi^{(\text{zero})}$.

Hybrid 2: We change the commitments $\rho_j^{(0)}$ so that they are now commitments of 0 instead of the randomness $r_j^{(0)}$.

$$\left\{ \rho_j^{(0)} := \text{Com}_{\text{ck}}(0) \right\}_{j \in \bar{T}} \quad , \quad \left\{ z_j \leftarrow \text{Enc}(0, 1 ; r_j^{(0)}) \right\}_{j \in \bar{T}}$$

$$\left\{ \pi_j^{(0)} \leftarrow \text{Sim}^{\text{Zero}}(\rho_j^{(0)}, z_j, \text{pk}) \right\} \quad , \quad \hat{c} := c + \sum_{i=1}^N z_i$$

$\text{View}_1 \stackrel{c}{\approx} \text{View}_2$ by the *hiding* property of the trapdoor commitment scheme.

Hybrid 3: In this hybrid, for $i \in T$ we let

$$r_i^{(0)} := \text{Ext}^{\text{Com}}(\rho_i^{(0)}) \quad \text{and} \quad \tilde{z}_i := \text{Enc}_{\text{pk}}(0, 1 ; r_i^{(0)})$$

and halt if $\tilde{z}_i \neq z_i$, where $\{z_i\}_{i \in T}$ are the ciphertexts sent by \mathcal{A} . We claim that $\text{View}_2 \stackrel{c}{\approx} \text{View}_3$. This is guaranteed by the *simulation-soundness* of $\Pi^{(\text{zero})}$ and the *extraction* property of the trapdoor commitment.

Hybrid 4: Let $j^* \in \bar{T}$, and for $j \in \bar{T}$, let $\tilde{z}_j = z_j$. We view \tilde{z}_{j^*} as a deterministic function of \hat{c} and the rest of the \tilde{z}_i 's: $z_{j^*} = \hat{c} - \sum_{i=1}^N \tilde{z}_i$. We now change \hat{c} to be a fresh encryption of y .

$$\left\{ \rho_j^{(0)} := \text{Com}_{\text{ck}}(0) \right\}_{j \in \bar{T}} \quad , \quad \left\{ z_j \leftarrow \text{Enc}(0, 1 ; r_j^{(0)}) \right\}_{j \in \bar{T}, j \neq j^*} \quad , \quad z_{j^*} = \hat{c} - \sum_{i=1}^N \tilde{z}_i$$

$$\left\{ \pi_j^{(0)} \leftarrow \text{Sim}^{\text{Zero}}(\rho_j^{(0)}, z_j, \text{pk}) \right\} \quad , \quad \hat{c} \leftarrow \text{Enc}_{\text{pk}}(y, 1)$$

We claim that $\text{View}_3 \stackrel{s}{\approx} \text{View}_4$. In both cases, z_{j^*} is the same deterministic function of \hat{c} and the rest of the \tilde{z}_j 's. The distribution of \hat{c} is statistically close to $\text{Enc}_{\text{pk}}(y, 1)$ since the noise in each z_i is a super-polynomial factor larger than the noise in c .

We have proved that $\text{View}_0 \stackrel{c}{\approx} \text{View}_4$. Since View_4 is the view that $\mathcal{S}_{\text{RAND}}^{\text{MAL}}$ produces, and View_0 is the view produced in an execution of $\Pi_{\text{RAND}}^{\text{MAL}}$, we have that $\text{IDEAL}_{\mathcal{F}, \mathcal{S}_{\text{RAND}}^{\text{MAL}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi_{\text{RAND}}^{\text{MAL}}, \mathcal{A}}(\vec{x})$. \square

A.3 Key Combination

Proof of Theorem 7.3: We describe the simulator $\mathcal{S}_{\text{KEYS}}^{\text{SH}}$ for a (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ in Figure 18. We prove that $\text{IDEAL}_{\Pi_{\text{KEYS}}^{\text{SH}}, \mathcal{S}_{\text{KEYS}}^{\text{SH}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{A}}(\vec{x})$ using a series of hybrids. The view of the \mathcal{A} consists of

$$\text{View} = \left(\{\text{sk}_i\}_{i \in T}, \{\text{pk}_i, \text{ek}_i\}_{i \in [N]}, \left\{ w_{d,\tau}^{(j,k)}, \tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]} \right)$$

Simulator $\mathcal{S}_{\text{KEYS}}^{\text{SH}}$

Round 1:

- Runs $\text{params} = \left(\mathbf{a}, \left\{ a_{d,\tau}^{(i)}, a'_{d,\tau}^{(i)} \right\}_{d,\tau,i} \right) \leftarrow \text{Setup}(1^\kappa, N)$. When a corrupted party P_i calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with sid , sends $(\text{sid}, \text{params})$.
- For $k \in [N]$,
 - Parses $\text{ek}_k := \left\{ \tilde{\mathbf{c}}_{0,d}^{(k)}, \tilde{\mathbf{c}}_{1,d}^{(k)}, \tilde{\mathbf{z}}_{0,d}^{(k)}, \tilde{\mathbf{z}}_{1,d}^{(k)} \right\}_{d \in [D]}$.
 - Verifies that $\tilde{\mathbf{c}}_{1,d}^{(k)} = (a_{d,1}^{(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{(k)})$ and $\tilde{\mathbf{z}}_{1,d}^{(k)} = (a'_{d,1}^{(k)}, \dots, a'_{d, \lfloor \log q_{d-1} \rfloor}^{(k)})$
 - For all $j \in \bar{T}$ and $k \in [N]$, samples $w_{d,\tau}^{(j,k)} \leftarrow R_{q_d}$ and sends $(\text{sid}, j, w_{d,\tau}^{(j,k)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Round 2: Upon receiving $(\text{sid}, i, w_{d,\tau}^{(i,k)})$ from \mathcal{F}_{BC} for all $i \in T, k \in [N], d \in [D], \tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$:

- For all $j \in \bar{T}$ and $k \in [N]$, samples $\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{q_d}^2$ and sends $(\text{sid}, j, \tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Output whatever \mathcal{A} outputs.

Figure 18: Simulator $\mathcal{S}_{\text{KEYS}}^{\text{SH}}$ for (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ in protocol $\Pi_{\text{KEYS}}^{\text{SH}}$

We simply prove that the joint distribution of these variables in the simulation is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol $\Pi_{\text{KEYS}}^{\text{SH}}$. We further prove that the distribution of ek as output by the simulator is computationally indistinguishable from the output of the algorithm `CombineKeys`. We let View_i be the joint distribution of these variables in hybrid i .

Hybrid 0: This is a real-world execution of $\Pi_{\text{KEYS}}^{\text{SH}}$. We have

$$\left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow \tilde{\mathbf{y}}_{d,\tau}^{(j,k)} + \text{Enc}_{\text{pk}}(0) \right\}_{j \in \bar{T}, k \in [N]}$$

Hybrid 1: We change how we compute $\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}$. We now choose each pair uniformly at random from $R_{q_d}^2$.

$$\left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]}$$

$\text{View}_0 \stackrel{c}{\approx} \text{View}_1$ because the ciphertext $\text{Enc}_{\text{pk}}(0)$ added to $\tilde{\mathbf{y}}_{d,\tau}^{(j,k)}$ in Hybrid 0 is *pseudorandom*, and therefore $\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}$ in Hybrid 0 is pseudorandom as well.

Hybrid 2: We change how we compute $\tilde{\mathbf{c}}_{d,\tau}^{(k)}$. We now choose each pair uniformly at random from R_{qd}^2 .

$$\left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow R_{qd}^2 \right\}, \quad \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]}$$

$\text{View}_1 \stackrel{c}{\approx} \text{View}_2$ because the ciphertext $\text{Enc}_{\text{pk}}(0)$ added to $\tilde{\mathbf{x}}_{d,\tau}^{(k)}$ in Hybrid 0 is *pseudorandom*, and therefore $\tilde{\mathbf{c}}_{d,\tau}^{(k)}$ in Hybrid 0 is pseudorandom as well.

Hybrid 3: We change how we compute $w_{d,\tau}^{(j,k)}$. We now choose it uniformly at random from R_{qd} .

$$\left\{ w_{d,\tau}^{(j,k)} \leftarrow R_{qd} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow R_{qd}^2 \right\}, \quad \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]}$$

The fact that $\text{View}_2 \stackrel{c}{\approx} \text{View}_3$ follows from a series of hybrid arguments, one per each $s_d^{(j)}$, in which we use the fact that $w_{d,\tau}^{(j,k)}$ is a PLWE sample with secret $s_d^{(j)}$ and is therefore pseudorandom.

We have proved that $\text{View}_0 \stackrel{c}{\approx} \text{View}_3$. Since View_3 is the view that $\mathcal{S}_{\text{KEYS}}^{\text{SH}}$ produces, and View_0 is the view produced in an execution of $\Pi_{\text{KEYS}}^{\text{SH}}$, we have that $\text{IDEAL}_{\mathcal{F}, \mathcal{S}_{\text{KEYS}}^{\text{SH}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi_{\text{KEYS}}^{\text{SH}}, \mathcal{A}}(\vec{x})$. Furthermore, the distribution of ek as output by the simulator is computationally indistinguishable from the output of CombineKeys . This again follows from a series of hybrid arguments, all based on the PLWE assumption with secret s_d . \square

Proof of Theorem 7.4: We describe the simulator $\mathcal{S}_{\text{KEYS}}^{\text{MAL}}$ for a (static) semi-honest adversary \mathcal{A} corrupting parties in $T \subseteq [N]$ in Figure 19. We prove that $\text{IDEAL}_{\Pi_{\text{KEYS}}^{\text{MAL}}, \mathcal{S}_{\text{KEYS}}^{\text{MAL}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\mathcal{F}_{\text{KEYS}}^{\mathcal{E}}, \mathcal{A}}(\vec{x})$ using a series of hybrids. The view of the \mathcal{A} consists of

$$\text{View} = \left(\left\{ \text{sk}_i \right\}_{i \in T}, \left\{ \text{pk}_i, \text{ek}_i \right\}_{i \in [N]}, \left\{ w_{d,\tau}^{(j,k)}, \pi_{d,\tau}^{(j,k)}, \tilde{\mathbf{c}}_{d,\tau}^{(k)}, \phi_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \varphi_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]} \right)$$

We simply prove that the joint distribution of these variables in the simulation is computationally indistinguishable from the joint distribution of the same variables in a real-world execution of protocol $\Pi_{\text{KEYS}}^{\text{SH}}$. We further prove that the distribution of ek as output by the simulator is computationally indistinguishable from the output of the algorithm CombineKeys . We let View_i be the joint distribution of these variables in hybrid i .

Hybrid 0: This is a real-world execution of $\Pi_{\text{KEYS}}^{\text{MAL}}$. We have

$$\begin{aligned} & \left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Prove}^{\text{Enc}}((w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}), (s_d^{(j)}, e_{d,\tau}^{(j,\tau)})) \right\}_{j \in \bar{T}, k \in [N]} \\ & \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Prove}^{\text{Rand}}((\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}), (1, r_{d,\tau}^{(k)})) \right\}_{j \in \bar{T}, k \in [N]} \\ & \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow \tilde{\mathbf{y}}_{d,\tau}^{(j,k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(j,k)}) \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Prove}^{\text{Rand}}((\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}), (s_d^{(j)}, r_{d,\tau}^{(j,k)})) \right\}_{j \in \bar{T}, k \in [N]} \end{aligned}$$

Simulator $\mathcal{S}_{\text{KEYS}}^{\text{MAL}}$

Round 1:

- Runs $\text{params} = \left(\mathbf{a}, \left\{ a_{d,\tau}^{(i)}, a'_{d,\tau}{}^{(i)} \right\}_{d,\tau,i} \right) \leftarrow \text{Setup}(1^\lambda, N)$. When a corrupted party P_i calls $\mathcal{F}_{\text{CRS}}^{\mathcal{D}}$ with sid , sends $(\text{sid}, \text{params})$.
- For $k \in [N]$,
 - Parses $\text{ek}_k := \left\{ \tilde{\mathbf{c}}_{0,d}^{(k)}, \tilde{\mathbf{c}}_{1,d}^{(k)}, \tilde{\mathbf{z}}_{0,d}^{(k)}, \tilde{\mathbf{z}}_{1,d}^{(k)} \right\}_{d \in [D]}$.
 - Verifies that $\tilde{\mathbf{c}}_{1,d}^{(k)} = (a_{d,1}^{(k)}, \dots, a_{d, \lfloor \log q_{d-1} \rfloor}^{(k)})$ and $\tilde{\mathbf{z}}_{1,d}^{(k)} = (a'_{d,1}{}^{(k)}, \dots, a'_{d, \lfloor \log q_{d-1} \rfloor}{}^{(k)})$
 - For all $j \in \bar{T}$ and $k \in [N]$, samples $w_{d,\tau}^{(j,k)} \leftarrow R_{q_d}$ and computes $\pi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Enc}}(w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(k)})$. Sends $(\text{sid}, j, w_{d,\tau}^{(j,k)}, \pi_{d,\tau}^{(j,k)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Round 2: Upon receiving $(\text{sid}, i, w_{d,\tau}^{(i,k)})$ from \mathcal{F}_{BC} for all $i \in T, k \in [N], d \in [D], \tau \in \{0, \dots, \lfloor \log q_{d-1} \rfloor\}$, for all $j \in \bar{T}$ and $k \in [N]$:

- Samples $\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{q_d}^2$
- Computes $\phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)})$ and $\varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(i,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)})$.
- Sends $(\text{sid}, j, \tilde{\mathbf{c}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \phi_{d,\tau}^{(k)}, \varphi_{d,\tau}^{(j,k)})$ to \mathcal{A} on behalf of \mathcal{F}_{BC} .

Output whatever \mathcal{A} outputs.

Figure 19: Simulator $\mathcal{S}_{\text{KEYS}}^{\text{MAL}}$ for (static) malicious adversary \mathcal{A} corrupting parties in $T \subsetneq [N]$ in protocol $\Pi_{\text{KEYS}}^{\text{MAL}}$

Hybrid 1: We change how the proofs $\phi_{d,\tau}^{(k)}, \varphi_{d,\tau}^{(j,k)}$ are created. Instead of computing them honestly using the $\text{Prove}^{\text{Rand}}$ algorithm, we compute simulated proofs using the Sim^{Rand} algorithm.

$$\left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Prove}^{\text{Enc}}((w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}), (s_d^{(j)}, e_{d,\tau}^{(j,\tau)})) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow \tilde{\mathbf{y}}_{d,\tau}^{(j,k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(j,k)}) \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}$$

$\text{View}_0 \stackrel{c}{\approx} \text{View}_1$ by the *zero-knowledge* property of Π^{rand} .

Hybrid 2: We change how we compute $\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}$. We now choose each pair uniformly at random from $R_{q_d}^2$.

$$\left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Prove}^{\text{Enc}}((w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}), (s_d^{(j)}, e_{d,\tau}^{(j,\tau)})) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow \tilde{\mathbf{x}}_{d,\tau}^{(k)} + \text{Enc}_{\text{pk}}(0; r_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}$$

$$\left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{q_d}^2 \right\}_{j \in \bar{T}, k \in [N]}, \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]}$$

$\text{View}_1 \stackrel{c}{\approx} \text{View}_2$ because the ciphertext $\text{Enc}_{\text{pk}}(0)$ added to $\tilde{\mathbf{y}}_{d,\tau}^{(j,k)}$ in Hybrid 1 is *pseudorandom*, and therefore $\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}$ in Hybrid 1 is pseudorandom as well.

Hybrid 3: We change how we compute $\tilde{\mathbf{c}}_{d,\tau}^{(k)}$. We now choose each pair uniformly at random from R_{qd}^2 .

$$\begin{aligned} \left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Prove}^{\text{Enc}}((w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}), (s_d^{(j)}, e_{d,\tau}^{(j,\tau)})) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \end{aligned}$$

$\text{View}_2 \stackrel{c}{\approx} \text{View}_3$ because the ciphertext $\text{Enc}_{\text{pk}}(0)$ added to $\tilde{\mathbf{x}}_{d,\tau}^{(k)}$ in Hybrid 2 is *pseudorandom*, and therefore $\tilde{\mathbf{c}}_{d,\tau}^{(k)}$ in Hybrid 2 is pseudorandom as well.

Hybrid 4: We change how the proofs $\phi_{d,\tau}^{(k)}, \pi_{d,\tau}^{(j,k)}$ are created. Instead of computing them honestly using the $\text{Prove}^{\text{Enc}}$ algorithm, we compute simulated proofs using the Sim^{Enc} algorithm.

$$\begin{aligned} \left\{ w_{d,\tau}^{(j,k)} := a_{d,\tau}^{(k)} s_d^{(j)} + 2e_{d,\tau}^{(j,k)} \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Enc}}(w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \end{aligned}$$

$\text{View}_3 \stackrel{c}{\approx} \text{View}_4$ by the *zero-knowledge* property of $\Pi^{(\text{enc})}$.

Hybrid 5: We change how we compute $w_{d,\tau}^{(j,k)}$. We now choose it uniformly at random from R_{qd} .

$$\begin{aligned} \left\{ w_{d,\tau}^{(j,k)} \leftarrow R_{qd} \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \pi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Enc}}(w_{d,\tau}^{(j,k)}, a_{d,\tau}^{(j,k)}) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{c}}_{d,\tau}^{(k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \phi_{d,\tau}^{(k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{c}}_{d,\tau}^{(k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \\ \left\{ \tilde{\mathbf{z}}_{d,\tau}^{(j,k)} \leftarrow R_{qd}^2 \right\}_{j \in \bar{T}, k \in [N]} & , \quad \left\{ \varphi_{d,\tau}^{(j,k)} \leftarrow \text{Sim}^{\text{Rand}}(\tilde{\mathbf{z}}_{d,\tau}^{(j,k)}, \tilde{\mathbf{x}}_{d,\tau}^{(k)}) \right\}_{j \in \bar{T}, k \in [N]} \end{aligned}$$

The fact that $\text{View}_4 \stackrel{c}{\approx} \text{View}_5$ follows from a series of hybrid arguments, one per each $s_d^{(j)}$, in which we use the fact that $w_{d,\tau}^{(j,k)}$ is a PLWE sample with secret $s_d^{(j)}$ and is therefore pseudorandom.

We have proved that $\text{View}_0 \stackrel{c}{\approx} \text{View}_5$. Since View_5 is the view that $\mathcal{S}_{\text{KEYS}}^{\text{MAL}}$ produces, and View_0 is the view produced in an execution of $\Pi_{\text{KEYS}}^{\text{MAL}}$, we have that $\text{IDEAL}_{\mathcal{F}, \mathcal{S}_{\text{KEYS}}^{\text{MAL}}}(\vec{x}) \stackrel{c}{\approx} \text{REAL}_{\Pi_{\text{KEYS}}^{\text{MAL}}, \mathcal{A}}(\vec{x})$. Furthermore, the distribution of ek as output by the simulator is computationally indistinguishable from the output of CombineKeys . This again follows from a series of hybrid arguments, all based on the PLWE assumption with secret s_d . \square