

Fixing Issues and Achieving Maliciously Secure Verifiable Aggregation in “VERIFL: Communication-Efficient and Fast Verifiable Aggregation for Federated Learning”*

Xiaojie Guo
Nankai University
xiaojie.guo@mail.nankai.edu.cn

August 24, 2022

Abstract

This work addresses the security flaw in the original VERIFL protocol and proposes a patched protocol. The patched protocol is secure against any static malicious adversary with a certain threshold and only introduces moderate modifications to the original protocol.

1 Introduction

In this work, we fix the security flaw in our original VERIFL protocol [GLL⁺21] and show that the patched VERIFL protocol can achieve malicious security. We do not claim the contribution of this work. The patched protocol preserves the efficiency feature (i.e., the overall communication between each party and the server is dominated by sending the masked input and receiving the masked aggregate result, no other dimension-dependent communication is required) claimed in the original work. This communication seems to be asymptotically optimal with respect to the input dimension since even *insecure* aggregation consumes the same dimension-dependent bandwidth.

The security flaw. We discuss the security flaw as follows. In brief, the published homomorphic hash values in the verification phase may help the adversary guess the input vector of an honest client if this vector itself does not have sufficient entropy (i.e., has only a few possible values). Note that a linearly homomorphic hash is a *deterministic* function of a input vector. This means that, in the ideal execution, the simulator without knowing the actual inputs of honest parties cannot simulate some hash values such that their distribution is exactly that in the real execution. Even if secret sharing scheme and commitment scheme help us hide the distribution of simulated hash values in the aggregation phase, these values still need to be revealed in the verification phase. Therefore, the distinguisher can distinguish the two executions, failing the security proof.

What CANNOT be promised by maliciously secure verifiable aggregation? The patched VERIFL protocol is maliciously secure in the sense of secure computation and does not prevent attacks that are out of the scope of secure computation, e.g., the poisoning attacks by changing the inputs of corrupt parties, or the inference attacks by using (securely) aggregate results and the inputs of corrupt parties. Moreover, since our protocol performs aggregation over \mathbb{Z}_q^d , we do not prevent the attacker from using crafted inputs to make aggregate results wrap around the modulus q . This wrap-around issue may be (inefficiently) addressed by additionally using range proof (see, e.g., [Bou00]) or other non-cryptographic mechanisms.

2 Preliminaries

2.1 Notations

We use κ to denote the computational security parameter. We use $[n]$ to denote the set $\{1, \dots, n\}$ for some integer $n \in \mathbb{N}$. We denote the set of integers by \mathbb{Z} and the quotient ring of integers modulo an integer $q \in \mathbb{N}$ by \mathbb{Z}_q . For a finite set \mathcal{X} , we use $|\mathcal{X}|$ to denote its size. We write $\text{negl}(\kappa)$ for an unspecified negligible function in κ .

2.2 Security Model

We consider the static malicious adversary in the standard simulation-based proof (see, e.g., [Gol04, Lin16]).

*On May 20, 2022, we found that there is a flaw in the security proof of VeriFL. We thank Jinhyun So for helping us spot this issue. Please cite this work as well if you would like to cite our original paper [GLL⁺21].

2.3 Linearly Homomorphic Hash

For dimension $d \in \mathbb{N}$, a linearly homomorphic hash (LHH) scheme with domain \mathbb{G}_{in}^d and range \mathbb{G}_{out} , where \mathbb{G}_{in} and \mathbb{G}_{out} are Abelian groups, consists of two PPT algorithms with the following syntax:

- $\text{LHHpp} \leftarrow \text{LHH.Setup}(1^\kappa)$. On input 1^κ , output a public parameter LHHpp . For simplicity, we assume that the other algorithm implicitly takes this public parameter as input.
- $h \leftarrow \text{LHH.Hash}(\mathbf{x})$. On input a vector $\mathbf{x} \in \mathbb{G}_{\text{in}}^d$, output its hash value $h \in \mathbb{G}_{\text{out}}$.

We require the above LHH scheme satisfies the following properties:

- **Linearity.** For any two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{G}_{\text{in}}^d$, it holds that

$$\Pr \left[h_{\mathbf{x}} + h_{\mathbf{y}} = \text{LHH.Hash}(\mathbf{x} + \mathbf{y}) \mid \begin{array}{l} \text{LHHpp} \leftarrow \text{LHH.Setup}(1^\kappa), \\ h_{\mathbf{x}} \leftarrow \text{LHH.Hash}(\mathbf{x}), h_{\mathbf{y}} \leftarrow \text{LHH.Hash}(\mathbf{y}) \end{array} \right] = 1.$$

- **Collision resistance.** For any $d \in \mathbb{N}$ and any PPT adversary \mathcal{A} , it holds that

$$\Pr \left[\text{Expt}_{\text{coll}}^{d, \mathcal{A}, \text{LHH}}(1^\kappa) = 1 \right] \leq \text{negl}(\kappa),$$

where the collision experiment $\text{Expt}_{\text{coll}}^{d, \mathcal{A}, \text{LHH}}(1^\kappa)$ is defined as follows:

Experiment $\text{Expt}_{\text{coll}}^{d, \mathcal{A}, \text{LHH}}(1^\kappa)$

1. Generate $\text{LHHpp} \leftarrow \text{LHH.Setup}(1^\kappa)$ and send it to \mathcal{A} .
2. Receive two distinct $\mathbf{x}_0, \mathbf{x}_1 \in \mathbb{G}_{\text{in}}^d$ from \mathcal{A} .
3. Output 1 if $\text{LHH.Hash}(\mathbf{x}_0) = \text{LHH.Hash}(\mathbf{x}_1)$; otherwise output 0.

In this work, we can instantiate this LHH scheme with [BGG95] under the hardness assumption of discrete logarithm over some finite group \mathbb{G} with prime order q . The resulting LHH scheme sets $\mathbb{G}_{\text{in}} = \mathbb{Z}_q$ and $\mathbb{G}_{\text{out}} = \mathbb{G}$.

2.4 Non-interactive Commitment with Equivocality and Straight-line Extractability

We focus on the non-interactive commitment scheme in the random oracle model (ROM) [BR93]. In ROM, a non-interactive commitment scheme with message space \mathcal{M} , randomness space \mathcal{R} , commitment space \mathcal{C} , and a random oracle $\mathcal{O} = \mathcal{O}_{\mathcal{M}, \mathcal{R}, \mathcal{C}}$, consists of two PPT algorithms with the access to \mathcal{O} and the following syntax:

- $\text{com} \leftarrow \text{COM.Commit}^{\mathcal{O}}(m, r)$. On input a message $m \in \mathcal{M}$ and a randomness $r \in \mathcal{R}$, output a commitment $\text{com} \in \mathcal{C}$.
- $\{0, 1\} \leftarrow \text{COM.Open}^{\mathcal{O}}(\text{com}, m', r')$. On input a commitment string $\text{com} \in \mathcal{C}$, a message $m' \in \mathcal{M}$, and a randomness $r' \in \mathcal{R}$, output whether or not the three transcripts are consistent, i.e., $\text{com} = \text{COM.Commit}^{\mathcal{O}}(m', r')$.

We require the above commitment scheme satisfies the following properties:

- **Hiding.** For any PPT adversary $\mathcal{A}^{\mathcal{O}}$, it holds that

$$\left| \Pr \left[\text{Expt}_{\text{hiding}}^{\mathcal{A}^{\mathcal{O}}, \text{COM}, 0}(1^\kappa) = 1 \right] - \Pr \left[\text{Expt}_{\text{hiding}}^{\mathcal{A}^{\mathcal{O}}, \text{COM}, 1}(1^\kappa) = 1 \right] \right| \leq \text{negl}(\kappa),$$

where the hiding experiment $\text{Expt}_{\text{hiding}}^{\mathcal{A}^{\mathcal{O}}, \text{COM}, b}(1^\kappa)$ for $b \in \{0, 1\}$ is defined as follows:

Experiment $\text{Expt}_{\text{hiding}}^{\mathcal{A}^{\mathcal{O}}, \text{COM}, b}(1^\kappa)$

1. Receive two distinct $m_0, m_1 \in \mathcal{M}$ from $\mathcal{A}^{\mathcal{O}}$, sample $r \leftarrow \mathcal{R}$, and send $\text{com} \leftarrow \text{COM.Commit}^{\mathcal{O}}(m_b, r)$ to $\mathcal{A}^{\mathcal{O}}$.
2. Output the guessing bit b' output by $\mathcal{A}^{\mathcal{O}}$.

- **Binding.** For any PPT adversary \mathcal{A}^O , it holds that

$$\Pr \left[\text{Expt}_{\text{binding}}^{\mathcal{A}^O, \text{COM}}(1^\kappa) = 1 \right] \leq \text{negl}(\kappa),$$

where the binding experiment $\text{Expt}_{\text{binding}}^{\mathcal{A}^O, \text{COM}}(1^\kappa)$ is defined as follows:

Experiment $\text{Expt}_{\text{binding}}^{\mathcal{A}^O, \text{COM}}(1^\kappa)$

1. Receive $(\text{com}, (m, r), (m', r'))$ from \mathcal{A}^O , where $m \neq m'$.
2. Output 1 if $(\text{COM.Open}^O(\text{com}, m, r) = 1) \wedge (\text{COM.Open}^O(\text{com}, m', r') = 1)$; otherwise output 0.

- **Equivocality.** In the ideal world, there exists a PPT equivocation algorithm COM.Equiv^O such that, for any $m, m' \in \mathcal{M}$,

$$\Pr \left[\text{COM.Open}^O(\text{com}, m', r') = 1 \mid \begin{array}{l} r \leftarrow \mathcal{R}, \text{com} \leftarrow \text{COM.Commit}^O(m, r), \\ r' \leftarrow \text{COM.Equiv}^O(\text{com}, (m, r), m') \end{array} \right] \geq 1 - \text{negl}(\kappa).$$

- **Straight-line extractability.** In the ideal world, there exists a PPT straight-line extractor COM.Extract^O such that, for any $m \in \mathcal{M}$ and $r \in \mathcal{R}$,

$$\Pr \left[m = m^* \mid \begin{array}{l} \text{com} \leftarrow \text{COM.Commit}^O(m, r), \\ m^* \leftarrow \text{COM.Extract}^O(\text{com}) \end{array} \right] \geq 1 - \text{negl}(\kappa).$$

It is known in [Pas03] that, by using $\mathcal{O} : \mathcal{M} \times \mathcal{R} \rightarrow \mathcal{C}$, the folklore commitment scheme, where $\text{COM.Commit}^O(m, r) := \mathcal{O}(m, r)$, is secure (i.e., hiding and binding) and straight-line extractable in ROM. This scheme also satisfies the equivocality due to the programmability of random oracle: given a well-formed commitment $\text{com} = \text{COM.Commit}^O(m, r)$ and a message m' , the equivocation algorithm COM.Equiv^O samples a fresh randomness $r' \leftarrow \mathcal{R}$ and programs \mathcal{O} such that $\mathcal{O}(m', r') = \text{com}$. It is clear that the equivocality holds if this programming succeeds. The probability that the programming fails is bounded by that for the event that com is not a uniform value or $\mathcal{O}(m', r')$ was defined. Taking a union bound, this probability is at most $2q/|\mathcal{R}| = \text{negl}(\kappa)$ since we only make $q = \text{poly}(\kappa)$ commitments/decommitments and $|\mathcal{R}|$ can be exponentially large in κ .

We instantiate this folklore commitment scheme with $\mathcal{M} = \mathcal{R} = \mathcal{C} = \{0, 1\}^\kappa$ and have the following lemma.

Lemma 1. For any \mathcal{C} such that $|\mathcal{C}| \approx 2^\kappa$ and any PPT algorithm $\mathcal{A}^O(1^\kappa)$ that returns a commitment $\text{com} \in \mathcal{C}$ and an opening string $(m, r) \in \mathcal{M} \times \mathcal{R}$, it holds for the folklore commitment scheme in ROM that

$$\Pr \left[\text{COM.Open}^O(\text{com}, m, r) = 1 \mid \begin{array}{l} (\text{com}, m, r) \leftarrow \mathcal{A}^O(1^\kappa), \\ ((m, r), \text{com}) \notin \mathcal{L} \end{array} \right] \leq \text{negl}(\kappa),$$

where \mathcal{L} denotes the list that records \mathcal{A}^O 's queries to \mathcal{O} and their responses.

Proof. Note that the event $\text{COM.Open}^O(\text{com}, m, r) = 1$ happens if and only if $\text{com} = \text{COM.Commit}^O(m, r) = \mathcal{O}(m, r)$, which in turn occurs with probability at most $1/|\mathcal{C}|$ (exactly equaling $1/|\mathcal{C}|$ if $\mathcal{O}(m, r)$ was not defined) according to the randomness of \mathcal{O} . By using an exponentially large \mathcal{C} , we can see this lemma holds. \square

2.5 Key Agreement

A key agreement scheme with secret key space \mathcal{K}_{sk} , public key space \mathcal{K}_{pk} , and agreed key space \mathcal{K} , consists of three PPT algorithms with the following syntax:

- $\text{KApp} \leftarrow \text{KA.Setup}(1^\kappa)$. On input 1^κ , output a public parameter KApp . For simplicity, we assume that the other algorithms implicitly take this public parameter as input.
- $(\text{sk}, \text{pk}) \leftarrow \text{KA.KeyGen}(1^\kappa)$. On input 1^κ , output a key pair $(\text{sk}, \text{pk}) \in \mathcal{K}_{\text{sk}} \times \mathcal{K}_{\text{pk}}$.
- $\text{ak}_{i,j} \leftarrow \text{KA.Agree}(\text{sk}_i, \text{pk}_j)$. On input a secret key sk_i and a public key pk_j , output an agreed key $\text{ak}_{i,j} \in \mathcal{K}$.

We require the above key agreement scheme satisfies the following properties:

- **Correctness.** It holds that

$$\Pr \left[\text{KA.Agree}(\text{sk}_i, \text{pk}_j) = \text{KA.Agree}(\text{sk}_j, \text{pk}_i) \mid \begin{array}{l} \text{KApp} \leftarrow \text{KA.Setup}(1^\kappa), \\ (\text{sk}_i, \text{pk}_i) \leftarrow \text{KA.KeyGen}(1^\kappa), (\text{sk}_j, \text{pk}_j) \leftarrow \text{KA.KeyGen}(1^\kappa) \end{array} \right] = 1.$$

- **Security.** For any PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Expt}_{\text{key}}^{\mathcal{A}, \text{KA}, 0}(1^\kappa) = 1 \right] - \Pr \left[\text{Expt}_{\text{key}}^{\mathcal{A}, \text{KA}, 1}(1^\kappa) = 1 \right] \right| \leq \text{negl}(\kappa),$$

where the experiment $\text{Expt}_{\text{key}}^{\mathcal{A}, \text{KA}, b}(1^\kappa)$ for $b \in \{0, 1\}$ is defined as follows:

Experiment $\text{Expt}_{\text{key}}^{\mathcal{A}, \text{KA}, b}(1^\kappa)$

1. Generate $\text{KApp} \leftarrow \text{KA.Setup}(1^\kappa)$, $(\text{sk}_i, \text{pk}_i) \leftarrow \text{KA.KeyGen}(1^\kappa)$, and $(\text{sk}_j, \text{pk}_j) \leftarrow \text{KA.KeyGen}(1^\kappa)$.
2. If $b = 0$, sample $\text{ak}_{i,j} \leftarrow \mathcal{K}$. If $b = 1$, generate $\text{ak}_{i,j} \leftarrow \text{KA.Agree}(\text{sk}_i, \text{pk}_j)$.
3. Send $(\text{KApp}, \text{pk}_i, \text{pk}_j, \text{ak}_{i,j})$ to \mathcal{A} and give \mathcal{A} the oracle access to $\text{KA.Agree}(\text{sk}_i, \cdot)$ and $\text{KA.Agree}(\text{sk}_j, \cdot)$.
4. Output the guessing bit b' output by \mathcal{A} .

This key agreement scheme can be instantiated with [BIK⁺17] under the two oracle Diffie-Hellman (2ODH) assumption over some finite group \mathbb{G} with order q . This scheme is with $\mathcal{K}_{\text{sk}} = \mathbb{Z}_q$, $\mathcal{K}_{\text{pk}} = \mathbb{G}$, and $\mathcal{K} = \{0, 1\}^\kappa$.

2.6 Secret Sharing

For threshold $t \in \mathbb{N}$ and number $n \in \mathbb{N}$ of parties where $t \leq n$, a (t, n) -secret sharing scheme with secret space \mathcal{M} , and share space \mathcal{S} , consists of three PPT algorithms with the following syntax:

- $\{\langle s \rangle_i\}_{i \in \mathcal{P}} \leftarrow \text{SS.Share}(s, \mathcal{P})$. On input a secret $s \in \mathcal{M}$ and an n -sized set \mathcal{P} of parties, output a set of secret shares $\{\langle s \rangle_i\}_{i \in \mathcal{P}} \subseteq \mathcal{S}$.
- $\{s', \perp\} \leftarrow \text{SS.Combine}(\{\langle s \rangle_i\}_{i \in \mathcal{P}'}, \mathcal{P}')$. On input a set of secret shares $\{\langle s \rangle_i\}_{i \in \mathcal{P}'} \subseteq \mathcal{S}$ and a set \mathcal{P}' of parties, output a secret $s' \in \mathcal{M}$ or an abort symbol \perp .

We require the above secret sharing scheme satisfies the following properties:

- **Correctness.** For any $t, n \in \mathbb{N}$ where $t \leq n$, any $s \in \mathcal{M}$, any n -sized set \mathcal{P} , and any $\mathcal{P}' \subseteq \mathcal{P}$ where $|\mathcal{P}'| \geq t$, it holds that

$$\Pr \left[s = s' \mid \begin{array}{l} \{\langle s \rangle_i\}_{i \in \mathcal{P}} \leftarrow \text{SS.Share}(s, \mathcal{P}), \\ s' \leftarrow \text{SS.Combine}(\{\langle s \rangle_i\}_{i \in \mathcal{P}'}, \mathcal{P}') \end{array} \right] = 1.$$

- **Reverse samplability.** For the secret sharing scheme SS, there also exists a PPT algorithm SS.RSample such that, for any $t, n \in \mathbb{N}$ where $t \leq n$, any $s \in \mathcal{M}$, any n -sized set \mathcal{P} , any subsets $\mathcal{P}', \mathcal{P}_\perp \subseteq \mathcal{P}$ where $|\mathcal{P}'| \geq t$ and $|\mathcal{P}_\perp| < t$, and any (unqualified) set of shares $\{\langle s \rangle_i\}_{i \in \mathcal{P}_\perp} \in \mathcal{S}^{|\mathcal{P}_\perp|}$, it holds that

$$\Pr \left[s = s' \mid \begin{array}{l} \{\langle s \rangle_i\}_{i \in \mathcal{P} \setminus \mathcal{P}_\perp} \leftarrow \text{SS.RSample}(s, \mathcal{P}, \mathcal{P}_\perp, \{\langle s \rangle_i\}_{i \in \mathcal{P}_\perp}), \\ s' \leftarrow \text{SS.Combine}(\{\langle s \rangle_i\}_{i \in \mathcal{P}'}, \mathcal{P}') \end{array} \right] = 1.$$

- **Security.** For any $t, n \in \mathbb{N}$ where $t \leq n$, and any PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Expt}_{\text{share}}^{t, n, \mathcal{A}, \text{SS}, 0}(1^\kappa) \right] - \Pr \left[\text{Expt}_{\text{share}}^{t, n, \mathcal{A}, \text{SS}, 1}(1^\kappa) \right] \right| \leq \text{negl}(\kappa),$$

where the experiment $\text{Expt}_{\text{share}}^{t, n, \mathcal{A}, \text{SS}, b}(1^\kappa)$ for $b \in \{0, 1\}$ is defined as follows:

Experiment $\text{Expt}_{\text{share}}^{t, n, \mathcal{A}, \text{SS}, b}(1^\kappa)$

1. Receive from \mathcal{A} a secret $s \in \mathcal{M}$, an n -sized set \mathcal{P} , and a subset $\mathcal{P}_\perp \subseteq \mathcal{P}$ where $|\mathcal{P}_\perp| < t$.
2. If $b = 0$, sample $\{\langle s \rangle_i\}_{i \in \mathcal{P}_\perp} \leftarrow \mathcal{S}^{|\mathcal{P}_\perp|}$; otherwise generate $\{\langle s \rangle_i\}_{i \in \mathcal{P}} \leftarrow \text{SS.Share}(s, \mathcal{P})$. Send $\{\langle s \rangle_i\}_{i \in \mathcal{P}_\perp}$ to \mathcal{A} .
3. Output the guessing bit b' output by \mathcal{A} .

This secret sharing scheme can be instantiated with the well-known Shamir's secret sharing [Sha79] with $\mathcal{M} = \mathcal{S} = \mathbb{F}^2$, where \mathbb{F} is a finite field such that $|\mathbb{F}| \geq \max(2^t, q)$.

2.7 Symmetric Encryption

A symmetric encryption scheme with key space \mathcal{K} , plaintext space \mathcal{M} , and ciphertext space \mathcal{C} , consists of three PPT algorithms with the following syntax:

- $k \leftarrow \text{SE.KeyGen}(1^\kappa)$. On input 1^κ , output a key $k \in \mathcal{K}$.
- $c \leftarrow \text{SE.Enc}(k, m)$. On input a key $k \in \mathcal{K}$ and a plaintext $m \in \mathcal{M}$, output a ciphertext $c \in \mathcal{C}$.
- $\{m', \perp\} \leftarrow \text{SE.Dec}(k', c)$. On input a key $k' \in \mathcal{K}$ and a ciphertext $c \in \mathcal{C}$, output a plaintext $m' \in \mathcal{M}$ or an abort symbol \perp .

We require the above symmetric encryption scheme satisfies the following properties:

- **Correctness.** For any $m \in \mathcal{M}$, it holds that

$$\Pr \left[m = m' \mid \begin{array}{l} k \leftarrow \text{SE.KeyGen}(1^\kappa), \\ c \leftarrow \text{SE.Enc}(k, m), m' \leftarrow \text{SE.Dec}(k, c) \end{array} \right] = 1.$$

- **IND-CCA security.** For any PPT adversary \mathcal{A} , it holds that

$$\left| \Pr \left[\text{Expt}_{\text{IND-CCA}}^{\mathcal{A}, \text{SE}, 0}(1^\kappa) = 1 \right] - \Pr \left[\text{Expt}_{\text{IND-CCA}}^{\mathcal{A}, \text{SE}, 1}(1^\kappa) = 1 \right] \right| \leq \text{negl}(\kappa),$$

where the IND-CCA experiment $\text{Expt}_{\text{IND-CCA}}^{\mathcal{A}, \text{SE}, b}(1^\kappa)$ for $b \in \{0, 1\}$ is defined as follows:

Experiment $\text{Expt}_{\text{IND-CCA}}^{\mathcal{A}, \text{SE}, b}(1^\kappa)$

1. Generate $k \leftarrow \text{SE.KeyGen}(1^\kappa)$ and give \mathcal{A} the oracle access to $\text{LR}_{k,b}(\cdot, \cdot)$ and $\text{SE.Dec}(k, \cdot)$, where (i) $\text{LR}_{k,b}(\cdot, \cdot)$ takes two $m_0, m_1 \in \mathcal{M}$ as input and outputs $\text{SE.Enc}(k, m_b)$, and (ii) \mathcal{A} cannot query $\text{SE.Dec}(k, \cdot)$ with any ciphertext output by $\text{LR}_{k,b}(\cdot, \cdot)$.
2. Output the guessing bit b' output by \mathcal{A} .

This symmetric encryption scheme can be instantiated with AES with 128-bit keys.

2.8 Digital signature

A digital signature scheme with message space \mathcal{M} , signing key space $\mathcal{K}_{\text{sigk}}$, verification key space $\mathcal{K}_{\text{verk}}$, and signature space \mathcal{S} , consists of three PPT algorithms with the following syntax:

- $(\text{sigk}, \text{verk}) \leftarrow \text{SIG.KeyGen}(1^\kappa)$. On input 1^κ , output a key pair $(\text{sigk}, \text{verk}) \in \mathcal{K}_{\text{sigk}} \times \mathcal{K}_{\text{verk}}$.
- $\Sigma \leftarrow \text{SIG.Sign}(\text{sigk}, m)$. On input a signing key $\text{sigk} \in \mathcal{K}_{\text{sigk}}$ and a message $m \in \mathcal{M}$, output a signature $\Sigma \in \mathcal{S}$.
- $\{0, 1\} \leftarrow \text{SIG.Ver}(\text{verk}, m, \Sigma)$. On input a verification key $\text{verk} \in \mathcal{K}_{\text{verk}}$, a message $m \in \mathcal{M}$, and a signature $\Sigma \in \mathcal{S}$, output whether or not Σ is a valid signature on m .

We require the above digital signature scheme satisfies the following properties:

- **Correctness.** For any $m \in \mathcal{M}$, it holds that

$$\Pr \left[\text{SIG.Ver}(\text{verk}, m, \Sigma) = 1 \mid \begin{array}{l} (\text{sigk}, \text{verk}) \leftarrow \text{SIG.KeyGen}(1^\kappa), \\ \Sigma \leftarrow \text{SIG.Sign}(\text{sigk}, m) \end{array} \right] = 1.$$

- **Security.** For any PPT adversary \mathcal{A} , it holds that

$$\Pr \left[\text{Expt}_{\text{forge}}^{\mathcal{A}, \text{SIG}}(1^\kappa) = 1 \right] \leq \text{negl}(\kappa),$$

where the experiment $\text{Expt}_{\text{forge}}^{\mathcal{A}, \text{SIG}}(1^\kappa)$ is defined as follows:

Experiment $\text{Expt}_{\text{forge}}^{\mathcal{A}, \text{SIG}}(1^\kappa)$

1. Generate $(\text{sigk}, \text{verk}) \leftarrow \text{SIG.KeyGen}(1^\kappa)$, send verk to \mathcal{A} , and give \mathcal{A} the oracle access to $\text{SIG.Sign}(\text{sigk}, \cdot)$.
2. Receive $(m, \Sigma) \in \mathcal{M} \times \mathcal{S}$ from \mathcal{A} .
3. Output 1 if $\text{SIG.Ver}(\text{verk}, m, \Sigma) = 1$ and \mathcal{A} never queried $\text{SIG.Sign}(\text{sigk}, m)$; otherwise output 0.

In this work, we assume that each party has registered its public verification key in a public bulletin board provided by some trusted public key infrastructure (PKI). The underlying digital signature scheme is standard.

Functionality $\mathcal{F}_{\text{SecVerAgg}}$

Parameters: The threshold t , the number N of parties, the batch size ℓ , and the domain \mathbb{Z}_q^d .

Aggregate: This procedure can be executed ℓ times. Upon receiving (aggregate, σ, i, \mathbf{v}_i), where $\mathbf{v}_i \in \mathbb{Z}_q^d$, from some party $i \in [N]$:

1. Assert $\sigma \in [\ell]$ is the identifier of the current execution. If this is NOT the first (aggregate) message from i under the identifier σ , ignore this message; otherwise, record $(\sigma, i, \mathbf{v}_i)$, send (σ, i) to the server and the adversary, and wait for the server's input:
 - If the server sends (ready, σ, \mathcal{U}_3) such that (i) $\mathcal{U}_3 \subseteq [N]$ and $|\mathcal{U}_3| \geq t$, and (ii) all $\{(\sigma, i, \mathbf{v}_i)\}_{i \in \mathcal{U}_3}$ tuples were recorded, compute $\mathbf{a} := \sum_{i \in \mathcal{U}_3} \mathbf{v}_i \in \mathbb{Z}_q^d$, send (σ, \mathbf{a}) to the adversary, and wait for the adversary's input for each honest party $i \in \mathcal{U}_3$:
 - If the server is corrupt and the adversary sends (deliver, σ, i, \mathbf{a}') for $\mathbf{a}' \neq \mathbf{a}$, set $\text{Cheat}[\sigma, i] = 1$ and send (σ, \mathbf{a}') to the party i .
 - If the adversary sends (deliver, σ, i, \mathbf{a}), set $\text{Cheat}[\sigma, i] = 0$ and send (σ, \mathbf{a}) to the party i .
 - Otherwise, send (abort, i) to the party i and stop sending/receiving messages to/from it.

After the message delivery, ignore future messages to the σ -th execution and move to the $(\sigma + 1)$ -th execution.

- Otherwise, ignore this message.

Verify: Upon receiving (verify, i) from a party $i \in [N]$, send (success, i) to the party i if $\forall \sigma \in [\ell] : \text{Cheat}[\sigma, i] = 0$; otherwise, send (abort, i) to the party i .

Abort: Upon receiving (abort, i) from the adversary, send (abort, i) to the party i and stop sending/receiving messages to/from it.

Figure 1: Functionality of secure verifiable aggregation.

3 Protocol

3.1 High-level Overview

The technique behind our malicious-secure patched protocol is “commit-publish-unmask”. More specifically, all parties should first commit their seeds used to mask their actual inputs and send their masked inputs as well as the LHH values of these masked inputs. Then, the server publishes the **masked aggregate result** to all parties, and the LHH values are broadcast for the future verification that this masked aggregate result is correctly computed. Since the LHH scheme is applied to masked inputs rather than actual inputs, we get rid of the information leakage in the original VERIFL protocol, where the adversary can see the LHH values dependent on the actual inputs of honest parties. Finally, all parties jointly unmask their received masked aggregate result by opening and using the seeds committed before publishing the masked aggregate result.

Roughly speaking, the commitments of these seeds ensure that the corrupt parties should choose *the sum of their inputs* before seeing the aggregate result, and they cannot change this sum after the aggregate result is implicitly determined in the published masked aggregate result without being detected with overwhelming probability. From the perspective of security proof, the use of commitment allows the simulator to extract the seeds to be used by the corrupt parties *before it opening the commitments of honest parties*. Using these seeds and the published masked aggregate result, the simulator can extract the sum of the corrupt parties' inputs, which is sufficient for the simulation.

3.2 Security Proof

We model the ideal secure verifiable aggregation in the functionality $\mathcal{F}_{\text{SecVerAgg}}$ in Figure 1. For simplicity, we focus on the dropout model where any dropped party will never be online again. We consider the case where the server is corrupt (note that the malicious security for an honest server is easier to prove since the simulator can extract the input of *each* corrupt party and the view of the adversary is much simpler). Our patched VERIFL protocol Π_{VERIFL^+} is given in Figure 2 and Figure 3. The security is stated in the following theorem.

Theorem 1. *For any prime modulus $q \approx 2^\kappa$ and any set C of corrupt parties such that $|C| < \min(t, 2t - N)$, Π_{VERIFL^+} securely realizes $\mathcal{F}_{\text{SecVerAgg}}$ against any static malicious adversary that also corrupts the server.*

Proof. Let $C \subseteq [N] \cup \{\text{server}\}$ denote the set of corrupt parties such that $|C \setminus \{\text{server}\}| \leq t$, and $\mathcal{H} := ([N] \cup \{\text{server}\}) \setminus C$ denote the complement. We focus on $\text{server} \in C$. Without loss of generality, we assume the deterministic real-world adversary \mathcal{A} that simply outputs the joint view of corrupt parties and the ideal-world simulator \mathcal{S} that internally runs \mathcal{A} . We prove that the ideal execution with respect to \mathcal{S} is computationally indistinguishable from the real execution with respect to \mathcal{A} via the following hybrid argument, which implicitly gives out the construction of \mathcal{S} .

- Hybrid₀. This is the real execution.
- Hybrid₁. Same as the previous one, except that all honest parties are played by a simulator that is given the inputs of honest parties, emulates the random oracle H on-the-fly, “interacts” with the real-world adversary \mathcal{A} internally run by itself, and

Protocol Π_{VERIFL^+} (Part 1)

Setup: The security parameter κ , the threshold t , the number N of parties, the batch size ℓ , and the domain \mathbb{Z}_q^d . All cryptographic primitives defined in Section 2 and two public parameters: $\text{LHHpp} \leftarrow \text{LHH.Setup}(1^\kappa)$ and $\text{KApp} \leftarrow \text{KA.Setup}(1^\kappa)$. A random oracle $\text{H} : \{0, 1\}^\kappa \rightarrow \mathbb{Z}_q^d$. Each party $i \in [N]$ has a signing key sigk_i and receives from PKI the other parties' verification keys $\{\text{verk}_j\}_{j \in [N] \setminus \{i\}}$.

Verify: In the σ -th execution, the party $i \in [N]$ has input $\mathbf{v}_i^\sigma \in \mathbb{Z}_q^d$ and interacts with the server as follows:

• **Round 0 (Advertising keys)**

For each party $i \in [N]$:

1. Generate $(\text{sk}_i^\sigma, \text{pk}_i^\sigma) \leftarrow \text{KA.KeyGen}(1^\kappa)$, $(\text{msk}_i^\sigma, \text{mpk}_i^\sigma) \leftarrow \text{KA.KeyGen}(1^\kappa)$, and $\Sigma_i^{\text{key}, \sigma} \leftarrow \text{SIG.Sign}(\text{sigk}_i, (\sigma, \text{pk}_i^\sigma, \text{mpk}_i^\sigma))$.
2. Send $(\text{pk}_i^\sigma, \text{mpk}_i^\sigma, \Sigma_i^{\text{key}, \sigma})$ to the server and move to the next round.

Server:

1. Receive messages from at least t parties; otherwise abort. Let $\mathcal{U}_1^\sigma \subseteq [N]$ denote the set of surviving parties.
2. Send to each party $i \in \mathcal{U}_1^\sigma$ the set $K_i^\sigma := \{(j, \text{pk}_j^\sigma, \text{mpk}_j^\sigma, \Sigma_j^{\text{key}, \sigma})\}_{j \in \mathcal{U}_1^\sigma \setminus \{i\}}$ and move to the next round.

• **Round 1 (Sharing metadata)**

For each party $i \in [N]$:

1. Receive K_i^σ from the server, deduce the set \mathcal{U}_1^σ implicit in K_i^σ , and **assert**

$$(\mathcal{U}_1^\sigma \subseteq [N]) \wedge (|\mathcal{U}_1^\sigma| \geq t) \wedge (\forall j \in \mathcal{U}_1^\sigma \setminus \{i\} : \text{SIG.Ver}(\text{verk}_j, (\sigma, \text{pk}_j^\sigma, \text{mpk}_j^\sigma), \Sigma_j^{\text{key}, \sigma}) = 1).$$

2. Sample a self-mask seed $b_i^\sigma \leftarrow \{0, 1\}^\kappa$.
3. Generate $\text{com}_i^{\text{seed}, \sigma} \leftarrow \text{COM.Commit}^O(b_i^\sigma, r_i^\sigma)$ and $\text{com}_i^{\text{msk}, \sigma} \leftarrow \text{COM.Commit}^O(\text{msk}_i^\sigma, s_i^\sigma)$, where $r_i^\sigma, s_i^\sigma \leftarrow \{0, 1\}^\kappa$.
4. Generate t -out-of- $|\mathcal{U}_1^\sigma|$ shares of (b_i^σ, r_i^σ) : $\{(b_{i,j}^\sigma, r_{i,j}^\sigma)\}_{j \in \mathcal{U}_1^\sigma} \leftarrow \text{SS.Share}((b_i^\sigma, r_i^\sigma), \mathcal{U}_1^\sigma)$.
5. Generate t -out-of- $|\mathcal{U}_1^\sigma|$ shares of $(\text{msk}_i^\sigma, s_i^\sigma)$: $\{(\text{msk}_{i,j}^\sigma, s_{i,j}^\sigma)\}_{j \in \mathcal{U}_1^\sigma} \leftarrow \text{SS.Share}((\text{msk}_i^\sigma, s_i^\sigma), \mathcal{U}_1^\sigma)$.
6. For $j \in \mathcal{U}_1^\sigma \setminus \{i\}$, compute $k_{i,j}^\sigma \leftarrow \text{KA.Agree}(\text{sk}_i^\sigma, \text{pk}_j^\sigma)$ and ciphertext $c_{i,j}^\sigma \leftarrow \text{SE.Enc}(k_{i,j}^\sigma, ((b_{i,j}^\sigma, r_{i,j}^\sigma), (\text{msk}_{i,j}^\sigma, s_{i,j}^\sigma)))$.
7. Generate $\Sigma_i^{\text{com}, \sigma} \leftarrow \text{SIG.Sign}(\text{sigk}_i, (\sigma, \text{com}_i^{\text{seed}, \sigma}, \text{com}_i^{\text{msk}, \sigma}))$.
8. Send $(\{(j, c_{i,j}^\sigma)\}_{j \in \mathcal{U}_1^\sigma \setminus \{i\}}, \text{com}_i^{\text{seed}, \sigma}, \text{com}_i^{\text{msk}, \sigma}, \Sigma_i^{\text{com}, \sigma})$ to the server and move to the next round.

Server:

1. Receive messages from at least t parties; otherwise abort. Let $\mathcal{U}_2^\sigma \subseteq \mathcal{U}_1^\sigma$ denote the set of surviving parties.
2. Send to each party $i \in \mathcal{U}_2^\sigma$ the set $C_i^\sigma := \{(j, c_{j,i}^\sigma, \text{com}_j^{\text{seed}, \sigma}, \text{com}_j^{\text{msk}, \sigma}, \Sigma_j^{\text{com}, \sigma})\}_{j \in \mathcal{U}_2^\sigma \setminus \{i\}}$ and move to the next round.

• **Round 2 (Publishing hash & masked aggregate result)**

For each party $i \in [N]$:

1. Receive C_i^σ from the server, deduce the set \mathcal{U}_2^σ implicit in C_i^σ , and **assert**

$$(\mathcal{U}_2^\sigma \subseteq \mathcal{U}_1^\sigma) \wedge (|\mathcal{U}_2^\sigma| \geq t) \wedge (\forall j \in \mathcal{U}_2^\sigma \setminus \{i\} : \text{SIG.Ver}(\text{verk}_j, (\sigma, \text{com}_j^{\text{seed}, \sigma}, \text{com}_j^{\text{msk}, \sigma}), \Sigma_j^{\text{com}, \sigma}) = 1).$$

2. For $j \in \mathcal{U}_2^\sigma \setminus \{i\}$, compute $\text{mak}_{i,j}^\sigma \leftarrow \text{KA.Agree}(\text{msk}_i^\sigma, \text{mpk}_j^\sigma)$.
3. Generate the masked input $\mathbf{p}_i^\sigma := \mathbf{v}_i^\sigma + \text{H}(b_i^\sigma) + \sum_{j \in \mathcal{U}_2^\sigma} \Delta_{i,j} \cdot \text{H}(\text{mak}_{i,j}^\sigma) \in \mathbb{Z}_q^d$, where $\Delta_{i,j} = 1$ if $i < j$, $\Delta_{i,j} = -1$ if $i > j$, and $\Delta_{i,j} = 0$ if $i = j$, the linearly homomorphic hash $h_i^\sigma \leftarrow \text{LHH.Hash}(\mathbf{p}_i^\sigma)$, and $\Sigma_i^{\text{input}, \sigma} \leftarrow \text{SIG.Sign}(\text{sigk}_i, (\sigma, h_i^\sigma))$.
4. Send $(\mathbf{p}_i^\sigma, h_i^\sigma, \Sigma_i^{\text{input}, \sigma})$ to the server and move to the next round.

Server:

1. Receive messages from at least t parties; otherwise abort. Let $\mathcal{U}_3^\sigma \subseteq \mathcal{U}_2^\sigma$ denote the set of surviving parties.
2. Send to each party $i \in \mathcal{U}_3^\sigma$ the transcript $H_i^\sigma := (\{(j, h_j^\sigma, \Sigma_j^{\text{input}, \sigma})\}_{j \in \mathcal{U}_3^\sigma \setminus \{i\}}, \mathbf{p}^\sigma)$, where $\mathbf{p}^\sigma := \sum_{j \in \mathcal{U}_3^\sigma} \mathbf{p}_j^\sigma$, and move to the next round.

Figure 2: Patched VERIFL protocol. If any **assertion** fails, a party will abort without output.

outputs whatever \mathcal{A} outputs. This hybrid is syntactically the same as the real execution, and the joint output distributions in the two hybrids are identical.

- Hybrid₂. Same as the previous one, except that the simulator additionally aborts if there exist two honest parties $i, j \in \mathcal{H}$ such that (i) K_i^σ contains different $(\text{pk}_j^\sigma, \text{mpk}_j^\sigma)$ from that used by the party j in the σ -th **Round 0**, and (ii) the party i 's

• **Round 3 (Checking consistency)**

For each party $i \in [N]$:

1. Receive H_i^σ from the server, deduce the set \mathcal{U}_3^σ implicit in H_i^σ , and **assert**

$$(\mathcal{U}_3^\sigma \subseteq \mathcal{U}_2^\sigma) \wedge (|\mathcal{U}_3^\sigma| \geq t) \wedge (\forall j \in \mathcal{U}_3^\sigma \setminus \{i\} : \text{SIG.Ver}(\text{verk}_j, (\sigma, h_j^\sigma), \Sigma_j^{\text{input}, \sigma}) = 1).$$

2. Define sync message $\text{syncMsg}_i^\sigma := (\{(j, h_j^\sigma, \text{com}_j^{\text{seed}, \sigma}, \text{mpk}_j^\sigma)\}_{j \in \mathcal{U}_3^\sigma}, \{(j, \text{com}_j^{\text{msk}, \sigma})\}_{j \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma})$.
3. Send $\Sigma_i^{\text{sync}, \sigma} \leftarrow \text{SIG.Sign}(\text{sigk}_i, (\sigma, \text{syncMsg}_i^\sigma))$ to the server and move to the next round.

Server:

1. Receive messages from at least t parties; otherwise abort. Let $\mathcal{U}_4^\sigma \subseteq \mathcal{U}_3^\sigma$ denote the set of surviving parties.
2. Send to each party $i \in \mathcal{U}_4^\sigma$ the set $S_i^\sigma := \{(j, \Sigma_j^{\text{sync}, \sigma})\}_{j \in \mathcal{U}_4^\sigma \setminus \{i\}}$ and move to the next round.

• **Round 4 (Opening commitments)**

For each party $i \in [N]$:

1. Receive S_i^σ from the server, deduce the set \mathcal{U}_4^σ implicit in S_i^σ , and **assert**

$$(\mathcal{U}_4^\sigma \subseteq \mathcal{U}_3^\sigma) \wedge (|\mathcal{U}_4^\sigma| \geq t) \wedge (\forall j \in \mathcal{U}_4^\sigma \setminus \{i\} : \text{SIG.Ver}(\text{verk}_j, (\sigma, \text{syncMsg}_j^\sigma), \Sigma_j^{\text{sync}, \sigma}) = 1).$$

2. For $j \in \mathcal{U}_4^\sigma \setminus \{i\}$, compute $(\langle (b_j^\sigma, r_j^\sigma) \rangle_i, \langle (\text{msk}_j^\sigma, s_j^\sigma) \rangle_i) \leftarrow \text{SE.Dec}(k_{i,j}^\sigma, c_{j,i}^\sigma)$.
3. Send $(\langle (b_i^\sigma, r_i^\sigma) \rangle, \{(j, \langle (b_j^\sigma, r_j^\sigma) \rangle_i)\}_{j \in \mathcal{U}_3^\sigma}, \{(j, \langle (\text{msk}_j^\sigma, s_j^\sigma) \rangle_i)\}_{j \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma})$ to the server and move to the next round.

Server:

1. Receive messages from at least t parties; otherwise abort. Let $\mathcal{U}_5^\sigma \subseteq \mathcal{U}_4^\sigma$ denote the set of surviving parties.
2. For $i \in \mathcal{U}_5^\sigma \setminus \mathcal{U}_3^\sigma$, reconstruct $(b_i^\sigma, r_i^\sigma) \leftarrow \text{SS.Combine}(\{\langle (b_i^\sigma, r_i^\sigma) \rangle_j\}_{j \in \mathcal{U}_5^\sigma}, \mathcal{U}_5^\sigma)$.
3. For $i \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma$, reconstruct $(\text{msk}_i^\sigma, s_i^\sigma) \leftarrow \text{SS.Combine}(\{\langle (\text{msk}_i^\sigma, s_i^\sigma) \rangle_j\}_{j \in \mathcal{U}_5^\sigma}, \mathcal{U}_5^\sigma)$.
4. Send to each party $i \in \mathcal{U}_5^\sigma$ the transcript $T_i^\sigma := (\{(j, b_j^\sigma, r_j^\sigma)\}_{j \in \mathcal{U}_3^\sigma}, \{(j, \text{msk}_j^\sigma, s_j^\sigma)\}_{j \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma})$.

• **Round 5 (Unmasking)**

For each party $i \in [N]$:

1. Receive T_i^σ from the server and **assert** that “the two sets in T_i^σ are indexed by \mathcal{U}_3^σ and $\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma$, respectively” and

$$(\forall j \in \mathcal{U}_3^\sigma \setminus \{i\} : \text{COM.Open}^O(\text{com}_j^{\text{seed}, \sigma}, b_j^\sigma, r_j^\sigma) = 1) \wedge (\forall j \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma : \text{COM.Open}^O(\text{com}_j^{\text{msk}, \sigma}, \text{msk}_j^\sigma, s_j^\sigma) = 1).$$

2. For $j \in \mathcal{U}_3^\sigma$ and $k \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma$, compute $\text{mak}_{j,k}^\sigma \leftarrow \text{KA.Agree}(\text{msk}_k^\sigma, \text{mpk}_j^\sigma)$.
3. Compute the aggregate result $\mathbf{a}^\sigma := \mathbf{p}^\sigma - \sum_{j \in \mathcal{U}_3^\sigma} \text{H}(b_j^\sigma) - \sum_{j \in \mathcal{U}_3^\sigma, k \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma} \Delta_{j,k} \cdot \text{H}(\text{mak}_{j,k}^\sigma)$ and output $(\sigma, \mathbf{a}^\sigma)$.

Verify: The party $i \in [N]$ locally proceeds as follows:

1. For $\sigma \in [\ell]$, compute $h^\sigma := \sum_{j \in \mathcal{U}_3^\sigma} h_j^\sigma$ and sample a random coefficient $\alpha_\sigma \leftarrow \mathbb{Z}_q$ (note that \mathbb{Z}_q is included in LHHpp).
2. Compute $h_p \leftarrow \text{LHH.Hash}(\sum_{\sigma \in [\ell]} \alpha_\sigma \cdot \mathbf{p}^\sigma)$ and $h'_p := \sum_{\sigma \in [\ell]} \alpha_\sigma \cdot h^\sigma$.
3. **Assert** $h_p = h'_p$ and output (success, i).

Figure 3: Patched VERIFL protocol. If any **assertion** fails, a party will abort without output.

signature verification for $(\sigma, \text{pk}_j^\sigma, \text{mpk}_j^\sigma)$ passes in the σ -th **Round 1**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the security of the digital signature scheme. As a failure event, the simulator aborts only with negligible probability. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₃. Same as the previous one, except that, for every two honest parties $i, j \in \mathcal{H}$, the simulator replaces the agreed key $k_{i,j}^\sigma$ with a uniformly random key in the σ -th **Round 1**.

Conditioned on that the simulator does not abort, $k_{i,j}^\sigma$ is identical between the two parties in the previous hybrid due to the correctness of the key agreement scheme. Under this condition, any adversary \mathcal{A} that can distinguish the two hybrids with non-negligible probability implies an efficient attacker against the security of the key agreement scheme. Using the law of

total probability, the joint output distributions in the two hybrids are computationally indistinguishable¹.

- Hybrid₄. Same as the previous one, except that, for every two honest parties $i, j \in \mathcal{H}$, the simulator uses an equivalent way to decrypt the ciphertext $c_{j,i}^\sigma$ on behalf of the party i in the σ -th **Round 4**: if $c_{j,i}^\sigma$ equals that generated by the simulator in the σ -th **Round 1**, the simulator retrieves $((b_j^\sigma, r_j^\sigma))_i, ((\text{msk}_j^\sigma, s_j^\sigma))_i$ from its memory without actually performing decryption; otherwise, the simulator honestly decrypts $c_{j,i}^\sigma$ as per the previous hybrid.

This hybrid is syntactically the same as the previous one due to the correctness of the symmetric encryption scheme.

- Hybrid₅. Same as the previous one, except that, for every two honest parties $i, j \in \mathcal{H}$, the simulator replaces the ciphertext $c_{j,i}^\sigma$ with an encryption of 0 in the σ -th **Round 1**.

From the view of \mathcal{A} , the only difference between the two hybrids is the distribution of $c_{j,i}^\sigma$. Any adversary \mathcal{A} that can distinguish the two hybrids with non-negligible probability implies an efficient attacker against the IND-CCA security of the symmetric encryption scheme. The joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₆. Same as the previous one, except that, for every two honest parties $i, j \in \mathcal{H}$, the simulator uses an equivalent way to define the share $((b_j^\sigma, r_j^\sigma))_i, ((\text{msk}_j^\sigma, s_j^\sigma))_i$ in the σ -th **Round 1**: instead of directly generating this share, the simulator “lazily samples” it by running SS.RSample on $((b_j^\sigma, r_j^\sigma), (\text{msk}_j^\sigma, s_j^\sigma))$ and the shares given to all corrupt parties.

This hybrid is syntactically the same as the previous one due to the reverse samplability of the secret sharing scheme.

- Hybrid₇. Same as the previous one, except that, for every honest party $j \in \mathcal{H}$ and corrupt party $i \in \mathcal{C}$, the simulator replaces the ciphertext $c_{j,i}^\sigma$ with an encryption of random shares in the σ -th **Round 1**.

From the view of \mathcal{A} , the only difference between the two hybrids is the distribution of the shares given to the corrupt parties. Since $|\mathcal{C}| < t$, any adversary \mathcal{A} that can distinguish the two hybrids with non-negligible probability implies an efficient attacker against the security of the secret sharing scheme. The joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₈. Same as the previous one, except that, for every honest party $i \in \mathcal{H}$, the simulator (i) generates

$$\text{com}_i^{\text{msk},\sigma} \leftarrow \text{COM.Commit}^O(0, s_i^\sigma)$$

in the σ -th **Round 1**, and (ii) replaces the lazy sampling for $(\text{msk}_i^\sigma, s_i^\sigma)$ with that for $(\text{msk}_i^\sigma, s_i'^\sigma)$, where

$$s_i'^\sigma \leftarrow \text{COM.Equiv}^O(\text{com}_i^{\text{msk},\sigma}, (0, s_i^\sigma), \text{msk}_i^\sigma).$$

The simulator additionally aborts if the equivocation algorithm fails.

The equivocality of the commitment scheme ensures that the simulator can obtain $s_i'^\sigma$ with overwhelming probability, i.e., the probability that the simulator aborts only increases by a negligible value. Conditioned on the successful equivocation, from the view of \mathcal{A} , the only difference between the two hybrids is the distribution of the commitment $\text{com}_i^{\text{msk},\sigma}$. Under this condition, any adversary \mathcal{A} that can distinguish two hybrids with non-negligible probability implies an efficient attacker against the hiding property of the commitment scheme. Using the law of total probability, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₉. Same as the previous one, except that the simulator additionally aborts if there exist two honest parties $i, j \in \mathcal{H}$ such that (i) C_i^σ contains different $(\text{com}_j^{\text{seed},\sigma}, \text{com}_j^{\text{msk},\sigma})$ from that used by the party j in the σ -th **Round 1**, and (ii) the party i 's signature verification for $(\sigma, \text{com}_j^{\text{seed},\sigma}, \text{com}_j^{\text{msk},\sigma})$ passes in the σ -th **Round 2**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the security of the digital signature scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₀. Same as the previous one, except that the simulator additionally aborts if there exist two honest parties $i, j \in \mathcal{H}$ such that (i) H_i^σ contains different h_j^σ from that used by the party j in the σ -th **Round 2**, and (ii) the party i 's signature verification for (σ, h_j^σ) passes in the σ -th **Round 3**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the security of the digital signature scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

¹Note that the transition in this hybrid is independent of the probability that the simulator aborts due to *the conditions in the previous hybrid* so that the established indistinguishability between any two previous adjacent hybrids will not be skewed by the transition. As a result, the simulator still aborts in this hybrid with the *same* negligible probability due to these conditions. This requirement simplifies the probability analysis, and it is also maintained in the subsequent hybrids by carefully arranging their orders.

- Hybrid₁₁. Same as the previous one, except that the simulator additionally aborts if there exist two honest parties $i, j \in \mathcal{H}$ such that (i) $\text{syncMsg}_i^\sigma \neq \text{syncMsg}_j^\sigma$ in the σ -th **Round 3**, and (ii) the party i 's signature verification for $(\sigma, \text{syncMsg}_i^\sigma)$ passes in the σ -th **Round 4**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the security of the digital signature scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

Note that the sync message syncMsg_i^σ is locally defined by the party i . For any two distinct honest parties $i, j \in \mathcal{H}$ such that $\text{syncMsg}_i^\sigma \neq \text{syncMsg}_j^\sigma$, they cannot both survive the assertion in the σ -th **Round 4** conditioned on that the simulator does not abort. Otherwise, there must be at least $t - |C|$ distinct honest parties are required to sign each sync message, contradicting the assumption $2(t - |C|) > N - |C|$.

That is, for the honest parties that do not abort and need to send messages to the server in the σ -th **Round 4**, their sync messages must be identical if the simulator does not abort. The transcripts in this identical sync message, which is denoted by syncMsg^σ , are authenticated so that \mathcal{A} cannot forge the entries of the honest parties. From syncMsg^σ , the simulator can define \mathcal{U}_2^σ and \mathcal{U}_3^σ that are consistent for the surviving honest parties *before* sending any message to \mathcal{A} in this round.

- Hybrid₁₂. Same as the previous one, except that the simulator additionally aborts if (i) there exists some honest party $i \in \mathcal{U}_3^\sigma \cap \mathcal{H}$ such that \mathcal{A} queries $H(b_i^\sigma)$ before the simulator sends any message in the σ -th **Round 4**, or (ii) there exists some honest party $i \in \mathcal{H} \setminus \mathcal{U}_3^\sigma$ such that \mathcal{A} queries $H(b_i^\sigma)$.

In these two hybrids, only the commitment $\text{com}_i^{\text{seed}, \sigma}$ depends on the concrete value of b_i^σ . Any adversary \mathcal{A} can trigger this abort with non-negligible probability must be able to guess the committed b_i^σ with non-negligible probability before seeing the valid opening string. Such an adversary implies an efficient attacker against the hiding property of the commitment scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₃. Same as the previous one, except that the simulator additionally aborts if there exist two distinct honest parties $i, j \in \mathcal{U}_3^\sigma \cap \mathcal{H}$ such that \mathcal{A} queries $H(\text{mak}_{i,j}^\sigma)$.

Conditioned on that the simulator does not abort, $\text{mak}_{i,j}^\sigma = \text{KA.Agree}(\text{msk}_i^\sigma, \text{mpk}_j^\sigma) = \text{KA.Agree}(\text{msk}_j^\sigma, \text{mpk}_i^\sigma)$ is well-defined for the two parties. Also, only the agreed key $\text{mak}_{i,j}^\sigma$ depends on the concrete values of msk_i^σ and msk_j^σ in these two hybrids. Any adversary \mathcal{A} can trigger this abort with non-negligible probability must be able to guess $\text{mak}_{i,j}^\sigma$ with non-negligible probability, i.e., implying an efficient attacker against the security of the key agreement scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₄. Same as the previous one, except that the simulator (i) sets a random $\mathbf{p}_i^\sigma \leftarrow \mathbb{Z}_q^d$ for every honest party $i \in \mathcal{H}$ in the σ -th **Round 2**, and (ii) programs $H(b_i^\sigma)$ for every honest party $i \in \mathcal{U}_3^\sigma \cap \mathcal{H}$ before sending any message to \mathcal{A} in the σ -th **Round 4** such that the following consistency holds:

$$H(b_i^\sigma) = \mathbf{p}_i^\sigma - \mathbf{v}_i^\sigma - \sum_{j \in \mathcal{U}_2^\sigma} \Delta_{i,j} \cdot H(\text{mak}_{i,j}^\sigma).$$

Conditioned on that the simulator does not abort, we claim that the joint output distributions in the two hybrids are identical. On the one hand, for $i \in \mathcal{U}_3^\sigma \cap \mathcal{H}$, \mathbf{p}_i^σ is uniform from the view of \mathcal{A} until receiving messages from the simulator in the σ -th **Round 4** due to (i) the never queried $H(b_i^\sigma)$ in the previous hybrid, or (ii) the uniform sampling in this hybrid. Since $H(b_i^\sigma)$ was not queried by \mathcal{A} before the simulator sends any message in the σ -th **Round 4**, the programming must succeed so that, when the opening string is reconstructed, the consistency holds in both two hybrids from the view of \mathcal{A} . On the other hand, for $i \in \mathcal{H} \setminus \mathcal{U}_3^\sigma$, no programming is performed to maintain the consistency. However, since \mathcal{A} never queries $H(b_i^\sigma)$, this inconsistency cannot be observed by \mathcal{A} , and the resulting \mathbf{p}_i^σ in the previous hybrid is as uniform as that in this hybrid.

It is worth noting that the identical distribution of \mathbf{p}_i^σ , which is a part of the view of \mathcal{A} and affects the abort probability of the simulator, ensures that the transition in this hybrid does not skew the previous computational indistinguishability. Using the law of total probability, the joint output distributions in the two hybrids are identical.

- Hybrid₁₅. Same as the previous one, except that, for every honest party $i \in \mathcal{U}_3^\sigma \cap \mathcal{H}$, the simulator instead programs

$$H(b_i^\sigma) = \mathbf{p}_i^\sigma - \text{sum}_{\text{honest}, i}^\sigma - \sum_{j \in \mathcal{U}_2^\sigma \setminus (\mathcal{U}_3^\sigma \cap \mathcal{H})} \Delta_{i,j} \cdot H(\text{mak}_{i,j}^\sigma),$$

where $\{\text{sum}_{\text{honest}, i}^\sigma\}_{i \in \mathcal{U}_3^\sigma \cap \mathcal{H}}$ are uniform conditioned on $\sum_{i \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \text{sum}_{\text{honest}, i}^\sigma = \sum_{i \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \mathbf{v}_i^\sigma$.

Conditioned on that the simulator does not abort, the joint output distributions in the two hybrids are identical. We define

$$\text{sum}_{\text{honest}, i}^\sigma := \mathbf{v}_i^\sigma + \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \Delta_{i,j} \cdot H(\text{mak}_{i,j}^\sigma),$$

which is uniform from the view of \mathcal{A} since it never queries $H(\text{mak}_{i,j}^\sigma)$ for any $j \neq i$ and $j \in \mathcal{U}_3^\sigma \cap \mathcal{H}$ under the condition. The above definition satisfies the condition on the sum given that $\sum_{i,j \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \Delta_{i,j} \cdot H(\text{mak}_{i,j}^\sigma) = \mathbf{0}$. Using the law of total probability, the joint output distributions in the two hybrids are identical.

- Hybrid₁₆. Same as the previous one, except that the simulator additionally aborts if (i) there exist some honest party $i \in \mathcal{H}$ and some corrupt party $j \in \mathcal{U}_3^\sigma \cap \mathcal{C}$ (or, $j \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{C}$) such that T_i^σ contains different b_j^σ (or, msk_j^σ) from that committed by the party j in the σ -th **Round 1**, and (ii) the party i does not abort in the σ -th **Round 5**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the binding property of the commitment scheme. Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₇. Same as the previous one, except that the simulator additionally aborts if (i) there exist two honest parties $i \in \mathcal{H}$ and $j \in (\mathcal{U}_3^\sigma \cap \mathcal{H}) \setminus \{i\}$ (or, $j \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{H}$) such that T_i^σ contains different b_j^σ (or, msk_j^σ) from that chosen by the simulator on behalf of the party j in the σ -th **Round 1**, and (ii) the party i does not abort in the σ -th **Round 5**.

Any adversary \mathcal{A} that can trigger this abort with non-negligible probability implies an efficient attacker against the binding property of the commitment scheme². Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₈. Same as the previous one, except that the simulator additionally aborts if (i) there exists some corrupt party $j \in \mathcal{U}_2^\sigma \cap \mathcal{C}$ such that syncMsg contains a *malformed* com_j^σ , which is not obtained from \mathcal{A} 's invocation of $\text{COM.Commit}^O(\cdot, \cdot)$, and (ii) all honest parties do not abort in the σ -th **Round 5**.

By Lemma 1, the probability that the simulator aborts due to this condition is negligible (i.e., at most the negligible probability in the lemma times a polynomial number of the opening strings of the corrupt parties' commitments). Using the difference lemma, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₁₉. Same as the previous one, except that the simulator is no longer given the inputs of honest parties and interacts with the functionality $\mathcal{F}_{\text{SecVerAgg}}$ to complete the simulation. More specifically,

- Whenever some honest party i played by the simulator aborts, the simulator sends (abort, i) to $\mathcal{F}_{\text{SecVerAgg}}$.
- Note that syncMsg is identical for all parties in $\mathcal{H}_{\text{survive}}^\sigma \subseteq \mathcal{U}_3^\sigma \cap \mathcal{H}$, the set of the honest parties that are still played by the simulator just *before* sending any message to \mathcal{A} in the σ -th **Round 4**. At this time, the simulator learns $\{h_i^\sigma\}_{i \in \mathcal{U}_3^\sigma}$ from syncMsg and checks, for each honest party $i \in \mathcal{H}_{\text{survive}}^\sigma$, the consistency

$$\sum_{i \in \mathcal{U}_3^\sigma} h_i^\sigma = \text{LHH.Hash}(\mathbf{p}^{\sigma,i}),$$

where $\mathbf{p}^{\sigma,i}$ is the masked aggregate result in H_i^σ received by the party i in the σ -th **Round 3**. Let $\mathcal{H}_{\text{pass}}^\sigma \subseteq \mathcal{H}_{\text{survive}}^\sigma$ denote the set of the honest parties whose $\mathbf{p}^{\sigma,i}$ pass this consistency check. There are three cases for each $\sigma \in [\ell]$:

- Case (i):** $\mathcal{H}_{\text{pass}}^\sigma = \emptyset$. After $\mathcal{H}_{\text{survive}}^\sigma$ and $\mathcal{H}_{\text{pass}}^\sigma$ are determined, the simulator sends $(\text{aggregate}, \sigma, j, \tilde{v}_j)$ to $\mathcal{F}_{\text{SecVerAgg}}$ on behalf of each corrupt party $j \in \mathcal{U}_3^\sigma \cap \mathcal{C}$, where \tilde{v}_j is uniform. Then, the simulator sends $(\text{ready}, \sigma, \mathcal{U}_3^\sigma)$ to $\mathcal{F}_{\text{SecVerAgg}}$ on behalf of the corrupt server, receives back (σ, \mathbf{a}) , and programs H in the σ -th **Round 4** using

$$\sum_{i \in \mathcal{U}_3^\sigma \cap \mathcal{H}} v_i^\sigma := \mathbf{a} - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \tilde{v}_j.$$

Finally, for each honest party $i \in \mathcal{H}_{\text{out}}^\sigma \subseteq \mathcal{H}_{\text{survive}}^\sigma$, where $\mathcal{H}_{\text{out}}^\sigma$ is the set of the honest parties surviving the assertion in the σ -th **Round 5**, the simulator sends $(\text{deliver}, \sigma, i, \mathbf{a}^{\sigma,i})$ to $\mathcal{F}_{\text{SecVerAgg}}$, where

$$\mathbf{a}^{\sigma,i} := \mathbf{p}^{\sigma,i} - \sum_{j \in \mathcal{U}_3^\sigma} H(b_j^\sigma) - \sum_{j \in \mathcal{U}_3^\sigma, k \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma} \Delta_{j,k} \cdot H(\text{KA.Agree}(\text{msk}_k^\sigma, \text{mpk}_j^\sigma))$$

and $(\{b_j^\sigma\}_{j \in \mathcal{U}_3^\sigma}, \{\text{msk}_k^\sigma\}_{k \in \mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma})$ is included in T_i^σ . Other honest parties in $(\mathcal{U}_3^\sigma \cap \mathcal{H}) \setminus \mathcal{H}_{\text{out}}^\sigma$ aborts.

- Case (ii):** $\exists i, j \in \mathcal{H}_{\text{pass}}^\sigma : \mathbf{p}^{\sigma,i} \neq \mathbf{p}^{\sigma,j}$. The simulator aborts in this case.
- Case (iii):** $(\mathcal{H}_{\text{pass}}^\sigma \neq \emptyset) \wedge (\forall i, j \in \mathcal{H}_{\text{pass}}^\sigma : \mathbf{p}^{\sigma,i} = \mathbf{p}^{\sigma,j})$. Let \mathbf{p}^σ denote this identical $\mathbf{p}^{\sigma,i}$ for $i \in \mathcal{H}_{\text{pass}}^\sigma$. The simulator is identical to that in **Case (i)** except that it defines \tilde{v}_j for each corrupt party $j \in \mathcal{U}_3^\sigma \cap \mathcal{C}$ as follows: First, the simulator runs the straight-line extractor COM.Extract^O on the corrupt parties' commitments

$$\{\text{com}_j^{\text{seed},\sigma}\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}, \{\text{com}_j^{\text{msk},\sigma}\}_{j \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{C}}$$

²The reduction uses the fact that, before sending T_i^σ to the party i , \mathcal{A} is given the commitment of the honest party $j \neq i$ and its opening string.

to extract the committed values in these commitments, i.e., $\{b_j^\sigma\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}$ and $\{\text{msk}_k^\sigma\}_{j \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{C}}$. If the extractor fails, the simulator immediately aborts all honest parties still played by it by sending (abort, \cdot) to $\mathcal{F}_{\text{SecVerAgg}}$ and terminates. Then, the simulator extracts the sum of the inputs of the corrupt parties in $\mathcal{U}_3^\sigma \cap \mathcal{C}$:

$$\begin{aligned} \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \mathbf{v}_j^\sigma := & \left(\mathbf{p}^\sigma - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \mathbf{p}_j^\sigma - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}, k \in \mathcal{U}_2^\sigma \cap \mathcal{H}} \Delta_{j,k} \cdot \text{H}(\text{KA.Agree}(\text{msk}_k^\sigma, \text{mpk}_j^\sigma)) \right) \\ & - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \text{H}(b_j^\sigma) - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}, k \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{C}} \Delta_{j,k} \cdot \text{H}(\text{KA.Agree}(\text{msk}_k^\sigma, \text{mpk}_j^\sigma)). \end{aligned}$$

Finally, the simulator samples random $\{\tilde{\mathbf{v}}_j\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}$ conditioned on that $\sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \tilde{\mathbf{v}}_j = \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \mathbf{v}_j^\sigma$.

We claim that the joint output distributions in the two hybrids are computationally indistinguishable. Note that all aborted parties behave identically in the two hybrids. First, we consider this distribution conditioned on that the simulator does not abort due to either the conditions in the previous hybrid or the condition in this hybrid. For any fixed $\sigma \in [\ell]$:

- **Case (i).** The joint distribution of the honest parties' outputs and the corrupt parties' views is identically distributed in the two hybrids *by the end of Aggregate phase* since (i) the simulator also uses the sum $\sum_{i \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \mathbf{v}_i^\sigma$, and (ii) the simulator sets the honest parties' outputs as per the previous hybrid. It is left to consider the cheating flags maintained in $\mathcal{F}_{\text{SecVerAgg}}$ and to be used in **Verify** phase. For each $i \in \mathcal{H}_{\text{out}}^\sigma$, $\text{Cheat}[\sigma, i] = 0$ if and only if

$$\mathbf{a}^{\sigma,i} := \mathbf{a} + \left(\mathbf{p}^{\sigma,i} - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{H}} \mathbf{p}_j^\sigma \right) - \sum_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}} \left(\tilde{\mathbf{v}}_j + \text{H}(b_j^\sigma) + \sum_{k \in \mathcal{U}_2^\sigma \setminus (\mathcal{U}_3^\sigma \cap \mathcal{C})} \Delta_{j,k} \cdot \text{H}(\text{KA.Agree}(\text{msk}_k^\sigma, \text{mpk}_j^\sigma)) \right) = \mathbf{a},$$

where “:=” comes from that $\{b_j^\sigma\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{H}}$ given to the party i are the same as those chosen by the simulator and used in the programming on $\text{H}(b_j^\sigma)$ for each $j \in \mathcal{U}_3^\sigma \cap \mathcal{H}$ (otherwise the simulator will abort). For any $i \in \mathcal{H}_{\text{out}}^\sigma$, $\text{Cheat}[\sigma, i] = 0$ happens with probability $1/q^d$ due to the randomness of $\{\tilde{\mathbf{v}}_j\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}$.

- **Case (iii).** We consider the following two sub-cases.

On the one hand, if \mathcal{A} never uses malformed commitments on behalf of the corrupt parties, the straight-line extractor always succeeds. The reason is that \mathcal{A} cannot open the corrupt parties' commitments to other values than those being committed in the σ -th **Round 1** (otherwise the simulator will abort) and thus there will be only one record in the list maintained by \mathcal{O} for each commitment. Like **Case (i)**, the joint output distribution by the end of **Aggregate** phase is identical. One can also check that, for $i \in \mathcal{H}_{\text{out}}^\sigma \cap \mathcal{H}_{\text{pass}}^\sigma$, $\mathbf{a}^{\sigma,i} = \mathbf{a}$ always holds due to (i) the successful straight-line extraction, (ii) the party-wise consistent \mathbf{p}^σ (by the case definition) and $\{\text{mpk}_j^\sigma\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}$ (by the consistent syncMsg for $\mathcal{H}_{\text{survive}}$), and (iii) the binding ($\{b_j^\sigma\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{C}}$, $\{\text{msk}_k^\sigma\}_{k \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{C}}$) and ($\{b_j^\sigma\}_{j \in \mathcal{U}_3^\sigma \cap \mathcal{H}}$, $\{\text{msk}_k^\sigma\}_{k \in (\mathcal{U}_2^\sigma \setminus \mathcal{U}_3^\sigma) \cap \mathcal{H}}$) (otherwise the simulator will abort). In this sub-case, $\text{Cheat}[\sigma, i] = 0$ happens with probability 1 for $i \in \mathcal{H}_{\text{out}}^\sigma \cap \mathcal{H}_{\text{pass}}^\sigma$; $\mathbf{p}^{\sigma,i} \neq \mathbf{p}^\sigma$ ensures that $\text{Cheat}[\sigma, i] = 0$ happens with probability 0 for $i \in \mathcal{H}_{\text{out}}^\sigma \setminus \mathcal{H}_{\text{pass}}^\sigma$.

On the other hand, if \mathcal{A} uses malformed commitments in the σ -th execution, all honest parties will abort in the σ -th **Round 5** from the view of \mathcal{A} . In the previous hybrid, this abort is due to the assertion in the σ -th **Round 5** (otherwise the simulator will abort). In this hybrid, this abort is due to the simulator sending (abort, \cdot) to $\mathcal{F}_{\text{SecVerAgg}}$ (since the straight-line extractor fails). In this sub-case, the joint output distribution in the two hybrids is **trivially** identical since all honest parties abort before **Verify** phase and the joint output distribution until the previous **Aggregate** phase is identical.

Now, we put these two cases together and, for non-triviality, consider that \mathcal{A} never uses malformed commitments so that $\mathcal{H}_{\text{out}}^\ell \cap \mathcal{H} \neq \emptyset$ for **Verify** phase. For each honest party $i \in \mathcal{H}_{\text{out}}^\ell \subseteq \dots \subseteq \mathcal{H}_{\text{out}}^1$, there must exist a set $\mathcal{I}_i \subseteq [\ell]$ such that $\forall \sigma \in \mathcal{I}_i : i \in \mathcal{H}_{\text{out}}^\sigma \setminus \mathcal{H}_{\text{pass}}^\sigma$ and $\forall \sigma \in [\ell] \setminus \mathcal{I}_i : i \in \mathcal{H}_{\text{out}}^\sigma \cap \mathcal{H}_{\text{pass}}^\sigma$. On the one hand, if $\mathcal{I}_i = \emptyset$, the party i always outputs $(\text{success}, i)$ in this hybrid since $\mathcal{I}_i = \emptyset$ implies, for every $\sigma \in [\ell]$, $i \in \mathcal{H}_{\text{out}}^\sigma \cap \mathcal{H}_{\text{pass}}^\sigma \neq \emptyset$ and $\text{Cheat}[\sigma, i] = 0$. In the previous hybrid, $\mathcal{I}_i = \emptyset$ also makes the party i output $(\text{success}, i)$ due to the definition of $\mathcal{H}_{\text{pass}}^\sigma$ and the linearity of the LHH scheme. On the other hand, it is left to consider $\mathcal{I}_i \neq \emptyset$. By union bound, the party i in this hybrid outputs $(\text{success}, i)$ with probability at most $(1/q^d)^{|\mathcal{I}_i|}$. In the previous hybrid with $\mathcal{I}_i \neq \emptyset$, the party i outputs $(\text{success}, i)$ with probability at most $1/q$, which is taken over the random coefficients. The statistical distance between the two joint distributions in **Verify** phase is bounded by $1/q$, which is negligible. In other words, conditioned on that the simulator does not abort, \mathcal{A} can only distinguish this hybrid from the previous one with negligible advantage.

Then, we bound the probability that the simulator aborts. We have seen from the previous hybrids that the simulator aborts due to the conditions in these hybrids with negligible probability, and this probability is not skewed by the transition in this hybrid. In this hybrid, the simulator additionally aborts due to **Case (ii)**. Note that **Case (ii)** happens with negligible probability; otherwise \mathcal{A} implies an efficient attacker against the collision resistance of the LHH scheme. Overall, the simulator still aborts with negligible probability.

Using the law of total probability, the joint output distributions in the two hybrids are computationally indistinguishable.

- Hybrid₂₀. Same as the previous one, except that the abort conditions for the simulator are removed. This hybrid is computationally indistinguishable from the previous one since it is known from the previous hybrid that the simulator aborts only with negligible probability. This simulator is essentially the simulator \mathcal{S} in the ideal execution.

The above hybrid argument completes this proof. □

References

- [BGG95] Mihir Bellare, Oded Goldreich, and Shafi Goldwasser. Incremental cryptography and application to virus protection. In *27th ACM STOC*, pages 45–56, Las Vegas, NV, USA, May 29 – June 1, 1995. ACM Press.
- [BIK⁺17] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 2017*, pages 1175–1191, Dallas, TX, USA, October 31 – November 2, 2017. ACM Press.
- [Bou00] Fabrice Boudot. Efficient proofs that a committed number lies in an interval. In Bart Preneel, editor, *EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 431–444, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.
- [GLL⁺21] Xiaojie Guo, Zheli Liu, Jin Li, Jiqiang Gao, Boyu Hou, Changyu Dong, and Thar Baker. Verifl: Communication-efficient and fast verifiable aggregation for federated learning. *IEEE Trans. Inf. Forensics Secur.*, 16:1736–1751, 2021.
- [Gol04] Oded Goldreich. *Foundations of Cryptography: Basic Applications*, volume 2. Cambridge University Press, Cambridge, UK, 2004.
- [Lin16] Yehuda Lindell. How to simulate it - A tutorial on the simulation proof technique. Cryptology ePrint Archive, Report 2016/046, 2016. <https://eprint.iacr.org/2016/046>.
- [Pas03] Rafael Pass. On deniability in the common reference string and random oracle model. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 316–337, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany.
- [Sha79] Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.