LERNA: Secure Single-Server Aggregation via Key-Homomorphic Masking

Hanjun Li[1], Huijia Lin[1], Antigoni Polychroniadou[2], and Stefano Tessaro[1]

[1] University of Washington, Seattle, WA, USA
{hanjul,rachel,tessaro}@cs.washington.edu
[2] J.P. Morgan AI Research & AlgoCRYPT CoE,
New York, NY, USA
antigoni.polychroniadou@jpmorgan.com

Abstract. This paper introduces LERNA, a new framework for single-server secure aggregation. Our protocols are tailored to the setting where multiple consecutive aggregation phases are performed with the same set of clients, a fraction of which can drop out in some of the phases. We rely on an initial secret sharing setup among the clients which is generated once-and-for-all, and *reused* in all following aggregation phases. Compared to prior works [Bonawitz et al. CCS'17, Bell et al. CCS'20], the reusable setup eliminates one round of communication between the server and clients per aggregation—i.e., we need two rounds for semi-honest security (instead of three), and three rounds (instead of four) in the malicious model. Our approach also significantly reduces the server's computational costs by only requiring the reconstruction of a *single* secret-shared value (per aggregation). Prior work required reconstructing a secret-shared value for each client involved in the computation.
We provide instantiations of LERNA based on both the Decisional Composite Residuosity (DCR) and (Ring) Learning with Rounding ((R)LWR) assumptions respectively and evaluate a version based on the latter assumption. In addition to savings in round-complexity (which result in reduced latency), our experiments show that the server computational costs are reduced by two orders of magnitude in comparison to the state-of-the-art. In settings with a large number of clients, we also reduce the computational costs up to twenty-fold for most clients, while a small set of "heavy clients" is subject to a workload that is still smaller than that of prior work.

Keywords: Secure Aggregation · Reusable Setup · Privacy Preserving Machine Learning

1 Introduction

A secure aggregation protocol allows a set of clients, each holding an input x_i, to interact with one or more servers, so that the latter learns the sum $\sum x_i$, but no additional information. The inputs x_i could be integers, often mod q, or vectors

of integers. In contrast to the usual setting of multi-party computation, which assumes point-to-point channels, here communication only occurs between each individual client and the server(s), i.e., there is no *direct* inter-client communication and clients can only communicate indirectly through the server(s).

Secure aggregation protocols are suitable for a broad range of applications, such as privacy-preserving telemetry in browsers [16], analytics in digital contact tracing [3], and Federated Machine Learning [10]. Practical *multi-server* protocols [17, 21] are, in fact, already being considered for standardization by IETF [28]. In this paper, however, we target the *single-server* setting. This setting is preferable whenever distributing trust among multiple non-colluding entities is not easily feasible. However, it is also more challenging, as protocols require multiple rounds of interaction and need to accommodate for potential *client dropouts*, whilst ensuring the correctness of aggregation and the privacy of clients' inputs against other colluding clients and/or the server. These protocols have emerged primarily in the context of Federated Machine Learning, starting from Bonawitz et al. [11], which underlies Google's Federated ML system [10], and its recent optimizations and extensions [8, 7].

This paper introduces a new general paradigm for single-server secure aggregation, which improves upon the state-of-the-art in terms of round and computational complexities. Our protocols are particularly advantageous in settings where *repeated aggregation* phases are performed with the same set of clients (some of which may drop out) as they only require two rounds per aggregation, in addition to an initial setup round, at the presence of semi-honest colluding clients and/or server. In comparison, prior protocols [11, 8] require three rounds per aggregation (without initial setup). In the malicious security model, all protocols require one additional round, namely, three rounds in our protocols and four rounds in prior works. Moreover, our approach also significantly improves the server workload by reducing the number of secret-sharing reconstructions.

Repeated Aggregation. While existing single server aggregation protocols mainly focus on running a single aggregation, many scenarios require running repeated aggregation sessions throughout a period of time, with the same set of clients. A prototypical application involves a number of sensors or nodes in a network reporting telemetry data. For example, a company of Internet of Things (IoT) devices may want to aggregate operation data from a certain area periodically to help understand how the devices are used throughout the day. Other examples include wireless sensor networks (WSN) [26], smart meters [4], and medical devices [29].

Our protocol leverages the repeated aggregation setting by having an initial setup round that generates correlated states among clients to facilitate the many aggregation phases later, reducing both round and computational complexity. The protocol is robust to drop-outs, as long as the fraction of drop-out clients is bounded at any point in time. Our main protocol focuses on the setting with a large number of clients, e.g. $M \geq 20K$. To reduce communication costs, it selects a committee of fixed size $O(\kappa^2)$ in the initial setup round to hold the correlated states. And the committee stays unchanged through out many ag-

gregation phases. The protocol guarantees the privacy of clients' inputs against statically corrupted clients that may collude with the server, provided that the total number of corrupted clients in the setup and all aggregation phases are bounded. Compared to protocols designed for single aggregation, we rely on the more stringent condition that the total number of corrupted clients is bounded across many aggregation sessions. However, one can alleviate this assumption by periodically rerunning the setup phase, generating fresh correlated states among clients. Different applications may refresh at a different frequency, say, every day, every week, or even longer, depending on how likely clients are corrupted. Viewed this way, our protocol offers a new tradeoff between the rate of corruption and efficiency gain.

Alternatively, when the number of client is small, e.g. $M \leq 80$, our protocol can avoid the committee in the initial setup round to guarantee stronger privacy: in this setting, the clients may be *adaptively* corrupted instead of statically as assumed above. (See Section E for details on this variant.)

Existing Single-server Secure Aggregation. It is helpful to first review the blueprint behind existing single-server aggregation protocols [11, 8]. Here, we restrict ourselves to the semi-honest setting for simplicity, but these protocols (along with ours) can be modified to support malicious corruption of server and clients.

The initial idea is to have each client $i \in [M]$ send a *masked* input $z_i = x_i + c_i$ to the server. To generate these masks, every pair of clients i, j establishes a shared key $k_{ij} = k_{ji} = \mathrm{PRG}(g^{s_i s_j})$, where g^{s_i} is a group element which acts as an ephemeral *public key* associated with each client $i \in [M]$, and which is shared in an initial round (through the server) with all other clients. The value s_i is kept secret by client i. Then, each client $i \in [M]$ uses the mask

$$c_i = \sum_{j<i} k_{ij} - \sum_{j>i} k_{ij} \ .$$

These masks satisfy in particular the *cancellation property* $\sum_i c_i = 0$, and consequently the server can simply output $\sum_i z_i = \sum_i x_i$.

A first concern is that this only works if each client remains alive and indeed submits its own masked input—a term $k_{ij} = k_{ji}$ included in client j's mask c_j is not canceled out without client i's contribution. To handle a dropout, each client additionally secret shares their own secret s_i, which is reconstructed in case they drop out, to then, in turn, derive all k_{ji}'s for $j \neq i$.

A second concern is that a slow client i could be prematurely labeled as a dropout, and their secret s_i reconstructed *before* the masked value z_i reaches the server, thus revealing x_i. To prevent this, each client initially shares a second random mask b_i, along with s_i, and sends instead the masked input $z_i = x_i + b_i + c_i$ to the server. Then, after receiving the masked inputs $\{z_i\}_{i \in I}$ from a subset $I \subseteq [M]$ of the clients, for each $i \in I$, the server reconstructs b_i, thus allowing the inclusion of $(x_i + c_i)$ in the final sum. In contrast, it reconstructs s_i for all $i \notin I$, thus enabling the computation of $\sum_{i \in I} x_i$ as discussed above. For every

client $i \notin I$, because b_i remains secret, the value x_i remains protected even if later z_i is obtained by the adversary.

Therefore, the overall protocol needs three rounds. An additional round is needed to tolerate a malicious server, and it forces the server to commit to a single set I of clients which are claimed not to have dropped out.

The Costs of Secret Sharing. The most expensive part in the above blueprint is the initial sharing of s_i and b_i, along with the later reconstruction of (one of) them for each client. This impacts both the round and computational complexity in several ways.

Foremost, secret sharing s_i and b_i takes one additional round of communication. While some initial setup round is somewhat inherent (e.g., to share keys to allow clients to communicate with each other via the server), this becomes a bigger concern in the repeated aggregation setting. Here, it is crucial that the values s_i and b_i are re-generated and re-shared *at each repeated session*, for otherwise dropping out at some later session may compromise the privacy of the inputs from prior sessions.

Moreover, the computation and communication costs due to secret sharing are high – $\Theta(M)$ for each client, and $\Theta(M^2)$ for the server. Crucially, the server needs to reconstruct one secret shared value—either s_i or b_i—for each client. In addition, for every client dropout, the server needs to perform $\Theta(M)$ exponentiations to recover the corresponding values k_{ij}. To reduce costs, Bell et. al. [8] proposed to have clients only secret share in a random neighborhood of size $\Theta(\log M + \kappa)$, where κ is the statistical security parameter. Though this idea reduces the client and server costs to $\Theta(\log M + \kappa)$ and $M(\Theta(\log M + \kappa))$, respectively, the improvement is at the cost of weakening the security guarantees at the presence of maliciously corrupted clients and/or server.[3]

Our Contributions. This paper proposes LERNA, a new lightweight approach to single-server secure aggregation which addresses the aforementioned issues. Foremost, it reduces the round complexity to two respectively three communication rounds for semi-honest and malicious security, respectively, in addition to an initial offline round which establishes a setup that can be re-used across multiple aggregations. Moreover, LERNA also features very small server costs, as the server only needs to perform a single reconstruction of a secret-shared value. We validate the performance of LERNA also by benchmarking a prototype implementation.

An important feature of our implementation is that it identifies a (random) subset of the clients as a *committee*. Our benchmarking shows that the computational costs of committee members are smaller than the client costs of prior solutions. However, LERNA is even more lightweight for clients outside of the committee. Indeed, in addition to participating in an initial setup stage, nonmembers only need to send a single message to the server to include their input in an aggregation session, and subsequent interaction within the same session only

[3] More specifically, using the protocol of Bell et. al., if the server is malicious, it may recover the sums of inputs of multiple subsets of clients.

involves committee members. Our benchmarking demonstrates up to twenty-fold performance improvement for these non-committee clients.

A drawback of our solution, as shown in our benchmark, is a relatively heavy communication cost in the initial offline round. This requires participating devices to have sufficient storage and network bandwidth. To amortize this one-time cost, an ideal application for LERNA runs repeated aggregation for large numbers of iterations, T, before rerunning the setup. We envision running LERNA for machine learning from data collected from a large number of relatively powerful devices, e.g. the payment terminals Amazon One, medical imaging devices, weather stations, etc.

Our protocols are built on top of a new primitive, which we call a *key-homomorphic masking* scheme, which allows clients to initially secret share a *re-usable* secret value (i.e., which can be reused across multiple computations) to the committee as part of the initial offline round. We provide two instantiations from, respectively, the DCR assumption [19] and (Ring) LWR assumption [6], with the latter being our main result.

Related Work. The same reduction in round complexity was very recently achieved by Guo et al. [22], also relying on a re-usable secret shared value. However, their solution performs the aggregation in the exponent of a discrete-log hard group, resulting essentially in the sever obtaining the value $g^{\sum_i x_i}$, where g is a group generator. In other words, the actual result can only be extracted by computing the discrete logarithm, which is feasible only if $\sum_i x_i$ is sufficiently small. This forces the computation to be over small domains accommodating Federated learning of models with small weights, such as quantized or compressed models. In contrast, most Federated ML tasks typically involve large values. LERNA does not suffer from this drawback. Our approach differs from [22] in that it relies on different mathematical structures (underlying the LWR and DCR assumptions) to obtain the aggregated sum *in the clear*. This, in turn, requires overcoming a few challenges, in particular, designing special secret-sharing schemes tailored to our requirements – linear reconstruction via small coefficients (for LWR) and working over the integers (for DCR).

The work of [24] proposed a semi-honest protocol, SASH+, using a seed-homomorphic PRG based on LWR similar to our key-homomorphic masking scheme. However, SASH+ exploits the homomorphic property in a different way from LERNA. At high-level, assuming LWR with dimension n, SASH+ reduces the problem of aggregating ℓ-dimension inputs to aggregating n-dimensional homomorphic PRG seeds, which is done using the protocol of [8]. This reduction reduces the computation cost of the server and each client by roughly a factor of (ℓ/n), but at the cost of increasing the round complexity from 3 to 4 per iteration, and introducing an error to the aggregation result that scales linearly with M. In comparison, LERNA reduces the round complexity from 3 to 2, and improves the computation cost at the same time. LERNA also computes the aggregation results exactly without error. As we'll discuss in our benchmarks, LERNA server, and non-committee clients are significantly faster than SASH+'s, while LERNA committee clients become slower than SASH+ clients for very large M.

The work of [31] focuses on the specific application of repeated aggregation in federated machine learning (FL), where the server selects a random subset of clients to aggregate at each iteration. It observes that the usual random client selection strategy in FL causes a leakage of client inputs when the model is close to converged. The paper proposes a new client selection algorithm to mitigate this leakage, assuming an honest server following this new algorithm. We note that LERNA can also be adapted to run repeated aggregation over a different subset of clients at each iteration. The mitigation strategy can then be orthogonally applied to the semi-honest version of LERNA. We stress that the client selection strategy is not to be confused with LERNA's committee selection. Client selection could be added on top of our protocol (but is not included explicitly), and would happen in every iteration, whereas committee selection is within our protocol, and happens only once during its setup phase.

A recent and concurrent work by Bell et al. [7] additionally considers the question of input validation. While this is extremely important, it is orthogonal to the issues studied by this paper. Their system also uses Ring-LWE for efficiency improvement, but still follows broadly the above blueprint without a re-usable setup.

Follow-up Work of [25]. We point interested readers also to the follow-up work [25].

1.1 Overview of LERNA

LERNA's approach differs from the existing protocols in [11, 8] whose core idea is hiding each input x_i with a masks that, as described above, satisfies the cancellation property. Instead, LERNA starts with a conceptually simpler solution, where each client i hides its input x_i with a (random) mask c_i as $z_i = x_i + c_i$, and sends the masked value z_i to the server. With the help of the clients, the server first recovers $c_U = \sum_{i \in U} c_i$, for the set of online clients U, and hence the aggregation result $x_U = \sum_{i \in U} z_i - c_U$. The key question we answer is how the clients securely help the server to compute c_U.

Straw Man Solution. The first naïve idea is to let every client secret share its mask c_i with all other clients using a *linear* secret sharing scheme $(\mathsf{Share}, \mathsf{Recon})$, such as Shamir's secret-sharing scheme. In particular, the Recon algorithm involves evaluating a linear function on the shares. As in prior works [11, 8] each client only has a private and authenticated channel with the server. They can also communicate with each other indirectly through the server. Assuming a PKI setup, such indirect communication can be private and authenticated.

In more detail, each client $i \in [M]$ sends (through the server) the j'th share c_j^i of c_i to each other client $j \in [M]$, before sending their masked input $z_i = c_i + x_i$. The server then finds the set of clients U who have completed both steps, and notifies them of the set U for aggregation. Each client j then locally aggregates the shares it has received from clients $i \in U$, obtaining $c_j^U = \sum_{i \in U} c_j^i$. By the linear homomorphism of the secret sharing, c_j^U is the j'th share of the aggregated mask c_U. As long as enough clients, say $j \in U' \subseteq U$, send their aggregated

shares c_j^U to the server, the latter can reconstruct $c_U = \mathsf{Recon}(\{c_j^U\}_{j \in U'})$, and then recover the aggregated input x_U.

This simple solution is, however inefficient: The step where each client i shares its mask c_i with all other clients has overall $\Omega(M^2)$ communication complexity per aggregation. To aggregate T times, the cost grows as $\Omega(M^2 \times T)$.

Key-homomorphic Masking Scheme. Somewhat informally a key-homomorphic masking scheme involves a pair of algorithms $\mathsf{Mask}, \mathsf{UnMask}$. The Mask algorithm takes an input x from some input space $\mathbb{Z}_p$, a masking key k from some key space $\mathcal{K}$, and a tag τ, and computes a masked message $z \leftarrow \mathsf{Mask}(k, \tau, x)$. The UnMask algorithm takes the above z and an "empty" mask $c \leftarrow \mathsf{Mask}(k, \tau, 0)$ under the same key k and tag τ, and recovers the message $x \leftarrow \mathsf{UnMask}(z, c)$.

Importantly, the scheme is *additively key-homomorphic* for masks with the *same* tag τ: $\mathsf{Mask}(k+k', \tau, x+x') \equiv \mathsf{Mask}(k, \tau, x) \boxplus \mathsf{Mask}(k', \tau, x')$, where $\boxplus$ represents homomorphic addition. We can generalize the additive homomorphism to evaluate any linear function L over masks $\{z_i \leftarrow \mathsf{Mask}(k_i, \tau, x_i)\}$:

$$\mathsf{Eval}(L, \{z_i\}) \equiv \mathsf{Mask}(L(\{k_i\}), \tau, L(\{x_i\})) ,$$

where the linear function L is evaluated respectively over k_i's in the key space $\mathcal{K}$ and over x_i's in the message space $\mathbb{Z}_p$.

Jumping ahead, our instantiation of the masking scheme under LWR will only achieve *approximate* key-homomorphism. We will explain below how we get around this limitation. For now, it is helpful to assume a perfect masking scheme to convey the main idea.

Sketch of the LERNA Protocol. We now describe the semi-honest protocol. Note that the following description depends on a commitment $Q \subseteq [M]$. One can easily think of this committee as containing all clients, although in our concrete instantiation below, we only include a (random) subset of the clients in Q

- Setup phase: The clients agree on a common committee $Q \subseteq [M]$ using public, common randomness. Every client i secret shares a fresh masking key k_i as $\{k_j^i\}_{j \in [Q]}$ and sends the j'th share k_j^i to committee member $j \in Q$.
- Online phase: In the t^{th} aggregation session,
  1. The clients sample a common tag $\tau \leftarrow \mathcal{H}(\mathsf{sid}, t)$ using a hash function $\mathcal{H}$, modeled as a random oracle. Every client P_i then computes a masked input $z_i \leftarrow \mathsf{Mask}(k_i, \tau, x_i)$ and sends z_i to the server.
     The server identifies the set U of online clients. It sends U to all committee members Q, indicating that it wants to aggregate the inputs in U.
  2. Upon receiving U, every committee member P_j aggregates the key shares k_j^i it received from clients $i \in U$, obtaining $k_j^U = \sum_{i \in U} k_j^i$, which by linear homomorphism, equals the j'th share of $k_U = \sum_{i \in U} k_i$. (Therefore, given enough shares $\{k_j^U\}_{j \in U'}$, for a large enough subset U', one can recover k_U.) Then, P_j computes an empty mask $c_j^U \leftarrow \mathsf{Mask}(k_j^U, \tau, 0)$, and sends it back to the server.

Upon receiving enough shares $\{c_j^U\}_{j \in U'}$ from a subset $U' \subseteq U$, the server homomorphically computes the aggregated mask

$$\begin{aligned} c_U &= \mathsf{Eval}(\mathsf{Recon}, \{c_j^U\}_{j \in U'}) \\ &\equiv \mathsf{Mask}(\mathsf{Recon}(\{k_j^U\}_{j \in U'}), \tau, 0) \equiv \mathsf{Mask}(k_U, \tau, 0) \end{aligned}$$

where the first equivalence uses the fact that the Recon algorithm is linear. Similarly,

$$\begin{aligned} z_U &= \sum_{i \in U} z_i = \sum_{i \in U} \mathsf{Mask}(k_i, \tau, x_i) \\ &\equiv \mathsf{Mask}(\sum_{i \in U} k_i, \tau, \sum_{i \in U} x_i) = \mathsf{Mask}(k_U, \tau, x_U) \end{aligned}$$

The server can now recover $x_U = \mathsf{UnMask}(z_U, c_U)$.

LWR-based Instantiation. Our main instantiation of the masking scheme is inspired by the simple seed-homomorphic PRG of [12]. The LWR assumption [6] is associated with two moduli $q > p$, where p is the modulus of the message space. A tag τ is an LWR public vector $\mathbf{a} \in \mathbb{Z}_q^n$, and the masking key k is an LWR secret $\mathbf{s} \in \mathbb{Z}_q^n$. A masked input z is simply an LWR sample rounded to p added with the message x, i.e.,

$$\text{LWR: } \tau = \mathbf{a} \in \mathbb{Z}_q^n, \qquad k = \mathbf{s} \in \mathbb{Z}_q^n, \qquad z = \lfloor \langle \mathbf{s}, \mathbf{a} \rangle \rfloor_p + x \in \mathbb{Z}_p \ .$$

The linear structure of LWR implies the key homomorphism property. However, it only holds *approximately* due to rounding errors. More specifically: *i)* additive key-homomorphism holds approximately with bounded error, and *ii)* linear key-homomorphism holds with bounded error if the linear function L evaluated has *small coefficients*. To see *i)*, consider two masks with keys $k_1 = \mathbf{s}_1, k_2 = \mathbf{s}_2$, inputs x_1, x_2, and a common tag $\tau = \mathbf{a}$. We have

$$\begin{aligned} z_1 + z_2 &= \lfloor \langle \mathbf{s}_1, \mathbf{a} \rangle \rfloor_p + x_1 + \lfloor \langle \mathbf{s}_2, \mathbf{a} \rangle \rfloor_p + x_2 \\ &= \lfloor \langle \mathbf{s}_1 + \mathbf{s}_2, \mathbf{a} \rangle \rfloor_p + x_1 + x_2 + \varepsilon \ , \end{aligned}$$

where ε is the rounding difference between $\lfloor \langle \mathbf{s}_1, \mathbf{a} \rangle \rfloor_p + \lfloor \langle \mathbf{s}_2, \mathbf{a} \rangle \rfloor_p$ and $\lfloor \langle \mathbf{s}_1 + \mathbf{s}_2, \mathbf{a} \rangle \rfloor_p$, which is bounded by 1. With regard to *ii)*, when evaluating a linear function L over the masks using the above approximate additive homomorphism, the error is scaled by the coefficients of L.

The approximate key homomorphism creates a technical issue in the protocol: when the server evaluates Recon homomorphically, it introduces an additive error in the aggregation result. To remove the error, our solution is to multiply the inputs with a scaling factor Δ, set to be larger than the noise.

If the coefficients of Recon are large – e.g., as in Shamir's secret sharing – then the error induced by homomorphic evaluation, and hence the scaling factor Δ becomes large, causing a significant overhead in the protocol. To minimize

this overhead, we will use a linear secret-sharing scheme whose reconstruction function has only $-1, 0, 1$ coefficients – referred to as the *flatness* property. An additional benefit of the flatness property is that Recon becomes computationally cheaper, involving only simple additions and subtractions.

Committee Based Flat Secret Sharing Scheme. As motivated above, we need a secret sharing scheme with small reconstruction coefficients. One solution appears to come from the work of [20], which transforms any monotone Boolean formula for the threshold function into a linear secret-sharing scheme with small coefficients, satisfying flatness. Unfortunately, however, known constructions of Boolean formulae for the threshold function with M inputs has a size $\Omega(M^{5.3})$ [32], which by the transformation of [20] gives a secret sharing consisting of $\Omega(M^{5.3})$ elements in total. This is prohibitively expensive and recent work [5] indicates several challenges in improving this.

Our committee-based construction follows the blueprint of [20], but drastically reduces the total share size from $\Omega(M^{5.3})$ to $\Theta(\kappa^2)$ where κ is the security parameter. Our key observation is that in the setting of secure aggregation, a much weaker secret sharing scheme (than that of [20]) suffices:

1. Instead of using a monotone Boolean formula for threshold functions, it suffices to consider *gap threshold* functions. Such a function outputs 1 if more than ρ fraction of the inputs are 1 and outputs 0 if less than $\gamma < \rho$ fraction of the inputs are 1 (and has no guarantees for inputs in between). The values of ρ and γ correspond to the reconstruction and privacy thresholds in the context of secret sharing.
2. Instead of using a single formula, we use a *distribution* $\mathcal{F}$ of formulae. Our secret-sharing scheme has a setup phase where a formula is sampled $f \leftarrow \mathcal{F}$. As such, the security and correctness of secret sharing only need to hold with overwhelming probability over the random choice of f.
   Sampling $f \leftarrow \mathcal{F}$ directly translates to sampling a committee of share holders in the secret sharing scheme, corresponding to the committee Q chosen in the setup phase of our protocol above.
3. In fact, we do not even need formulae that compute exactly the gap threshold function. Instead, it suffices if for every "promised" input x, a random formula $f \leftarrow \mathcal{F}$ computes the correct output with overwhelming probability. That is,

$$\begin{aligned}&\forall x \text{ with hamming weight } < \gamma M \text{ or } > \rho M,\\ &\Pr\left[f(x) \text{ correct } | f \leftarrow \mathcal{F}\right] > 1 - \mathrm{negl}(\kappa).\end{aligned}$$

These relaxations allow us to modify the randomized construction of formulae for threshold function in [32] to obtain a distribution of formulae with sizes $\Theta(\kappa^2)$ satisfying the above. The transformation of [20] then gives a committee based secret sharing consisting of only $\Theta(\kappa^2)$ elements in total, with a $\Theta(\kappa^2)$ size committee.

Server Efficiency. LERNA admits very efficient server computation. Upon collecting all the masked inputs z_i and all the mask shares c_i^U, the server simply computes a sum $\sum_{i \in U} z_i$, reconstruction over shares c_i^U, and finally unmasks.

Since our secret sharing has 0/1 coefficients, reconstruction is also computing a sum. LERNA server is 100× faster than that of prior work [8], where the server needs to perform $\Theta(M)$ reconstruction of Shamir's secret sharing, and $\Theta(M(\log M + \kappa))$ group exponentiations. See Section 5 and Section 6 for asymptotic comparisons and experimental data.

Static vs. Adaptive Corruptions. One consequence of the above approach is that the random choices involved in sampling the formula (i.e., the committee of share holders) need to be independent of corruptions and dropouts in an execution of the protocol, which we expect to be chosen non-adaptively (for dropouts, in fact, we only require an overall set of potential dropouts to be fixed non-adaptively, but when individual parties drop out can be chosen adaptively). We stress that this assumption already inherently underlies the optimized aggregation protocol from [8], which relies on choosing a random graph independently of corruption and dropout patterns.

For the setting where the number of clients is small, e.g. $M \leq 80$, we show an alternative instantiation of LERNA that doesn't involve sampling a committee of share holders in Section E. In this variant, LERNA tolerates adaptive corruption.

2 Preliminaries

In this section, we explain the system and failure models of LERNA, and give an overview of LERNA's security requirements. We provide a formal security definition in the UC framework in Appendix B.

System Model. LERNA is a framework for secure aggregation involving M clients and a single server. Different from the systems in [11, 8], LERNA has a one-time setup phase followed by many, T, online phases (also referred to as aggregation sessions). The setup phase creates correlated secrets $s_1, \ldots, s_M$ among the M clients, which are re-used in all following online phases. During each online phase, the server computes the aggregation over fresh inputs $\mathbf{x}_1, \ldots, \mathbf{x}_M$ from the same set of clients. The inputs to the clients $\mathbf{x}_i \in \mathbb{Z}^\ell$ are large integer vectors from a bounded (but potentially exponentially large) range, and the aggregation results are computed coordinate-wise over the integers.

Communication Model. Similar to prior work, LERNA has a simple communication pattern. During the online phases, each client communicates only with the server through private and authenticated channels. During the setup phase, the clients communicate indirectly with each other through the server, also in a private and authenticated way. This can be achieved by assuming a PKI setup, or, to avoid the PKI setup, the clients can run pairwise key-agreement through the server at the beginning of the setup phase. We need to assume (similarly to [11, 8]) the server behave honestly in the key-agreement round.

The LERNA protocol proceeds in rounds. In each round, each client may send one message to the server, and may receive a reply message from the server. For simplicity, we assume synchronized communication channels.

Failure Model. LERNA is designed to be robust against two types of failures, corruption and dropout. For the first type of failure, a subset of the parties, may or may not include the server, collude to try to learn the individual input of the other clients. We further differentiate static and adaptive corruptions. In static corruption, the adversary selects a subset of corrupted parties at the beginning of the protocol execution. In adaptive corruption, the adversary is free to choose which party to corrupt at any stage of the protocol execution. Our main protocol in Section 4, suitable for running with large number of clients, tolerates static corruption. The variant described in Section E for running with small number of clients tolerates adaptive corruption. In the semi-honest setting, the adversary learns the inputs and the internal states of the corrupted parties, throughout the setup phase and all online phases. In the malicious setting, the adversary controls the actions of the corrupted parties entirely.

For the second type of failure, a potentially different subset of clients drop out from each online phase (and may come back in the future). We model no clients dropout during the setup phase. This is equivalent to saying only the set of clients who complete the setup phase is considered during the following online phases. More precisely, we model the dropout failure by allowing the adversary to choose a set of potential dropout clients D_t for each online phase t, all at the beginning of the protocol. The adversary is allowed to adaptively decide whether and when each client $P_i \in D_t$ (from the potential set) actually drops out during the online phase t.

Security Definition. The security of LERNA has two aspects: correctness and privacy. They are parameterized by a constant fraction δ, which represents the fraction of dropout clients tolerated by LERNA.

Correctness guarantees that in the semi-honest setting, the server computes the correct result in a session, as long as less than δM clients drop out in that session. In contrast, in the malicious setting, a corrupted client may arbitrarily "pollute" the aggregation result or cause it to be $\perp$, indicating an error.

For privacy, we consider an adversary that *statically* corrupts at most a γ fraction of the clients, before the aggregation protocol begins. We tolerate any fraction $0 \leq \gamma < 1 - \delta$. The adversary may additionally corrupt the server. The following privacy guarantee applies to both the semi-honest and the malicious settings.

In the simpler case, where only clients but not the server are corrupted, the adversary learns only the corrupted clients' inputs in each aggregation session and nothing else. In the case where the server is also corrupted, the adversary learns the corrupted clients' inputs, as well as a single sum of the honest clients' inputs in a sufficiently large set $U \subseteq [M]$, where $|U| > (1 - \delta)M$.

For comparison, the privacy guarantee of [8] is weaker. In the case where both the server and a subset of the clients are corrupted, an adversary may learn *multiple* non-overlapping sums of the honest inputs in each aggregation session. Their security guarantees that each such sum contains at least $\Omega(\log M)$ inputs, which provides a weaker degree of anonymity.

Formally, we define the security of LERNA in the UC framework [15]. Details of the UC framework and our formal security definition are deferred to Appendix B.

3 Technical Tools

In this section, we construct two technical tools, a key-homomorphic masking scheme and a flat secret sharing scheme. As outlined in the technical overview (Section 1.1), the masking scheme is used for hiding clients' input vectors, and the secret sharing scheme is used for sharing each client's secret masking key.

3.1 Key-homomorphic Masking

We first introduce the syntax of a key-homomorphic masking scheme.

- $\mathsf{Setup}(1^\lambda, \ell, B_{\mathrm{msg}})$: takes as inputs the security parameter λ, a message dimension ℓ, and a lower bound B_{msg} on the message modulus. It outputs public parameters pp, which defines a key space $\mathcal{K}$, a message space $\mathbb{Z}_{p_m}^\ell$ with some modulus $p_m \geq B_{\mathrm{msg}}$, and a mask space $\mathbb{Z}_q^\ell$ with some modulus q.

In our framework, we assume every client enters the setup phase (Figure 1) with common correctly generated public parameters pp. If the Setup algorithm is deterministic, or public-coin, then this assumption is simply a notational convenience, since each client can compute the common pp on its own, using a random oracle to derive common public randomness if necessary.

- $\mathsf{KeyGen}(\mathsf{pp})$: outputs a masking key $k \in \mathcal{K}$.
- $\mathsf{TagGen}(\mathsf{pp})$: outputs a tag τ.

In our framework, each client P_i derives its secret masking key k_i in the setup phase, and re-uses it during all online phases. In contrast, it derives a fresh tag τ for each online phase, using common public randomness. We only require the key-homomorphic property to hold for masks under the same tag τ. While the tag is public to all clients, the masking keys must remain secret.

- $\mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{m})$: takes as inputs a masking key $k \in \mathcal{K}$, a tag τ, and a message $\mathbf{m} \in \mathbb{Z}_{p_m}^\ell$, and outputs a masked message $\mathbf{c}_m$.
- $\mathsf{UnMask}(\mathsf{pp}, \mathbf{c}_m, \mathbf{c}_0)$: takes as inputs a masked message $\mathbf{c}_m$, and an "empty" mask $\mathbf{c}_0$ (of message $\mathbf{0}$) under the same key and tag. It recovers a message $\mathbf{m}^*$ or $\perp$.

The UnMask algorithm is a bit unusual, as it doesn't take the masking key k or the tag τ to recover the message. Instead, it asks the caller to first compute an empty mask $\mathbf{c}_0$ using the key k and tag τ, and then feed $\mathbf{c}_0$ to the algorithm. We define such a syntax because in our framework, the caller of UnMask is the server. The clients jointly help the server compute the empty mask $\mathbf{c}_0$, instead of revealing their masking keys, so that the keys remain secret during each online phase.

– $\mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i\})$: takes as inputs a linear function L with d integer coefficients and d masks $\{\mathbf{c}_i\}_{i\in[d]}$. It homomorphically evaluates L on the masks and outputs the result $\mathbf{c}_L$.

As mentioned earlier, the input masks $\{\mathbf{c}_i\}$ to the Eval algorithm should be masked under a common tag τ. Evaluating L on the masks roughly translates to evaluating L on both the masking keys, over the key space $\mathcal{K}$, and over the the messages, over the message space $\mathbb{Z}_{p_m}^{\ell}$. We define this property below as key-homomorphism.

Correctness. Formally, we define the correctness and the key-homomorphism requirement as follows.

Definition 1 (correctness)**.** *For all public parameters* pp*, tags* τ*, and keys* k *output by* Setup*,* TagGen *and* KeyGen*, and for all messages* $\mathbf{m} \in \mathbb{Z}_{p_m}^{\ell}$*, the following holds.*

$$\Pr\left[\mathsf{UnMask}(\mathsf{pp}, \mathbf{c}_m, \mathbf{c}_0) = \mathbf{m} \;\middle|\; \begin{array}{c} \mathbf{c}_m \leftarrow \mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{m}), \\ \mathbf{c}_0 \leftarrow \mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{0}). \end{array}\right] = 1.$$

Definition 2 (key-homomorphism)**.** *Consider any linear function* L*, represented by* d *integer coefficients. For all public parameters* pp*, tag* τ*, and keys* $k_1, .., k_\ell$ *output by* Setup*,* TagGen*, and* KeyGen*, and all messages* $\mathbf{m}_1, ..., \mathbf{m}_d \in \mathbb{Z}_{p_m}^{\ell}$*, the following holds.*

$$\begin{aligned} &\left\{\tilde{\mathbf{c}}_L \leftarrow \mathsf{Mask}\big(\mathsf{pp}, L(\{k_i\}), \tau, L(\{\mathbf{m}_i\})\big)\right\} \\ \equiv &\left\{\mathbf{c}_L \leftarrow \mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i\}) \mid \mathbf{c}_i \leftarrow \mathsf{Mask}(\mathsf{pp}, k_i, \tau, \mathbf{m}_i)\right\}, \end{aligned}$$

where $L(\{\mathbf{m}_i\})$ *is evaluated over* $\mathbb{Z}_{p_m}^{\ell}$ *and* $L(\{k_i\})$*, over the key space* $\mathcal{K}$*.*

The key-homomorphism definition above requires the evaluated mask $\mathbf{c}_L$ to have the same distribution as the "target" mask $\tilde{\mathbf{c}}_L$. We next introduce a relaxation to this rather strong property. Roughly, the evaluated mask $\mathbf{c}_L$ should be distributed close to the target mask $\tilde{\mathbf{c}}_L$. In other words, through homomorphic evaluation we obtain the target mask with some bounded additive noise.

Our framework requires two additional properties from an approximate key-homomorphic scheme. First, when computing UnMask on a masked input $\mathbf{c}_m$ and an empty mask $\mathbf{c}_0$, any additive noises in them translate to additive noises in the recovered message. Second, when computing Eval on noisy masks, the additive noises translates to an additive noise in the evaluated mask, with bounded magnitude. We formalize the above requirements as follows.

Definition 3 (ε-approximate key-homomorphism)**.** *Consider any linear function* L*, with* d *integer coefficients whose absolute values are bounded by some* $B_L \in \mathbb{N}$*.*

– *Let* $\tilde{\mathbf{c}}_L$*,* $\mathbf{c}_L$ *be the evaluated and the "target" masks as defined in Definition 2. We require*

$$\|\tilde{\mathbf{c}}_L - \mathbf{c}_L\|_\infty \le \varepsilon d B_L.$$

– *Let* $\mathbf{c}_m, \mathbf{c}_0$ *be the masked input and the empty mask as defined in Definition 1. For all integer noise vectors* $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{Z}^\ell$*, we require*

$$\mathsf{UnMask}(\mathsf{pp}, \mathbf{c}_m + \mathbf{e}_1, \mathbf{c}_0 + \mathbf{e}_2) = \mathbf{m} + \mathbf{e}_1 + \mathbf{e}_2 \bmod p_m.$$

– *Let* pp*,* $\{\mathbf{c}_i\}$ *be the public parameters and the masks as defined in Definition 2. For all integer noise vectors* $\{\mathbf{e}_i\}$*, whose values are bounded by some* $B_e \in \mathbb{N}$ *we require*

$$\|\mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i\}) - \mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i + \mathbf{e}_i\})\|_\infty \leq B_e d B_L.$$

Security. For security, we require a mask under a randomly chosen key hides its message. We further require this holds for a polynomial number of adaptive sessions, each with a fresh tag sampled with public randomness, reusing the same key.

Definition 4 (security)**.** *Let* λ *be the security parameter. The masking scheme is* secure *if for all input dimension* $\ell = \ell(\lambda) \leq \mathrm{poly}(\lambda)$ *and message modulus lower bound* $B_{\mathrm{msg}} = B_{\mathrm{msg}}(\lambda) \leq 2^{\mathrm{poly}(\lambda)}$*, any efficient adversary* $\mathcal{A}$ *has negligible advantage in distinguishing the experiments* $\mathsf{Exp}^{\mathcal{A},b}_{\mathsf{Mask}}(1^\lambda)$ *defined as follows:*

– *The challenger computes* $\mathsf{pp} \leftarrow \mathsf{Setup}(1^\lambda, \ell, B_{\mathrm{msg}})$*, and samples a masking key* $k \leftarrow \mathsf{KeyGen}(\mathsf{pp})$*. It launches* $\mathcal{A}(1^\lambda)$*, sends* pp *to* $\mathcal{A}$*, and repeats the following steps until* $\mathcal{A}$ *outputs a bit* b'*.*
  1. *Run* $\tau \leftarrow \mathsf{TagGen}(\mathsf{pp}; r)$ *using fresh randomness* r*, and send* (τ, r) *to* $\mathcal{A}$*.* $\mathcal{A}$ *replies with a message* $\mathbf{m} \in \mathbb{Z}^\ell_{p_m}$*.*
  2. *If* $b = 1$*, compute* $\mathbf{c}_1 \leftarrow \mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{m})$*. Otherwise, compute* $\mathbf{c}_0 \leftarrow \mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{0})$*. Send* c_b *to* $\mathcal{A}$*.*

Construction Based on LWR. We construct a 1-approximate key-homomorphic masking scheme based on the learning with rounding (LWR) assumption[6]. The construction is a slight modification to the almost seed homomorphic PRG based on LWR in [12].

Definition 5 (LWR [6])**.** *let* λ *be the security parameter,* $n = n(\lambda)$*,* $q = q(\lambda)$*,* $p = p(\lambda)$ *be integers. The* $LWR_{n,q,p}$ *assumption states that for any* $m = \mathrm{poly}(n)$ $A \leftarrow \mathbb{Z}_q^{m \times n}$*,* $\mathbf{s} \leftarrow \mathbb{Z}_q^n$*,* $\mathbf{u} \leftarrow \mathbb{Z}_q^m$*, the following indistinguishability holds:*

$$(A, \lfloor A \cdot \mathbf{s} \rceil_p) \approx^c (A, \lfloor \mathbf{u} \rceil_p),$$

where $\lfloor \cdot \rceil_p$ *is the rounding function defined as* $\lfloor \cdot \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p : x \mapsto \lfloor (p/q) \cdot x \rceil$*.*

Construction 1 (key-homomorphic masking by LWR)**.**

– $\mathsf{Setup}(1^\lambda, \ell, p_m)$: *deterministically* choose a modulus q and dimension n such that LWR_{n,q,p_m} is assumed to be hard. Output $\mathsf{pp} = (\ell, p_m, q, n)$. The key space is $\mathcal{K} = \mathbb{Z}_q^n$, the message space, $\mathbb{Z}^\ell_{p_m}$, which is the same as the mask space.
– $\mathsf{KeyGen}(\mathsf{pp})$: sample a vector $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, and output $k = \mathbf{s}$.

- $\mathsf{TagGen}(\mathsf{pp})$: sample a matrix $A \leftarrow \mathbb{Z}_q^{n\times\ell}$, and output $\tau = A$.
- $\mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{m})$: parse the key and tag as $k, \tau = \mathbf{s}, A$. Output the masked message $\mathbf{c}_m = \lfloor A \cdot \mathbf{s} \rfloor_{p_m} + \mathbf{m} \in \mathbb{Z}_{p_m}^{\ell}$.
- $\mathsf{UnMask}(\mathsf{pp}, \mathbf{c}_m, \mathbf{c}_0)$: output the message $\mathbf{m}^* = \mathbf{c}_m - \mathbf{c}_0 \in \mathbb{Z}_{p_m}^{\ell}$.
- $\mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i\})$: parse L as d integer coefficients $u_1, \ldots, u_d$. Output the evaluated mask $\mathbf{c}_L = \sum_{i\in[d]} u_i \mathbf{c}_i \in \mathbb{Z}_{p_m}^{\ell}$.

The idea of the construction is simple. A masking key is an LWR secret $k = \mathbf{s}$, and a tag is a random LWR public matrix $\tau = A$. Given a masking key $\mathbf{s}$, a tag A, and a message $\mathbf{m}$ as inputs, the Mask algorithm hides the message $\mathbf{m}$ with a fresh LWR sample $\lfloor A \cdot \mathbf{s} \rfloor_{p_m}$.

Lemma 1. *Construction 1 is a 1-approximate key-homomorphic masking scheme under the LWR_{n,q,p_m} assumption.*

Proof (of Lemma 1). The correctness of the above construction is clear, and the security of the above construction follows immediately from the LWR assumption. We show that the above construction has 1-approximate key-homomorphism.

Consider first the simple case of adding two masks of dimension $\ell = 1$. Let $\tau = \mathbf{a} \in \mathbb{Z}_q^n$ be any tag, $k_1 = \mathbf{s}_1, k_2 = \mathbf{s}_2 \in \mathbb{Z}_q^n$ be any two keys, and $m_1, m_2 \in \mathbb{Z}_{p_m}$ be any two messages. The two masks c_1, c_2, and the target mask $\tilde{c}_{\text{sum}}$ are computed as

$$c_1 = \lfloor \langle \mathbf{s}_1, \mathbf{a} \rangle \rfloor_{p_m} + m_1, \quad c_2 = \lfloor \langle \mathbf{s}_2, \mathbf{a} \rangle \rfloor_{p_m} + m_1,$$
$$\tilde{c}_{\text{sum}} = \lfloor \langle \mathbf{s}_1 + \mathbf{s}_2, \mathbf{a} \rangle \rfloor_{p_m} + m_1 + m_2.$$

The difference between the evaluated mask $c_{\text{sum}} = c_1 + c_2$ and the target mask $\tilde{c}_{\text{sum}}$ is simply the rounding error

$$e = \lfloor \langle \mathbf{s}_1, \mathbf{a} \rangle \rfloor_{p_m} + \lfloor \langle \mathbf{s}_2, \mathbf{a} \rangle \rfloor_{p_m} - \lfloor \langle \mathbf{s}_1 + \mathbf{s}_2, \mathbf{a} \rangle \rfloor_{p_m} \in \{0, 1\}.$$

Generalizing the above, we can conclude that multiplying a mask c by a coefficient u causes an error $|e| \leq |u|$. And evaluating a linear function over d masks with coefficients bounded by B_L causes an error $|e| \leq dB_L$. If the masks already contain errors bounded by B_e, then the evaluation amplifies it to an additional error $|e| \leq B_e dB_L$. We have verified that Construction 1 has 1-approximate key-homomorphism per Definition 3. □

Choosing Parameters q, n. It is proved in [6] that under the Learning With Error (LWE) assumption with dimension n, modulus q, and any noise distribution bounded by B, the LWR assumption also holds with dimension n and moduli q, p_m such that $q \geq B p_m n^{\omega(1)}$. It's commonly believed that the LWE assumption holds for sufficiently large $B = \text{poly}(n)$, and sub-exponential modulus-to-noise ratio $\alpha = q/B \leq 2^{\sqrt{n}}$. Therefore, given a message modulus $p_m \in \mathbb{N}$, it suffices to set $n = (\log p_m + \Omega(\lambda))^2$, and $q = B p_m n^{\log \lambda}$.

Extension to Ring LWR. The above scheme can also be instantiated using the Ring LWR assumption introduced together with LWR in [6]. We implement the

more computationally efficient version with Ring LWR and present experiment data in Section 6.

Construction Based on DCR. In Appendix A, we also construct an exact key-homomorphic masking scheme under the decisional composite residuosity (DCR) assumption.

3.2 Flat Secret Sharing

A threshold secret-sharing scheme with M parties normally has two algorithms $\mathsf{Share}, \mathsf{Recon}$, and is parameterized by privacy and reconstruction thresholds γ, ρ, where $0 < \gamma < \rho < 1$. Running Share on a secret value creates M shares. Running Recon on any subset of more than ρM shares recovers the secret. Any subset of less than γM shares contains no information about the secret.

Secret Sharing in Our Framework. Our framework uses the scheme in an unusual way. In the setup phase, the clients run Share to create shares of their masking keys. In the online phase, the server runs Recon *not* over the key shares, but homomorphically over empty masks created under the key shares. As long as Recon is a linear function, the key-homomorphism property (Definition 2) ensures that running Recon over the masks translates to over the underlying key shares. The two thresholds γ, ρ guarantees the masking keys are hidden when at most γM clients are corrupted, and Recon succeeds when at least ρM clients are online.

This approach creates a technical challenge when the masking scheme has only approximate key-homomorphism (Definition 3). Namely, evaluating Recon homomorphically creates an additive noise, which grows with the magnitude of the coefficients of Recon. The noise then propagates into the aggregation result.

To help remove the noise, each client input is multiplied with a scaling factor Δ, set larger than the noise. To accommodate the factor Δ in the clients inputs, the message modulus of the masking scheme is in turn increased by $\log \Delta$ bits. This overhead motivates us to construct a secret-sharing scheme with small coefficients in Recon, which we call a *flat* secret sharing scheme.

Overview of Our Scheme. Our starting point is the linear secret sharing scheme [20] that has $0, 1$ coefficients. However, using the scheme has a prohibitive overhead: the *total* share size scales polynomially in the population M, namely $\Omega(M^{5.3})$.

A first attempt at reducing the share size is to run the scheme in a small committee, sampled during the setup phase. If client corruption and dropout happen independently to the committee sampling, then the fractions of corruption and dropout in the committee roughly equal the true fractions in the population. This is true in our framework, where the set of corrupted clients, and potential dropout clients are decided statically at the beginning.

That is, we add a Setup algorithm to the scheme, which samples a committee $Q \subseteq [M]$ at random. It can be shown that when the fractions $0 < \gamma < \rho < 1$ has a constant gap, a committee of size $O(\kappa)$ suffices, with a $O(2^{-\kappa})$ statistical error.

Running [20] as a blackbox with a committee of size $O(\kappa)$ reduces the total share size from $O(M^{5.3})$ to $O(\kappa^{5.3})$. But we are able to further improve it to $O(\kappa^2)$, with a $O(\kappa^2)$-size committee, by re-visiting the analysis of [32], and constructing a committee version of [20] in a non-blackbox way. We summarize the syntax of our committee-based scheme for some secret space $\mathcal{M}$ below.

- $\mathsf{Setup}(1^\kappa, M)$ takes as inputs the statistical security parameter κ, and the population size M. It outputs a committee Q of share holders, and public parameters pp.
- $\mathsf{Share}(\mathsf{pp}, s)$ outputs shares $\{s_j\}_{j \in Q}$ computed from $s \in \mathcal{M}$.
- $\mathsf{Recon}(\mathsf{pp}, W, \{s_j\}_{j \in W})$ takes as inputs a set W indicating which shares are received, and the set of shares $\{s_j\}_{j \in W}$. It outputs a recovered secret s^* or $\perp$.

Correctness and Security. Formally, we define the correctness requirements as follows.

Definition 6 (ρ-reconstruction)**.** *Let κ be the statistical security parameter. For all population size $M \in \mathbb{N}$, secret $s \in \mathcal{M}$, and subset $T \subseteq [M]$ with size $|T| > \rho M$ the following holds.*

$$\Pr\left[\begin{array}{c|c} \mathsf{Recon}(\mathsf{pp}, W, \{s_j\}_W) & (Q, \mathsf{pp}) \leftarrow \mathsf{Setup}(1^\kappa, M), \\ = s & W = T \cap Q, \\ & \{s_j\}_Q \leftarrow \mathsf{Share}(\mathsf{pp}, s) \end{array}\right] \geq 1 - \mathrm{negl}(\kappa).$$

The usual security requires that, for any corruption set $C \subseteq [M]$ below the threshold, i.e. $|C| \leq \gamma M$, corrupted shares $\{s_j\}_{Q \cap C}$ contain no information about the secret s.

We need a stronger property (which implies the usual one) to prove security of our framework: given corrupted shares $\{s_j\}_{Q \cap C}$ of 0, there is algorithm Ext that "extents" them to a full set of shares $\{s_j\}_Q$ for any secret s. The shares $\{s_j\}_Q$ distribute statistically close to shares of s. This is analogous to the property that, given a corrupted subset of Shamir's shares, one can interpolate the rest of the shares to any secret s. We formalize this requirement as follows.

Definition 7 (γ-simulation-privacy)**.** *Let κ be the statistical security parameter. There exists an efficient deterministic algorithm Ext such that for all population size $M \in \mathbb{N}$, secret $s \in \mathcal{M}$, and subset $C \subseteq [M]$ with size $|C| < \gamma M$ the following two distributions are statistically close.*
They share the same public parameters $(Q, \mathsf{pp}) \leftarrow \mathsf{Setup}(1^\kappa, M)$.

1. *$\{s_j\}_Q$ is computed normally as $\{s_j\}_Q \leftarrow \mathsf{Share}(\mathsf{pp}, s)$.*
2. *$\{\tilde{s}_j\}_Q = \{s'_j\}_{Q \cap C} \cup \{\tilde{s}_j\}_{Q \cap \overline{C}}$ is computed by*
   $\{s'_j\}_Q \leftarrow \mathsf{Share}(\mathsf{pp}, 0)$ and $\{\tilde{s}_j\}_{Q \cap \overline{C}} = \mathsf{Ext}(\mathsf{pp}, C, \{s'_j\}_{Q \cap C}, s)$.

Flatness. As explained in "Secret Sharing in Our Framework", we require the Recon algorithm to have small coefficients as a linear function over the input shares. This minimizes the noise introduced by evaluating Recon homomorphically over empty masks. A similar situation arises in the security proof of our

framework, where the simulator needs to evaluate Ext (Definition 7) homomorphically over noisy masks. We therefore additionally require Ext to have small coefficients as a linear function over the input shares and the secret. We summarize the above requirements as “flatness”.

Definition 8 (flatness)**.** *Let κ be the statistical security parameter. A flat secret sharing scheme satisfies the following.*

- *The* Recon *algorithm, when not outputting $\perp$, can be written as a linear function over the input shares, with integer coefficients bounded by $O(1)$.*
- *The* Ext *algorithm can be written as a linear function over the input shares and the secret, with integer coefficients bounded by $O(\log \kappa)$.*

Construction Details. We start by recalling the result of [9] and [20], summarized in the following theorem.

Theorem 1 (formula to secret sharing [9],[20])**.** *For secrets over $\mathcal{M} = \mathbb{Z}_q$ for any modulus q or $\mathcal{M} = \mathbb{Z}$, there exists an efficient algorithm that translates any monotone Boolean formula $f : \{0,1\}^M \to \{0,1\}$, over variables $x_1, \ldots, x_M$, of size $d = |f|$, to a pair of secret sharing algorithms $\mathsf{Share}_f, \mathsf{Recon}_f$ satisfy the following:*

- $\mathsf{Share}_f(s)$ *computes d share units, each corresponding to a literal in f. For each share holder $i \in [M]$, its share s_i consists of all units corresponding to x_i.* $\mathsf{Share}_f(s)$ *outputs the shares $\{s_i\}$.*
  If $\mathcal{M} = \mathbb{Z}_q$, each share unit is an element in $\mathbb{Z}_q$. If $\mathcal{M} = \mathbb{Z}$, with secrets bounded by B, each unit is an integer bounded by $B2^\kappa$.
- *For any subset $T \subseteq [M]$, let $\mathbf{a}_T \in \{0,1\}^M$ denote the assignment where $a_i = 1$ iff $i \in T$. For every subset of the shares $\{s_j\}_T$, reconstruction $\mathsf{Recon}(T, \{s_j\}_T)$ succeeds iff $f(\mathbf{a}_T) = 1$.*
  For any subset $\{s_j\}_C$ that fails to reconstruct, there exists a simulation algorithm Ext *defined analogously to Definition 7.*
- *The algorithms $\mathsf{Share}_f, \mathsf{Recon}_f$ satisfy “flatness” per Definition 8.*

With Theorem 1, constructing a flat secret sharing scheme for any access structure reduces to finding a corresponding formula f:

- Setup constructs a formula f as pp, and defines the committee Q as the set of distinct literals in f.
- $\mathsf{Share}, \mathsf{Recon}$ simply run $\mathsf{Share}_f, \mathsf{Recon}_f$ given by Theorem 1.

Below we first describe the result of [32], which shows the existence of a formula f_t, of size $O(M^{5.3})$, for any t-threshold function. (Note that for any $\gamma M < t < \rho M$, f_t satisfies our requirement.)

Construction 2 (t-threshold monotone Boolean formula [32])**.**
In [32], f_t (over M variables) is implicitly constructed through a formulae distribution F_t satisfying the following:

$$\forall \mathbf{a} \in \{0,1\}^M, \quad \Pr\left[f(\mathbf{a}) = \mathsf{Thresh}_t(\mathbf{a}) \mid f \leftarrow F_t\right] > 1 - 2^M, \tag{1}$$

where Thresh_t denotes the t-threshold function. Applying the union bound over all 2^M values for $\mathbf{a}$, we have

$$\Pr\left[\forall \mathbf{a} \in \{0,1\}^M,\ f(\mathbf{a}) = \mathsf{Thresh}_t(\mathbf{a}) \mid f \leftarrow F_t\right] > 0.$$

Hence, there exists a formula f_t in F_t that computes Thresh_t exactly.

Further, note that for any threshold $0 < t < M$, the function Thresh_t over M inputs is equivalent to $\mathsf{Thresh}_{M'/2}$ over $M' = M + D \leq 2M$ inputs, with $D \leq M$ dummy variables always set to 1 or 0, respectively for the case of $t < M/2$ or $t \geq M/2$. For technical reasons, we always choose M' to be odd. Therefore, it remains to construct a formulae distribution $F_{M/2}$, for any *odd* M.

The construction is recursive. In the base case, $F^{(0)}$ is defined as

$$F^{(0)} := \begin{cases} x_j \text{ for a uniform } j \xleftarrow{\$} [M] & \text{w/ prob. } p = 3 - \sqrt{5} \\ 0 & \text{w/ prob. } (1-p). \end{cases}$$

For $i \geq 1$, the formulae distribution F^i is defined inductively

$$F^{(i)} := (F_1^{(i-1)} \vee F_2^{(i-1)}) \wedge (F_3^{(i-1)} \vee F_4^{(i-1)}),$$

where $F_1^{(i-1)}, F_2^{(i-1)}, F_3^{(i-1)}, F_4^{(i-1)}$ are distributions independent and identical to $F^{(i-1)}$. It's shown that after $k = O(1) + 2.65 \log M$ recursion steps, the distribution $F_{M/2} = F^{(k)}$ satisfies Equation 1.

Correctness and Efficiency of Construction 2. According to Equation 1, we examine the probability that, for any assignment $\mathbf{a} \in \{0,1\}^M$, a sample $f^{(i)} \leftarrow F^{(i)}$ computes the *incorrect* result.

- When $\mathbf{a}$ has less than $M/2$ ones, $f^{(i)}(\mathbf{a})$ is supposed to output 0, but instead (incorrectly) outputs 1. Let $p_s^{(i)}$ denote this probability, i.e., $f^{(i)}(\mathbf{a}) = 1$. By construction, we have

$$p_s^{(i)} = \left(1 - (1 - p_s^{(i-1)})^2\right)^2. \quad (2)$$

- When $\mathbf{a}$ has at least $M/2$ ones, let $p_c^{(i)}$ denote the probability that $f^{(i)}(\mathbf{a})$ (incorrectly) outputs 0. Similarly, we have

$$p_c^{(i)} = 1 - (1 - \left(p_c^{(i-1)}\right)^2)^2. \quad (3)$$

By construction of $F^{(0)}$, and that M is odd, we also have

$$p_s^{(0)} < p(\frac{1}{2} - \frac{1}{2M}), \quad p_c^{(0)} \leq (1-p) + p(\frac{1}{2} - \frac{1}{2M}).$$

It remains to show that $p_s^{(k)}, p_c^{(k)} < 2^M$ for $k = O(1) + 2.65 \log M$, which follows from the technical claims below, which are taken directly from [32].

Claim 1 (phase 1)**.** *For the recurrence relations specified by Equation 2, 3 with any initial values satisfying* $p_s^{(0)} < p/2 - p/(2M)$*,* $p_c^{(0)} < 1 - p/2 - p/(2M)$*, it holds that* $p_s^{(k_1)} \leq p/2 - \Omega(1)$*, and* $p_c^{(k_1)} \leq 1 - p/2 - \Omega(1)$ *for* $k_1 = 1.65 \log M$*.*

Claim 2 (phase 2)**.** *For the recurrence relations specified by Equation 2, 3 with any initial values satisfying* $p_s^{(0)} < p/2 - \Omega(1)$*, and* $p_c^{(0)} < 1 - p/2 - \Omega(1)$*, it holds that* $p_s^{(k_2)}, p_c^{(k_2)} < 2^M$ *for* $k_2 = O(1) + \log M$*.*

Intuitively, a formula sampled from $F^{(0)}$ fails with probability close to (but less than) $p/2$ and $1 - p/2$ respectively in the two cases. Each recursive step "shifts" them further away from the starting points towards 0. Claim 1 shows that it takes $k_1 = O(\log M)$ steps to start at $\Theta(1/M)$-away and shift to $\Omega(1)$-away from the starting points. Claim 2 shows that it takes additional $k_2 = O(\log M)$ steps to shift exponentially close to 0.

Since each recursive step multiplies the formula size by 4, after $k = k_1 + k_2 = O(1) + 2.65 \log M$ steps, the formulas in $F^{(k)}$ has size $4^{O(1)+2.65 \log M} = O(M^{5.3})$.

Reducing the Size of Construction 2. Our first observation is instead of the formula f_t, we only need a a formula $f_{\rho,\gamma}$ that 1) computes 1 if the inputs have $> \rho M$ ones, 2) computes 0 if the inputs have $< \gamma M$ ones, and 3) may otherwise compute either. We denote this (ρ, γ)-threshold function $\mathsf{Thresh}_{\rho,\gamma}$. A similar trick reduces computing $\mathsf{Thresh}_{\rho,\gamma}$ over M variables to $\mathsf{Thresh}_{1/2+\delta,1/2-\delta}$ over $M' \leq 2M$ variables for some constant fraction $\delta = (\rho - \gamma)/4$.

This observation allows us to calculate the initial failure probability for $f^{(0)} \leftarrow F^{(0)}$ differently from above.

- When $\mathbf{a}$ has less than $M(1/2 - \delta)$ ones, $f^{(0)}$ fails (i.e., computes 1) with probability $p_s^{(0)} < p(1/2 - \delta) < p/2 - \Omega(1)$.
- When $\mathbf{a}$ has more than $M(1/2 + \delta)$ ones, $f^{(0)}$ fails with probability $p_c^{(0)} < (1 - p) + p(1/2 + \delta) < 1 - p/2 - \Omega(1)$.

Since the initial values of $p_s^{(0)}, p_c^{(0)}$ already satisfies the condition for Claim 2, we indeed only need $k_2 = O(1) + \log M$ recursive steps! This observation already let us reduce the size of the formula from $4^{O(1)+2.56 \log M} = O(M^{5.3})$ to $4^{O(1)+\log M} = O(M^2)$.

Our second observation is that in the static corruption model, the set of corrupted and the reconstructing share holders C, T_i at each iteration i is fixed before the secret sharing Setup algorithm. Therefore, instead of finding an exact formula $f_{\rho,\gamma}$ that's correct on all assignments, it suffices to sample $f \leftarrow F_{\rho,\gamma}$ during Setup that's correct on the (poly(κ) many) fixed assignments $\mathbf{a}_C$ and $\mathbf{a}_{T_i}$.

In particular, we can avoid taking the union bound over 2^M values for $\mathbf{a}$, and only construct a distribution $F_{\rho,\gamma}$ (equivalently, $F_{1/2+\delta,1/2-\delta}$) such that

$$\forall \mathbf{a} \in \{0,1\}^M, \quad \Pr\left[f(\mathbf{a}) = \mathsf{Thresh}_{\rho,\gamma}(\mathbf{a}) \mid f \leftarrow F_{\rho,\gamma}\right] > 1 - 2^{\kappa}.$$

By Claim 2, we now only need $k_2' = O(1) + \log \kappa$ recursive steps, which further reduces the formula size to $O(\kappa^2)$!

To summarize, we obtain the following lemma.

Lemma 2 (flat secret sharing)**.** *For any population size $M \in \mathbb{N}$, constant fractions $0 < \gamma < \rho < 1$, integer modulus q and dimension ℓ, there exists a flat secret-sharing scheme* Setup, Share, Recon *for secrets space $\mathcal{M} = \mathbb{Z}_q^\ell$ or $\mathcal{M} = \mathbb{Z}^\ell$, with privacy and reconstruction thresholds γ, ρ. Furthermore,*

- *It has committee size $|Q| = O(\kappa^2)$, where the constant depends on the thresholds γ, ρ.*
- *The* Recon *algorithm, when written as a linear function, has $O(\kappa)$ non-zero coefficients, which are 1 or -1.*

Concrete Algorithm for Theorem 1. When sharing a secret according to a formula f, Share_f views f as a tree with AND, OR on the intermediate nodes, and literals x_i on the leaf nodes. It assigns a share to each node of this tree: i) Upon reaching an AND node, split the current share s into two additive shares of s, and assign them to the children. ii) Upon reaching an OR node, duplicate s and assign them to the children. iii) Upon reaching a literal x_i, assign s to share holder i. Reconstruction according to f follows a similar recursive algorithm.

4 The LERNA Framework

In this section, we describe our abstract secure aggregation protocol assuming the existence of the two technical tools introduced in Section 3:

- An ε-approximate key-homomorphic masking scheme HM = (HM.Setup, KeyGen, TagGen, Mask, UnMask, Eval) setup properly with HM.pp, specifying a message space $\mathbb{Z}_{p_m}^\ell$, mask space $\mathbb{Z}_q^\ell$ and key space $\mathcal{K}$.
- A flat secret sharing scheme SS = (SS.Setup, Share, Recon) for sharing the masking keys in the above key space $\mathcal{K}$.

The protocol additionally assumes a public key encryption scheme and two hash functions $\mathcal{H}_1, \mathcal{H}_2$ modeled as random oracles. We assume the hash functions $\mathcal{H}_1, \mathcal{H}_2$ output exactly the numbers of random bits required by the algorithms SS.Setup, and TagGen.

The protocol runs with M clients $\{P_i\}$ and a single server S for T iterations. During each iteration $t \in [T]$, every client P_i obtains a fresh integer vector $\mathbf{x} \in \mathbb{Z}^\ell$ from a bounded range $[0, B_x]$. To avoid wrap-around in the aggregation results, we setup the masking scheme with a modulus lower bound $B_{\text{msg}} = \Delta M B_x$, where Δ is a message scaling factor introduced in the protocol.

The protocol is further parameterized by two thresholds $\gamma, \delta \in (0, 1)$, specifying the maximum fractions of corrupted clients and dropout clients, respectively, under the restriction that $\gamma + \delta < 1$. We set the privacy threshold of the secret sharing scheme to γ, and the reconstruction threshold to $\rho = 1 - \delta$.

In the online phase, the protocol uses a noise bound B_e and a message scaling factor Δ, which we specify in Section 4.4 and Appendix A for concrete instantiations under LWR and DCR.

4.1 The Semi-honest Protocol

We start with the simpler, semi-honest variant of the protocol, given in Figure 1, and Figure 2. We prove the correctness and sketch the privacy of the semi-honest protocol in Appendix 4.3. We describe the additional steps to obtain the malicious variant in the next subsection, and defer the more formal (in the UC framework) security proof for the malicious protocol to Appendix C.

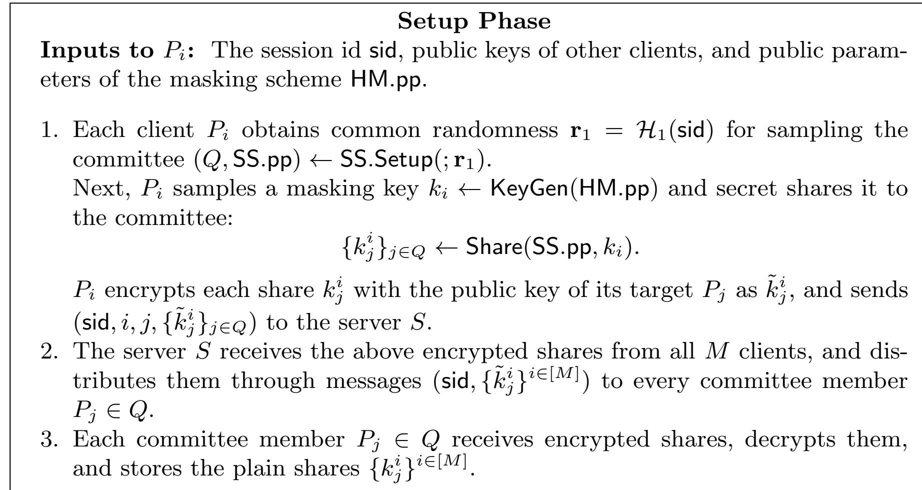

Setup Phase

Inputs to P_i: The session id sid, public keys of other clients, and public parameters of the masking scheme $\mathsf{HM.pp}$.

1. Each client P_i obtains common randomness $\mathbf{r}_1 = \mathcal{H}_1(\mathsf{sid})$ for sampling the committee $(Q, \mathsf{SS.pp}) \leftarrow \mathsf{SS.Setup}(;\mathbf{r}_1)$.
   Next, P_i samples a masking key $k_i \leftarrow \mathsf{KeyGen}(\mathsf{HM.pp})$ and secret shares it to the committee:
   $$\{k_j^i\}_{j\in Q} \leftarrow \mathsf{Share}(\mathsf{SS.pp}, k_i).$$
   P_i encrypts each share k_j^i with the public key of its target P_j as $\tilde{k}_j^i$, and sends $(\mathsf{sid}, i, j, \{\tilde{k}_j^i\}_{j\in Q})$ to the server S.
2. The server S receives the above encrypted shares from all M clients, and distributes them through messages $(\mathsf{sid}, \{\tilde{k}_j^i\}^{i\in[M]})$ to every committee member $P_j \in Q$.
3. Each committee member $P_j \in Q$ receives encrypted shares, decrypts them, and stores the plain shares $\{k_j^i\}^{i\in[M]}$.

Fig. 1: LERNA protocol for the setup phase.

Setup Phase. During the setup phase, the clients first agree on a small committee Q, computed using public common randomness $\mathbf{r}_1$. They each sample a secret masking key k_i, and secret share it to the committee Q, using the server to distribute those shares. To keep the shares secret from the server, the clients encrypt each share using the public key of its target share-holder.

Note that the clients only run the setup phase once, followed by T online phases. In each online phase, each client P_i uses the same masking key k_i to mask its fresh input vector $\mathbf{x}_i$. Reusing the masking key may seem like a privacy concern. To address this, we ensure that in each online phase, the clients sample a *fresh tag* τ used for computing the mask. The randomness of the tag τ protects the input vector $\mathbf{x}_i$, as long as the masking key remains secret.

Online Phase.

Step 1: Every client runs the key-homomorphic masking scheme $\mathsf{HM.Mask}$ to obtain a masked input vector $\mathbf{z}_i$, and sends it to the server S. It's important to note that key-homomorphism only holds for masks computed using the same tag τ. Therefore, the clients sample the tag using public common randomness $\mathbf{r}_2$.

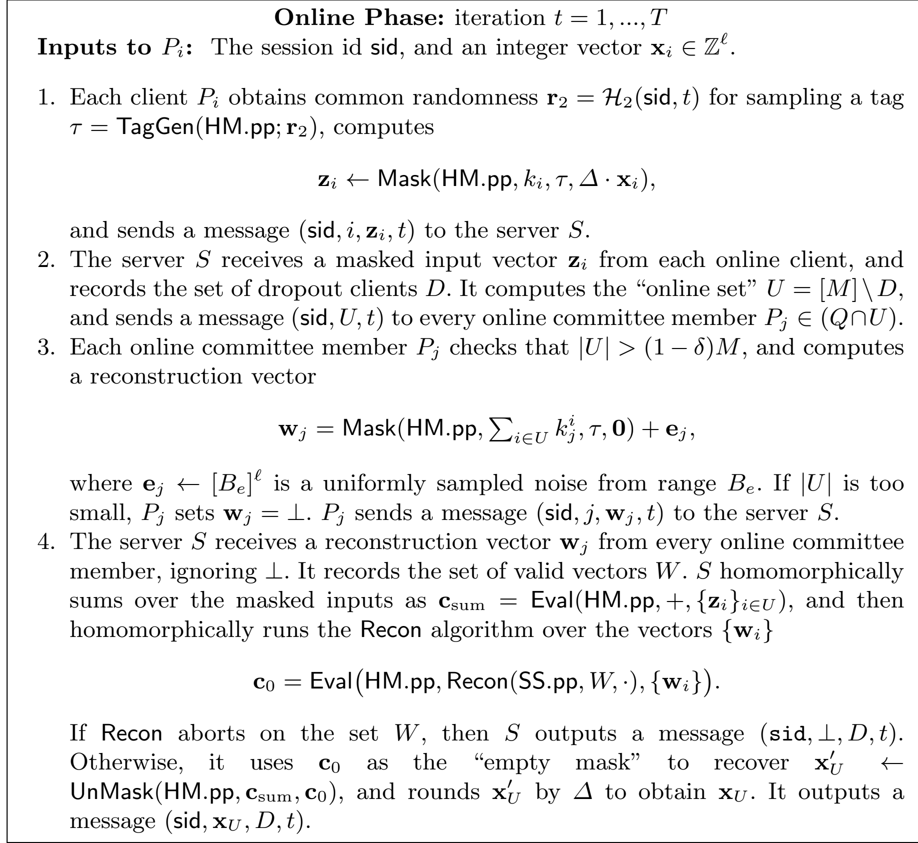

Online Phase: iteration $t = 1, ..., T$

Inputs to P_i: The session id sid, and an integer vector $\mathbf{x}_i \in \mathbb{Z}^\ell$.

1. Each client P_i obtains common randomness $\mathbf{r}_2 = \mathcal{H}_2(\mathsf{sid}, t)$ for sampling a tag $\tau = \mathsf{TagGen}(\mathsf{HM.pp}; \mathbf{r}_2)$, computes
$$\mathbf{z}_i \leftarrow \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \Delta \cdot \mathbf{x}_i),$$
and sends a message $(\mathsf{sid}, i, \mathbf{z}_i, t)$ to the server S.
2. The server S receives a masked input vector $\mathbf{z}_i$ from each online client, and records the set of dropout clients D. It computes the "online set" $U = [M] \setminus D$, and sends a message (sid, U, t) to every online committee member $P_j \in (Q \cap U)$.
3. Each online committee member P_j checks that $|U| > (1 - \delta)M$, and computes a reconstruction vector
$$\mathbf{w}_j = \mathsf{Mask}(\mathsf{HM.pp}, \textstyle\sum_{i \in U} k_j^i, \tau, \mathbf{0}) + \mathbf{e}_j,$$
where $\mathbf{e}_j \leftarrow [B_e]^\ell$ is a uniformly sampled noise from range B_e. If $|U|$ is too small, P_j sets $\mathbf{w}_j = \perp$. P_j sends a message $(\mathsf{sid}, j, \mathbf{w}_j, t)$ to the server S.
4. The server S receives a reconstruction vector $\mathbf{w}_j$ from every online committee member, ignoring $\perp$. It records the set of valid vectors W. S homomorphically sums over the masked inputs as $\mathbf{c}_{\mathrm{sum}} = \mathsf{Eval}(\mathsf{HM.pp}, +, \{\mathbf{z}_i\}_{i \in U})$, and then homomorphically runs the Recon algorithm over the vectors $\{\mathbf{w}_i\}$
$$\mathbf{c}_0 = \mathsf{Eval}\big(\mathsf{HM.pp}, \mathsf{Recon}(\mathsf{SS.pp}, W, \cdot), \{\mathbf{w}_i\}\big).$$
If Recon aborts on the set W, then S outputs a message $(\mathtt{sid}, \perp, D, t)$. Otherwise, it uses $\mathbf{c}_0$ as the "empty mask" to recover $\mathbf{x}'_U \leftarrow \mathsf{UnMask}(\mathsf{HM.pp}, \mathbf{c}_{\mathrm{sum}}, \mathbf{c}_0)$, and rounds $\mathbf{x}'_U$ by Δ to obtain $\mathbf{x}_U$. It outputs a message $(\mathsf{sid}, \mathbf{x}_U, D, t)$.

Fig. 2: LERNA protocol for the online phase (semi-honest).

Step 2: The server receives the masked input vectors $\{\mathbf{z}_i\}$ from the online clients, and replies the online set U to each committee member. Note that non-committee member clients don't need to send anything in the rest of the online phase.

Step 3: Every committee member P_j aggregates locally its shares of masking keys from the online set U to obtain an aggregated key share k_j^U, uses it to compute an "empty mask" as its reconstruction vector $\mathbf{w}_j$, and sends it to the server S.

Step 4: The server S receives reconstruction vectors $\{\mathbf{w}_j\}$ from the online committee members. It proceeds to locally recover the aggregation result.

First, it homomorphically aggregates the masked input vectors $\mathbf{z}_i$ to obtain $\mathbf{c}_{\mathrm{sum}}$. By key-homomorphism, the vector $\mathbf{c}_{\mathrm{sum}}$ approximately equals running $\mathsf{HM.Mask}$ on the scaled aggregation result $\mathbf{x}'_U = \Delta \cdot \sum_U \mathbf{x}_i$ under the key $k_U = \sum_U k_i$. It remains to obtain an "empty mask" $\mathbf{c}_0$ under the same key k_U,

with which the server can recover the scaled aggregation result $\mathbf{x}'_U$, and then the actual aggregation result $\mathbf{x}_U = \lfloor \mathbf{x}'_U / \Delta \rceil$ through rounding.

To obtain the empty mask $\mathbf{c}_0$ under the key k_U, the server homomorphically runs the algorithm $\mathsf{SS.Recon}$ over the reconstruction vectors $\mathbf{w}_j$. By key-homomorphism, the result indeed approximately equals $\mathbf{c}_0$. Note that approximate key-homomorphism causes some errors in the recovered result $\mathbf{x}'_U$. But we set the scaling factor Δ sufficiently large to make sure such errors are removed by the rounding step.

Alternative to the PKI Setup. The setup phase of our protocol requires the clients to encrypt their secret shares under the public keys of the target shareholders. For simplicity, our protocol assumes a public key infrastructure (PKI), and that each client enters the setup phase knowing every other client's public key.

An alternative approach is to let the clients run pairwise key agreement at the beginning of the setup phase, as described in the "Communication Model" paragraph (Section 2).

Committee Members and Non-members. Note that in each online phase, non-member clients only have one task: send masked input vectors to the server. The rest of the reconstruction steps are handled by committee member clients.

This separation of responsibility suggests an alternative aggregation model, where during each phase, only a small, potentially random, subset among the non-member clients is required to provide inputs. Our protocol can be adapted straightforwardly to guarantee: as long as not too many committee members drop out during the session, the server can securely compute the aggregation result. This scenario can be useful for stochastic federated learning algorithms that benefit from a large input population, but only learns from a random subset at each iteration.

4.2 Achieving Malicious Security

To achieve malicious security, we keep the setup phase (Figure 1) unchanged, and only modify the online phase (Figure 2) starting from step 2. The modifications follow similar ideas to prior work [11, 8]. The modified online phase is given in Figure 3, where the changes are highlighted in blue.

To see why we need the additional steps in the malicious setting, consider the following corrupted server. Recall that in the semi-honest online protocol, the server sends an online set U to online committee members to recover an aggregation result $\mathbf{x}_U = \sum_U \mathbf{x}_i$. A corrupted server instead sends different online sets, $U \neq U'$, to two subsets of online committee members. As long as both subsets are large enough, the correctness of the semi-honest protocol guarantees the successful recovery of both results $\mathbf{x}_U$ and $\mathbf{x}_{U'}$ by the server. This obviously violates our security definition, which requires only a single sum of honest inputs is leaked in each online phase.

The additional steps 3 - 4 in Figure 3 roughly ask each client, including corrupted ones, to "vote" on an online set U by signing a hash h_U. The server

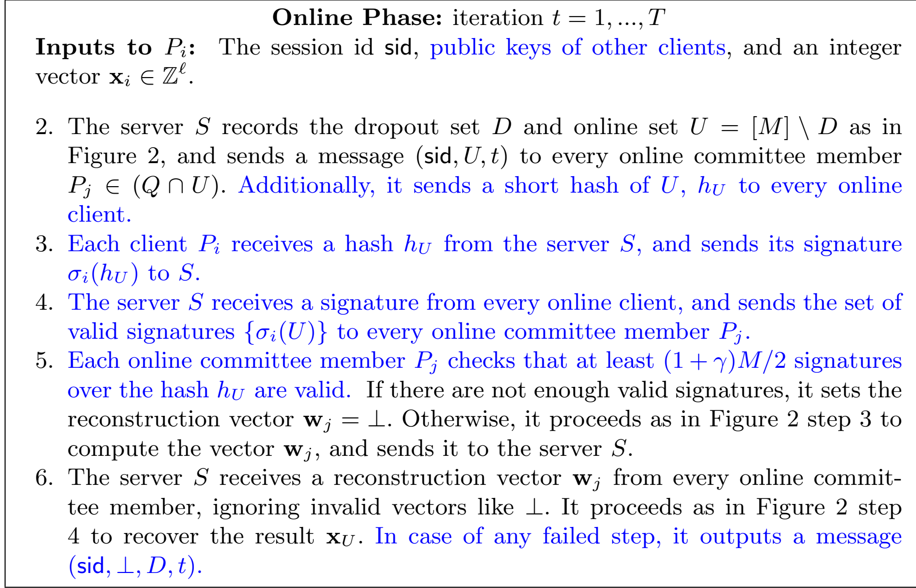

Online Phase: iteration $t = 1, ..., T$

Inputs to P_i: The session id sid, public keys of other clients, and an integer vector $\mathbf{x}_i \in \mathbb{Z}^\ell$.

2. The server S records the dropout set D and online set $U = [M] \setminus D$ as in Figure 2, and sends a message (sid, U, t) to every online committee member $P_j \in (Q \cap U)$. Additionally, it sends a short hash of U, h_U to every online client.
3. Each client P_i receives a hash h_U from the server S, and sends its signature $\sigma_i(h_U)$ to S.
4. The server S receives a signature from every online client, and sends the set of valid signatures $\{\sigma_i(U)\}$ to every online committee member P_j.
5. Each online committee member P_j checks that at least $(1+\gamma)M/2$ signatures over the hash h_U are valid. If there are not enough valid signatures, it sets the reconstruction vector $\mathbf{w}_j = \perp$. Otherwise, it proceeds as in Figure 2 step 3 to compute the vector $\mathbf{w}_j$, and sends it to the server S.
6. The server S receives a reconstruction vector $\mathbf{w}_j$ from every online committee member, ignoring invalid vectors like $\perp$. It proceeds as in Figure 2 step 4 to recover the result $\mathbf{x}_U$. In case of any failed step, it outputs a message $(\mathsf{sid}, \perp, D, t)$.

Fig. 3: LERNA protocol for the online phase (malicious).

collects those signatures as unforgeable votes and sends them to the committee members. The threshold in step 5 is set such that at most one online set U^* can have enough votes. Therefore, the above attack is prevented.

Preventing Abort Attacks. While setting the threshold for valid signatures in Step 5 to $(1+\gamma)M/2$ guarantees that at most one online set U^* has enough votes, it creates an opportunity for malicious clients to abort the protocol, even when the server is honest, by not sending enough valid signatures. To avoid this issue, we need enough honest clients so that their signatures alone are enough for the threshold. Restricting the corruption and dropout threshold γ, δ such that $(3\gamma + 2\delta) < 1$ suffices.

Claim 3. *Assuming $(3\gamma + 2\delta) < 1$, and the server is honest, then every honest committee member always collects at least $(1+\gamma)M/2$ valid signatures in Step 5.*

Proof. By the assumption, there are at least $(1-\gamma-\delta)M$ honest clients in each iteration that remain online, and will send a valid signature in Step 3 on the hash h_U received from an honest server. Calculation shows $(1-\gamma-\delta)M \geq (1-\gamma)M/2$ iff $1 \geq (3\gamma + 2\delta)$. $\square$

By the above claim, an honest server is guaranteed to receive non-$\perp$ reconstruction messages from all honest online committee members in Step 6. By ρ-reconstruction ($\rho = 1 - \delta$) of the secret sharing, the server succeeds in computing the empty mask $\mathbf{c}_0$.

Finally, the server may still abort if $\mathsf{UnMask}(\mathsf{HM.pp}, \mathbf{c}_{\text{sum}}, \mathbf{c}_0)$ fails. However, in our LWR masking scheme (Construction 1), the UnMask algorithm simply computes a subtraction modulo p_m, which always succeeds. We note that this is not true for the DCR masking scheme (Construction 3), where UnMask may fail on maliciously generated input $\widetilde{\mathbf{c}}_0$.

Overhead of the Malicious Protocol. As highlighted in Figure 3, the communication and computation overhead of the malicious variant consists of the server sending valid signatures $\{\sigma_i(U)\}$ in step 4, and each committee member verifying those signatures in step 5, respectively.

For ease of presentation, the variant shown in Figure 3 requires every client to send a signature in step 3. However, it can be shown that at the cost of a $O(2^{-\kappa})$ statistical error in privacy, only committee members need to send signatures. Note that the number of signatures is at most the committee size $|Q| = O(\kappa^2)$, which is independent of the number of clients M, or the input dimension ℓ. Therefore, when M or ℓ is large, sending and checking those signatures incur only negligible communication and computation overheads over the semi-honest variant.

4.3 Correctness and Privacy

To capture the security of LERNA formally, we define a secure aggregation functionality $\mathcal{F}_{\text{SecAgg}}$ in the UC framework (See Figure 8, and Appendix B). We state and prove the UC security of LERNA in Theorem 2 and Appendix C, which captures both correctness and privacy guarantees.

Below, we illustrate correctness by proving Lemma 3 (as a special case of Theorem 2). We then informally argue the privacy of the semi-honest variant, which already contains most of the key ideas.

Correctness.

Lemma 3 (correctness)**.** *If less than δM clients dropout in an online session t, then the server outputs the correct aggregation result with overwhelming probability in the semi-honest setting.*

Proof (sketch). Looking at the reconstruction step (online step 4), we first argue that the aggregated mask $\mathbf{c}_{\text{sum}}$ is distributed close to a mask over the aggregation result. By ε-approximate key homomorphism (Definition 3), we have

$$\|\mathbf{c}_{sum} - \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i \in U} k_i, \tau, \Delta \sum_{i \in U} \mathbf{x}_i)\|_\infty \leq \varepsilon M.$$

For the UnMask algorithm to work correctly, we need to argue the reconstructed mask $\mathbf{c}_0$ is distributed close to an empty mask under the key $\sum_{i \in U} k_i$. To this end, we first argue that the Recon algorithm succeeds over the shares from the set W with overwhelming probability. By assumption, online set U computed by the server at the online step 2 has size $|U| > (1 - \delta)M$. Therefore, all online committee members send reconstruction vectors $\mathbf{w}_j$ at online step 3. Let

the online set at online step 3 be $U' \subseteq U$. The set of valid reconstruction vectors W equals $W = U' \cap Q$. By assumption, we have $|U'| > (1-\delta)M$. Therefore, by $(1-\delta)$-reconstruction, the algorithm Recon indeed succeeds with overwhelming probability.

By flatness (Definition 8), the function $\mathsf{Recon}(\mathsf{SS.pp}, W, \cdot)$ is linear with $O(1)$ coefficients. Therefore, by ε-approximate key homomorphism, we have

$$\begin{aligned}&\|\mathbf{c}_0 - \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i\in U} \underbrace{\mathsf{Recon}(\mathsf{SS.pp}, W, \{k_j^i\}_{j\in W})}_{k_i}, \tau, \mathbf{0})\|_\infty \\ &\le O(\varepsilon B_e |Q|),\end{aligned}$$

where B_e is the bound on the noises $\mathbf{e}_j$ in the vectors $\mathbf{w}_j$.

Finally, we conclude that the UnMask algorithm on masks $\mathbf{c}_{\text{sum}}$ and $\mathbf{c}_0$ returns a noisy result $\mathbf{x}'_U = \Delta \sum_U \mathbf{x}_i + \mathbf{e}$, where the noise has entries bounded by $\|\mathbf{e}\|_\infty = O(\varepsilon(M + B_e|Q|))$. As long as the message scaling factor Δ is sufficiently large $\Delta \ge 2\|\mathbf{e}\|_\infty$, the server indeed recovers the correct result through rounding by Δ. □

(Semi-honest) Privacy. To argue privacy informally, we sketch an efficient simulator $\mathcal{S}$ which emulates the protocol execution with an adversary $\mathcal{A}$, without knowing each client's inputs. We focus on the more interesting case where the server and a subset C of up to γM clients are corrupted, and the simulator $\mathcal{S}$ simulates the remaining set $H = [M] \setminus C$ of honest clients.

The simulator is allowed to query the following leakage function once per online iteration t

$$f(U, t) = \begin{cases} \mathbf{x}_{H\cap U} = \sum_{i\in H\cap U} \mathbf{x}_i & \text{if } |U| > (1-\delta)M \\ \perp & \text{otherwise,} \end{cases}$$

which outputs the sum of honest inputs (at iteration t) over any sufficiently large set U. Intuitively, this shows that the adversary $\mathcal{A}$ doesn't learn any information beyond the above leakage during each iteration. Specifically, the leakage per iteration is a single sum over at least $(1-\delta-\gamma)M$ honest inputs.

The simulator is defined implicitly through the following hybrid experiments from the real protocol execution to the emulated execution.

H1. We briefly summarize the messages between honest clients and the corrupted server, hence equivalently the adversary $\mathcal{A}$.
During setup, an honest client P_i shares its masking key k_i to the committee Q as $\{k_j^i\}_{j\in Q} \leftarrow \mathsf{Share}(\mathsf{SS.pp}, k_i)$, and sends the encrypted shares to $\mathcal{A}$. By the security of encryption, effectively only the shares directed to corrupted clients $\{k_{j'}^i\}_{j'\in C\cap Q}$ are seen by the $\mathcal{A}$. An honest committee member P_j further receives shares from $\mathcal{A}$ as $\{k_j^{i'}\}_{i\in C}$.
During each online phase, an honest client P_i sends its masked input vector $\mathbf{z}_i \leftarrow \mathsf{Mask}(\mathsf{pp}, k_i, \tau, \Delta\mathbf{x}_i)$ to $\mathcal{A}$. An honest committee member P_j further receives an online set U and sends a reconstruction vector $\mathbf{w}_i \leftarrow \mathsf{Mask}(\mathsf{pp}, \sum_{i\in U} k_j^i, \tau, \mathbf{0})$ to $\mathcal{A}$.

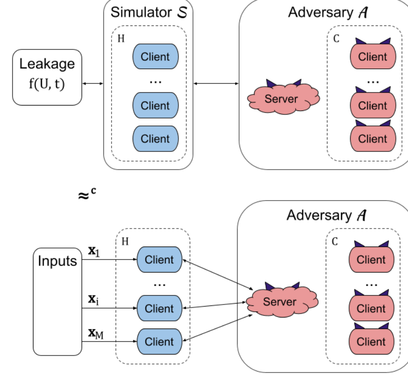

Fig. 4: Illustration of the simulator.

H2. In this hybrid, we apply γ-simulation privacy (Definition 7) of the secret sharing scheme to compute the key shares of k_i separately as corrupted shares and honest shares, using the Ext algorithm:
$\{\tilde{k}^i_j\}_{j\in C\cap Q} \leftarrow \mathsf{Share}(\mathsf{SS.pp}, 0)$,
$\{\tilde{k}^i_j\}_{j\in H\cap Q} = \mathsf{Ext}(\mathsf{SS.pp}, C, \{\tilde{k}^i_j\}_{j\in C\cap Q}, k_i)$.
Flatness (Definition 8), further says the Ext algorithm can be written as a linear function with coefficients bounded by $O(\log\kappa)$, where κ is the statistical security parameter. Therefore, for every $i \in H$ and $j \in H \cap Q$, we can write

$$\tilde{k}^i_j = c_j k_i + \mathsf{Ext}'_j(\{\tilde{k}^i_j\}_{j\in C\cap Q}), \tag{4}$$

where c_j is an integer bounded by $O(\log\kappa)$.

H3. In this hybrid, we apply approximate key-homomorphism (Definition 3) to simulate the reconstruction vectors as follows

$$\begin{aligned}\mathbf{w}_j &\approx \mathsf{Mask}(\mathsf{pp}, \textstyle\sum_{i\in H\cap U} \tilde{k}^i_j, \sum_{i'\in C\cap U} k^{i'}_j + \tau, \mathbf{0}) + \mathbf{e}_j \\ &\approx \mathsf{Mask}(\mathsf{pp}, \textstyle\sum_{i\in H\cap U} \tilde{k}^i_j, \tau, \mathbf{0}) \quad // \text{ first term (FT)} \\ &\quad + \mathsf{Mask}(\mathsf{pp}, \textstyle\sum_{i'\in C\cap U} k^{i'}_j, \tau, \mathbf{0}) + \mathbf{e}_j.\end{aligned}$$

We further simulate the first term (FT) using approximate key-homomorphism and Equation. (4) as

$$\begin{aligned}&\mathrm{FT} + \mathbf{e}_j \\ &\approx c_j \textstyle\sum_{i\in H\cap U} \mathsf{Mask}(\mathsf{pp}, k_i, \tau, \mathbf{0}) \quad // \text{ sum term (ST)} \\ &\quad + \mathsf{Mask}(\mathsf{pp}, \textstyle\sum_{i\in H\cap U} \mathsf{Ext}'_j(\{\tilde{k}^i_j\}_{C\cap Q}), \tau, \mathbf{0}) + \mathbf{e}_j.\end{aligned}$$

Finally, we simulate the sum term (ST) using approximate key-homomorphism as

$$\begin{aligned}&\mathrm{ST} + \mathbf{e}_j \\ &\approx c_j\big(\textstyle\sum_{H\cap U} \mathbf{z}_i - \mathsf{Mask}(\mathsf{pp}, 0, \tau, \Delta \overbrace{\textstyle\sum_{H\cap U} \mathbf{x}_i}^{\mathbf{x}_{H\cap U}})\big) + \mathbf{e}_j.\end{aligned} \tag{5}$$

As long as the smudging noise $\mathbf{e}_j$ is sampled from a sufficiently large range $B_e = O(\varepsilon \cdot \log \kappa \cdot M \cdot 2^\kappa)$, this hybrid is statistically close to the previous one.

H4. In this hybrid, we apply the security (Definition 4) of the masking scheme to simulate the masked input vectors as $\tilde{\mathbf{z}}_i \leftarrow \mathsf{Mask}(\mathsf{pp}, k_i, \tau, \mathbf{0})$. This hybrid is exactly how the simulator $\mathcal{S}$ interacts with $\mathcal{A}$. The only outside information needed by $\mathcal{S}$ is the sum $\mathbf{x}_{H\cap U}$, used for simulating the sum term (ST) (Equation 5). This can be obtained through its one access to the leakage function f at each iteration.

4.4 Instantiation Under LWR

Concretely, we instantiate the LERNA protocol with the LWR-based 1-approximate homomorphic masking scheme in Construction 1.

We set the noise bound $B_e = O(\log(\kappa) M 2^\kappa)$ as required by the security proofs in Section 4.3 and Appendix C, where κ is the statistical security parameter. We set the message scaling factor $\varDelta = O(M + B_e|Q|)$ as required by the correctness proof of Lemma 3, where $|Q| = O(\kappa^2)$ is the committee size of the flat secret sharing scheme, as described in Section 3.2. Under these settings, our protocol sets up the LWR-based masking scheme with message modulus (which is the same as the mask modulus) $p_m = \varDelta M B_x$, which has bit length $\log p_m < O(1) + 3\log\kappa + \kappa + 2\log M + \log B_x$.

The LWR-based masking scheme has keyspace $\mathcal{K} = \mathbb{Z}_q^n$, where the dimension n and modulus q is chosen such that LWR_{n,q,p_m} is assumed to be hard. We therefore instantiate a flat secret-sharing scheme with secret space $\mathcal{M} = \mathcal{K} = \mathbb{Z}_q^n$.

We present communication and computation efficiency analysis for the LWR instantiation in Appendix 5, and we summarize the comparisons with [8] in Table 1 and Table 2.

5 Efficiency Analysis

In this section, we analyze the communication and computation efficiency of LERNA, running with M clients and input vectors from $[0, B_x]^\ell$ for every aggregation session. We use the instantiation under LWR, with dimension n and moduli q, p_m set in Section 4.4, as a concrete example. We compare with the state-of-art secure aggregation protocol [8] in the end.

Communication. Since the server communication equals the sum of the client communications, we focus on the client side. In the following, We differentiate the committee members in $Q \subset [M]$, where $|Q| = O(\kappa^2)$ and the non-members. We start with the semi-honest variant.

Non-member client.

– Setup: It sends $|Q| = O(\kappa^2)$ encrypted shares of its masking key, where each share has the same size as the masking key itself. In the LWR instantiation, this takes $O(\kappa^2 n \log q)$ bits.

– Online: It sends a masked input vector, which takes $O(\ell \log p_m)$ bits, where $\log p_m = O(\kappa + \log M + \log B_x)$.

<u>*Committee member client.*</u>

– Setup: It receives M encrypted shares, which takes $O(Mn \log q)$ bits.
– Online: It receives an online set U (M bits), and sends an empty mask, which takes $O(\ell \log p_m)$ bits.

In the malicious variant, during each online phase each client additionally receives a short hash from the server, and sends a signature back. We count both messages as $O(\lambda)$ bits, where λ is the computational security parameter. A committee member additionally receives $O(M)$ valid signatures. Assuming the signature scheme allows aggregation, we count the aggregated signature also as $O(\lambda)$ bits. We summarize the above in Table 1.

Phase	Setup	Online
Non-member	$O(\kappa^2 n \log q)$	$O(\ell(\kappa + \log M + \log B_x) + \lambda)$
Member	$O((\kappa^2 + M)n \log q)$	$O(\ell(\kappa + \log M + \log B_x) + \lambda + M)$
[8] client	$O(\log M + \kappa)$	$O(\ell(\log M + \log B_x) + (\log M + \kappa)\lambda)$

Table 1: Client communication (bits) of LERNA and [8]

Computation. In the following asymptotic analysis, we count addition, multiplication, and rounding of ring elements as $\widetilde{O}(1)$ for simplicity. We benchmark the concrete computation efficiency of our prototype implementation in Section 6. We again start with the semi-honest variant.

<u>*Non-member client.*</u>

– Setup: It first compute $|Q| = O(\kappa^2)$ shares of its masking key. In our scheme, computing each share takes $O(\log \kappa) = \widetilde{O}(1)$ additions in the secret space. Therefore, in total it takes $\widetilde{O}(\kappa^2 n)$ time in the LWR instantiation. It then encrypts each share, which in total takes $\widetilde{O}(\kappa^2 n)$ time.
– Online: It computes an ℓ dimension masked input vector. Under the more computationally efficient [4] Ring LWR assumption, our scheme takes $O(\ell \log n) = \widetilde{O}(\ell)$ time to compute such a mask.

<u>*Committee member client.*</u>

– Setup: It decrypts M received shares, which takes $O(Mn)$ time.
– Online: It aggregates received shares over an online set U, where $|U| < M$, which takes $O(Mn)$ time. It then computes an empty ℓ dimension mask, which takes $\widetilde{O}(\ell)$ time.

[4] using NTT for polynomial multiplication

Server.

– Setup: It has no significant computation except forwarding messages from the clients to the committee members.
– Online: It first homomorphically adds the received dimension masked input vectors, which takes $O(M\ell)$ in $\mathbb{Z}_{p_m}$. It next homomorphically computes Recon over the received empty masks. In our scheme, Recon can be written as a linear function, with $O(\kappa)$ non-zero coefficients, where the coefficients take $O(\kappa^2)$ time to find. In total, running Recon over the masks takes $\widetilde{O}(\kappa\ell + \kappa^2)$ time.

In the malicious variant, we ignore the time for the server to compute a hash of the online set U, as well as the time for each client to sign the hash. The server additionally aggregates upto M signatures, and each committee member checks their validity. We count both as $\widetilde{O}(M)$ time. We summarize the above in Table 2.

Phase	Setup	Online
Non-member	$\widetilde{O}(\kappa^2 n)$	$\widetilde{O}(\ell)$
Member	$\widetilde{O}((\kappa^2 + M)n)$	$\widetilde{O}(\ell + Mn)$
[8] client	$\widetilde{O}(1)$	$\widetilde{O}(\kappa^2 + \kappa\ell)$
LERNA Sever	$\widetilde{O}(1)$	$\widetilde{O}((\kappa + M)\ell + \kappa^2)$
[8] server	$\widetilde{O}(\kappa M)$	$\widetilde{O}(\kappa M\ell + \kappa^2 M)$

Table 2: Computation time of LERNA and [8]

Comparison with [8]. For comparison, we consider the protocol of [8], adapted to our setting, multiple online aggregation sessions with a PKI setup. Its high-level (semi-honest) structure consists of one setup round and three online rounds.

– Setup: The server generates a communication graph where each client has $k = O(\log M + \kappa)$ neighbors, and sends to each client P_i its list of neighbors N_i.
– Online:
  1. Each client P_i generates a key-agreement pair $\mathsf{sk}_i, \mathsf{pk}_i$ and a PRG seed b_i. It secret shares the secret key and the seed sk_i, b_i to its neighbors using Shamir, and distributes encrypted shares, as well as its public key pk_i through the server.
  2. Each client P_i runs key-agreement with its neighbors to obtain pairwise shared secrets $\{k_{ij}\}_{j \in N_i}$. It masks its input vector by expanding the pairwise secrets k_{ij} and its own seed b_i using the PRG. It sends the masked inputs to the server, which replies to the online set among its neighbors.
  3. Each client P_i collects shares of b_i for its online neighbors and shares of sk_i for its dropout neighbors. It sends those shares to the server, which recovers the aggregation result by re-computing various masks through the PRG.

In the malicious variant, the protocol lets each client sample its own neighbors during setup, and adds an extra "signature check" round similar to Figure 3. W Note that in both the semi-honest and the malicious variants, [8] takes one more round than LERNA *per iteration* and achieves a weaker privacy guarantee, as explained in Section 2.

We reproduce its communication and computation efficiencies reported in [8] in Table 1 and Table 2.

6 Experimental Evaluation

We benchmark the concrete efficiency of the LERNA framework by implementing the semi-honest protocol instantiated under the (Ring) LWR assumption (cf. Section 4.4 for a description).

As our baseline, we compare our protocol design with the semi-honest protocol from [8], adapted naturally to the multi-session setting following the description in Section 5. In particular, the baseline server uses the setup phase to randomly sample a communication graph, and inform each client of its set of neighbors. Baseline clients re-use the same communication graph throughout the following online phases.

Our benchmarks clearly highlight the lightweight server computation during each online phase.

6.1 Implementation Details

Our prototypes are implemented in Python. The protocol simulations are run locally, using the ABIDES simulation framework [14]. Our implementations use the following libraries for heavy computations:

- SEAL [30] and PySEAL [5] for polynomial arithmetics required by Ring LWR.
- Gmpy2 [6] for large integer arithmetics.
- M2Crypto [7] as an interface to AES for implementing a PRG and a random oracle.
- PyNaCl [8] for public key encryption and key-agreement.

Setting Parameters. In the LERNA framework, we need to set two security parameters, $\lambda = 128, \kappa = 40$. Computationally secure primitives (e.g., encryption and the masking scheme) are set to have $\lambda = 128$ bits of security, and statistically secure primitives (e.g., the flat secret sharing scheme) are set to have $\kappa = 40$ bits of security. The concrete committee size equals $|Q| = 2^{14} = 16384$ for $\kappa = 40$.

In our prototype, the message modulus p_m for the (Ring) LWR based key-homomorphic masking scheme is set as described in Section 4.4, which ranges

[5] https://github.com/Lab41/PySEAL
[6] https://gmpy2.readthedocs.io/
[7] https://m2crypto.readthedocs.io/
[8] https://pynacl.readthedocs.io/

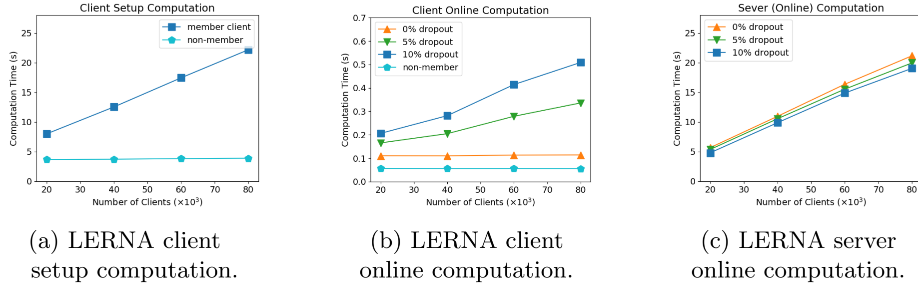

(a) LERNA client setup computation.

(b) LERNA client online computation.

(c) LERNA server online computation.

Fig. 5: LERNA computation time vs. number of clients (M), with fixed input dimension $\ell = 10K$.

from 142 to 145 bits in our benchmark settings. We set the RLWR dimension to be 2^{11}, and the modulus to $q = p_m \cdot 2^{254}$ to guarantee at least 128 bits of security, according to the hardness estimator [9] of [2].

In the baseline prototype, we set the field size for Shamir's secret sharing to be a 257 bit prime, because the secrets are 256 bit curves used in key-agreement. To set the neighborhood size k and privacy threshold t of Shamir's secret sharing, we follow Theorem 3.10 in [8] (section 3.5). In our settings where the number of parties ranges from $M = 400, \ldots, 80K$, the neighborhood size ranges from $k = 109, \ldots, 126$, and the privacy threshold ranges from $t = 55, \ldots, 63$ to achieve $2^{-\kappa} = 2^{-40}$ statistical error.

6.2 Benchmarks

Our benchmarks are run on a desktop machine with 32 Gigabyte of memory and with a single core CPU speed 3.9GHz. Our prototype implementations do not take advantage of multiple cores. For computation time measurements, we report an average over 10 experiment runs.

Computation Efficiency. We first benchmark the computation time of our LERNA prototype with increasing numbers of clients $M = 20K \ldots, 80K$. We run the prototype with $\ell = 10K$ dimension inputs vector with random entries from $[0, 2^{64}]$, and fix the corruption threshold at $\gamma = 10\%$. In Figure 5a, 5b, and 5c, we respectively plot our client runtime during the setup and the online phases, and our server runtime during the online phase. Comparing Figure 5a and 5b, we observe that the setup phase is much heavier compared to the online phases.

In Figure 5b, and 5c, we observe that the dropout rate affects the computation time of both committee member clients and the server. This is because our committee member needs to aggregate masking key shares over the dropout set, which becomes larger both under higher dropout rates and with a larger number of clients. Our server similarly aggregates masked input vectors over the online set, which becomes smaller under higher dropout rates.

[9] running code provided at `https://lwe-estimator.readthedocs.io`

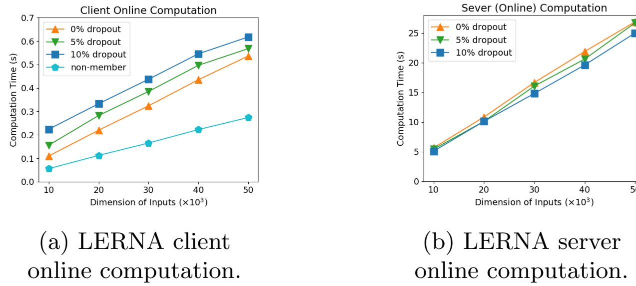

(a) LERNA client online computation.

(b) LERNA server online computation.

Fig. 6: LERNA computation time vs. input dimension (ℓ), with fixed number of clients $M = 20K$. The plot for client setup computation is omitted, as it doesn't depend on ℓ.

In Appendix D, we give more detailed numbers about the running time of different components of our protocols.

We next benchmark the computation time of our protocol with increasing input dimensions $\ell = 10K \ldots, 50K$. We run the prototype with $M = 20K$ clients, and fix the corruption threshold again at $\gamma = 10\%$. In Figure 6a, and Figure 6b, we respectively plot our client and server during the online phase. Since the clients and the server during the setup phase are independent of input dimensions, we omit their plots.

Communication Efficiency. In Table 3 we report the communication sizes of our client with increasing input dimensions $\ell = 10K \ldots, 50K$. The server communication can be deduced as the sum of all clients. Hence we omit its table. We run the prototype with $M = 20K$ clients, and input entries from $[0, 2^{64}]$. We fix the corruption threshold and the dropout rate both at 10%.

Phase	$\ell = 10K$	$\ell = 30K$	$\ell = 50K$
Non-member setup	2.00 (GB)	2.00 (GB)	2.00 (GB)
Member setup	4.44 (GB)	4.44 (GB)	4.44 (GB)
Non-member online	0.18 (MB)	0.54 (MB)	0.91 (MB)
Member online	0.37 (MB)	1.09 (MB)	1.82 (MB)

Table 3: Client communication sizes.

The total offline communication of our clients is indeed heavy, as reported in Table 3. Each client sends encrypted shares of its masking key to the server. Due to the large Ring LWR dimension (2048) and modulus (∼400 bits), this phase requires large communication (2 GB) from each client. Each committee member additionally receives the encrypted shares from the clients.

Thankfully, the entire offline phase doesn't need to be synchronized, which eases the bandwidth requirement. If needed, each client can send a share of its masking key to a committee member one at a time.

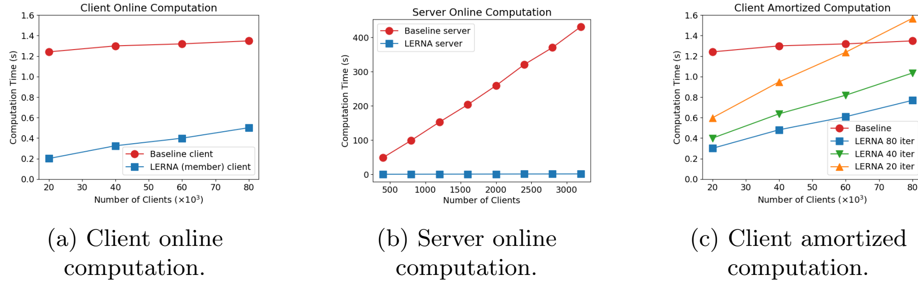

(a) Client online computation.

(b) Server online computation.

(c) Client amortized computation.

Fig. 7: Computation time comparison between LERNA and [8], with fixed input dimension $\ell = 10K$, and dropout rate $\gamma = 10\%$. 7b compares the server computation at a smaller number of clients $M = 400, \ldots, 3200$ due to the high cost of the baseline server. 7c compares the amortized computation time of a single setup phase plus 20/40/80 online phases. Since the baseline client has negligible computation during setup, its amortized time equals that shown in 7a.

Comparing with the Baseline. To compare with the baseline, we run both prototypes with $10K$ dimension inputs vectors with random entries from $[0, 2^{64}]$. We fix the corruption rate and the dropout rate at $\gamma = 10\%$.

As discussed in the introduction, we assume a statically corrupted set of clients throughout the repeated T sessions. A larger T, means a stronger assumption on the staticness and the fraction of corruption. On the flip side, since our protocol enjoys a re-usable setup across T sessions, a larger T gives better efficiency. In comparing with the baseline, we not only compare the computation time of each online iteration (Figure 7a, 7b), but also the amortized time over different settings of T (Figure 7c). The client computation and communication cost of running our setup phase (, where a fresh committee is formed and secret masking keys are shared,) are shown in Figure 5a and Table 3. The server costs of setup for our server and for the baseline solution are negligible. Hence we omit reporting them here. From Figure 7a, we observe that even our slower committee member client runs faster than the baseline during each online iteration for $M = 20K$ to $M = 80K$. As expected, its running time grows faster with M than the baseline because our committee member needs to aggregate masking key shares over the dropout set. If the dropout rate is a non-zero constant, as set in our experiment, then the committee client's work grows linearly in M. In comparison, the computation of the baseline depends linearly in its neighborhood size in the communication graph, which is $O(\log M)$.

In Figure 7c, we compare the clients' amortized running time (showing the heavy member clients for LERNA) of a single setup phase followed by $T = 20/40/80$ online iterations. Since the baseline client has negligible computation during setup, its amortized time equals its online computation time, which doesn't change with T. We observe an advantage, even for the member clients, over the baseline when amortized over more than $T = 40$ online sessions. For example, at $M = 80K$, the total client computation time of 40 LERNA iterations

equals $22 + 0.5 \cdot 40 = 42(s)$, according to Table 2. The total time of 40 baseline iterations is at least $1.2 \cdot 40 = 48(s)$, according to the plot.

In Figure 7b we are only able to compare the server's performance at moderate numbers of clients $M = 400 \ldots 3200$, because the baseline server runs too long when M reaches $10K$. But this is enough to illustrate LERNA's advantage (concretely, more than $100\times$) in server computation times.

Comparing with SASH+[24]. As mentioned in "Related Work", the protocol SASH+ from [24] reduces aggregating ℓ-dimension inputs to aggregating n-dimensional homomorphic PRG seeds, where n is the LWR dimension. SASH+ then runs [8] for the latter. Asymptotically, SASH+ reduces the computation cost of [24] from $\widetilde{O}(\kappa^2+\kappa\ell)$ to $\tilde{O}(\kappa^2+\kappa n+\ell)$ for the clients, and from $\widetilde{O}(\kappa M\ell+\kappa^2 M)$ to $\widetilde{O}(\kappa Mn+\kappa^2 M+\ell M)$ for the server. (See Appendix 5 for details on the formulas, and a comparison with LERNA's asymptotic efficiency.) We optimistically estimate that SASH+ reduces the computation cost of [8] by a factor of (ℓ/n).

In our benchmarks, $\ell = 10K$, and the LWR dimension $n = 2048$. We estimate the server and client computational costs of SASH+ to be 5x smaller than [8] (in reality, the improvement is smaller due to other computation steps that remain constant). Under this estimation, we observe that the LERNA server (Figure 7b) and non-committee member clients (Figure 5b) still significantly outperforms the SASH+ server and SASH+ clients. However, the cost of a LERNA committee member (Figure 7a) becomes comparable to (when M is relatively small e.g. $20K$) or slower than (when M is larger) a SASH+ client.

Acknowledgement. This paper was prepared in part for information purposes by the Artificial Intelligence Research group and AlgoCRYPT CoE of JPMorgan Chase & Co and its affiliates ("JP Morgan"), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful. 2023 JP Morgan Chase & Co. All rights reserved.

Hanjun Li was supported by a NSF grant CNS-2026774 and a Cisco Research Award.

Huijia Lin was supported by NSF grants CNS-1936825 (CAREER), CNS-2026774, a JP Morgan AI Research Award, a Cisco Research Award, and a Simons Collaboration on the Theory of Algorithmic Fairness.

Stefano Tessaro was supported in part by NSF grants CNS-2026774, CNS-2154174, a JP Morgan Faculty Award, a CISCO Faculty Award, and a gift from Microsoft.

References

1. Damiano Abram, Ivan Damgård, Claudio Orlandi, and Peter Scholl. An algebraic framework for silent preprocessing with trustless setup and active security. In Yevgeniy Dodis and Thomas Shrimpton, editors, *CRYPTO 2022, Part IV*, volume 13510 of *LNCS*, pages 421–452. Springer, Heidelberg, August 2022.
2. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046, 2015. `https://eprint.iacr.org/2015/046`.
3. Apple and Google. Exposure notification privacy-preserving analytics (ENPA), 2021.
4. Muhammad Rizwan Asghar, György Dán, Daniele Miorandi, and Imrich Chlamtac. Smart meter data privacy: A survey. *IEEE Commun. Surv. Tutorials*, 19(4):2820–2835, 2017.
5. Marshall Ball, Alper Çakan, and Tal Malkin. Linear threshold secret-sharing with binary reconstruction. In Stefano Tessaro, editor, *2nd Conference on Information-Theoretic Cryptography, ITC 2021, July 23-26, 2021, Virtual Conference*, volume 199 of *LIPIcs*, pages 12:1–12:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
6. Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 719–737. Springer, Heidelberg, April 2012.
7. James Bell, Adrià Gascón, Tancrède Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and Cathie Yun. Acorn: Input validation for secure aggregation. Cryptology ePrint Archive, Paper 2022/1461, 2022. `https://eprint.iacr.org/2022/1461`.
8. James Henry Bell, Kallista A. Bonawitz, Adrià Gascón, Tancrède Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly)logarithmic overhead. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020*, pages 1253–1269. ACM Press, November 2020.
9. Josh Cohen Benaloh and Jerry Leichter. Generalized secret sharing and monotone functions. In Shafi Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 27–35. Springer, Heidelberg, August 1990.
10. Kallista A. Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloé Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards federated learning at scale: System design. In Ameet Talwalkar, Virginia Smith, and Matei Zaharia, editors, *Proceedings of Machine Learning and Systems 2019, MLSys 2019, Stanford, CA, USA, March 31 - April 2, 2019*. mlsys.org, 2019.
11. 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. ACM Press, October / November 2017.
12. Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 410–428. Springer, Heidelberg, August 2013.

13. Lennart Braun, Ivan Damgård, and Claudio Orlandi. Secure multiparty computation from threshold encryption based on class groups. *IACR Cryptol. ePrint Arch.*, page 1437, 2022.
14. David Byrd, Maria Hybinette, and Tucker Hybinette Balch. ABIDES: towards high-fidelity multi-agent market simulation. In *Proceedings of the 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, SIGSIM-PADS 2020, Miami, FL, USA, June 15-17, 2020*, pages 11–22, 2020.
15. Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
16. Henry Corrigan-Gibbs. Privacy-preserving firefox telemetry with prio, 2020.
17. Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In Aditya Akella and Jon Howell, editors, *14th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2017, Boston, MA, USA, March 27-29, 2017*, pages 259–282. USENIX Association, 2017.
18. Ronald Cramer and Victor Shoup. Universal hash proofs and a paradigm for adaptive chosen ciphertext secure public-key encryption. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 45–64. Springer, Heidelberg, April / May 2002.
19. Ivan Damgård and Mats Jurik. A generalisation, a simplification and some applications of Paillier's probabilistic public-key system. In Kwangjo Kim, editor, *PKC 2001*, volume 1992 of *LNCS*, pages 119–136. Springer, Heidelberg, February 2001.
20. Ivan Damgård and Rune Thorbek. Linear integer secret sharing and distributed exponentiation. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006*, volume 3958 of *LNCS*, pages 75–90. Springer, Heidelberg, April 2006.
21. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 640–658. Springer, Heidelberg, May 2014.
22. Yue Guo, Antigoni Polychroniadou, Elaine Shi, David Byrd, and Tucker Balch. Microfedml: Privacy preserving federated learning for small weights. Cryptology ePrint Archive, Paper 2022/714, 2022. `https://eprint.iacr.org/2022/714`.
23. Brett Hemenway and Rafail Ostrovsky. Extended-DDH and lossy trapdoor functions. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 627–643. Springer, Heidelberg, May 2012.
24. Zizhen Liu, Si Chen, Jing Ye, Junfeng Fan, Huawei Li, and Xiaowei Li. SASH: efficient secure aggregation based on SHPRG for federated learning. In James Cussens and Kun Zhang, editors, *Uncertainty in Artificial Intelligence, Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022, Eindhoven, The Netherlands*, volume 180 of *Proceedings of Machine Learning Research*, pages 1243–1252. PMLR, 2022.
25. Yiping Ma, Jess Woods, Sebastian Angel, Antigoni Polychroniadou, and Tal Rabin. Flamingo: Multi-round single-server secure aggregation with applications to private federated learning. In *44th IEEE Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, May 21-25, 2023*, pages 477–496. IEEE, 2023.
26. Suat Özdemir and Yang Xiao. Secure data aggregation in wireless sensor networks: A comprehensive overview. *Comput. Networks*, 53(12):2022–2037, 2009.

27. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In Jacques Stern, editor, *EUROCRYPT'99*, volume 1592 of *LNCS*, pages 223–238. Springer, Heidelberg, May 1999.
28. Christopher Patton, Richard Barnes, and Phillipp Schoppmann. Verifiable Distributed Aggregation Functions. Internet-Draft draft-patton-cfrg-vdaf-01, Internet Engineering Task Force, March 2022. Work in Progress.
29. Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus H. Maier-Hein, Sébastien Ourselin, Micah J. Sheller, Ronald M. Summers, Andrew Trask, Daguang Xu, Maximilian Baust, and M. Jorge Cardoso. The future of digital health with federated learning. *CoRR*, abs/2003.08119, 2020.
30. Microsoft SEAL (release 4.0). `https://github.com/Microsoft/SEAL`, March 2022. Microsoft Research, Redmond, WA.
31. Jinhyun So, Ramy E Ali, Basak Guler, Jiantao Jiao, and Salman Avestimehr. Securing secure aggregation: Mitigating multi-round privacy leakage in federated learning. *arXiv preprint arXiv:2106.03328*, 2021.
32. Leslie G. Valiant. Short monotone formulae for the majority function. *J. Algorithms*, 5(3):363–366, 1984.

A LERNA Under DCR

A.1 Key Homomorphic Masking based on DCR

In this section, we construct an (exact) key-homomorphic masking scheme based on the decisional composite residuosity assumption [27, 19], which considers a modulus that is a product of two distinct primes, $N = p \cdot q$. If p, q are safe primes, i.e., $p = 2p' + 1, q = 2q' + 1$ for some primes p', q', we say the modulus N is *admissible*. We recall some results about the group $\mathbb{Z}^*_{N^{r+1}}$ below.

Lemma 4 ([19])**.** *For all $\lambda, r \in \mathbb{N}$, for all $N = pq$, a product of two distinct λ bit safe primes, and for all messages $m \in \mathbb{Z}_{N^r}$, there exists an efficient algorithm (i.e. polynomial time in λ and r)* Dec *such that*

$$\Pr\left[m = \mathsf{Dec}(h_m, N, r) \,\middle|\, h_m = (1+N)^m \mod N^{r+1}\right] = 1$$

Lemma 5 ([18, 19])**.** *Let λ be the security parameter. Consider r to be any polynomial of λ, and any $N = pq$, a product of two distinct λ bit safe primes. Let $X = \{a | a \in \mathbb{Z}^*_{N^{r+1}}, (a|N) = 1\}$ be the subgroup of $\mathbb{Z}^*_{N^{r+1}}$ with jacobi symbol 1 modulo N. Let $L = \{a^{N^r} | a \in X\}$ be the subgroup of $(N^r)^{th}$ powers of X. We have the following.*

- *The subgroups X, L are cyclic, each with order $2N^r p'q'$, and $2p'q'$. Knowing only N, one can efficiently sample (statistically close to uniform) a generator of L.*
- *X is a direct product of $X = H \cdot L$, where $H = \{(1+N)^a \mod N^{r+1} : a \in \mathbb{Z}_{N^r}\}$.*
- *Assuming the decisional composite residuosity (DCR) assumption, we have*

$$\{a \leftarrow X\} \approx_c \{b \leftarrow L\}.$$

We choose to use the subgroups X and L instead of $\mathbb{Z}^*_{N^{r+1}}$ and the $(N^r)^{th}$ powers because X, L are cyclic, which is convenient for proving security. In the following, we construct a masking scheme for messages in $\mathbb{Z}_{N^r}$ for any admissible N and polynomial $r = r(\lambda)$, based on the DCR assumption.

Construction 3 (key-homomorphic masking by DCR)**.**

- $\mathsf{Setup}(1^\lambda, \ell, p_m)$: sample an admissible 2λ bit modulus $N = (2p'+1)\cdot(2q'+1)$, and choose r the smallest integer such that $N^r > p_m$. Output $\mathsf{pp} = (N, r)$. The key space is $\mathcal{K} = \mathbb{Z}_{2p'q'}$, the message space, $\mathbb{Z}_{N^r}$, and the mask space $\mathbb{Z}^*_{N^{r+1}}$.
  Note that his Setup algorithm is *not* public coin. To use this masking scheme in LERNA, we need each client enter the protocol with public parameters $\mathsf{HM.pp}$ correctly computed by a trusted setup.
- $\mathsf{KeyGen}(\mathsf{pp})$: sample a random value $k \leftarrow \lfloor N/2 \rfloor$ and output k.
  Observe that the order of L is $2p'q'$, and that $k \mod 2p'q'$ is statistically close to uniform over $\mathbb{Z}_{2p'q'}$.

– $\mathsf{TagGen}(\mathsf{pp})$: sample ℓ random elements (forming a vector) $\mathbf{v} \in L^\ell$, each as follows. First sample $u_i \leftarrow \mathbb{Z}^*_{N^{r+1}}$, and compute the jacobi symbol $(u_i|N)$. If it's 1, then compute $v_i = u^{N^r}$. Otherwise, let a be any fixed constant with Jacobi symbol -1, and compute a component $v_i = (au)^{N^r}$. Output $\tau = \mathbf{v}$.
– $\mathsf{Mask}(\mathsf{pp}, k, \tau, \mathbf{m})$: parse τ as $\mathbf{v} \in \mathbb{Z}^{*\ell}_{N^{r+1}}$, and output the masked message $\mathbf{c}_m$, with each component $c_{m,i} = (1+N)^{m_i} v_i^k \mod N^{r+1}$.
– $\mathsf{UnMask}(\mathsf{pp}, \mathbf{c}_m, \mathbf{c}_0)$: compute (component-wise) $\mathbf{h}_m = \mathbf{c}_m / \mathbf{c}_0 \mod N^{r+1}$, and then output the message as $\mathbf{m} \leftarrow \mathsf{Dec}(\mathbf{h}_m, N, r)$. If Dec fails, then output $\perp$.
– $\mathsf{Eval}(\mathsf{pp}, L, \{\mathbf{c}_i\})$ parse L as d integer coefficients $u_1, \ldots, u_d$. Output the evaluated mask $\mathbf{c}_L = \prod_{i \in [d} \mathbf{c}_i^{u_i} \mod N^{r+1}$.

We briefly verify the correctness of our scheme in the 1-dimensional case ($\ell = 1$). Let $\tau = v$ be any tag and k be any key, and $m \in \mathbb{Z}_{N^r}$ be any message. Then the masked message c_M and a mask for zero are computed

$$c_m = (1+N)^m v^k \in \mod N^{r+1}, \quad c_0 = v^k \in \mod N^{r+1}.$$

UnMask first computes $h_m = c_m / c_0 = (1+N)^m \mod N^{r+1}$, and then $m \leftarrow \mathsf{Dec}(h_m, N, r)$. By Lemma 4, $\mathsf{Dec}(h_m, N, r)$ exactly recovers the message m, hence correctness holds.

To prove the security of the above construction, we first state and prove the following lemma. The proof argument is similar to that of Theorem 2 (EDDH from DCR) in [23].

Lemma 6. *Let λ be the security parameter. Consider any r that's polynomial in λ, and any $N = pq$, a product of two λ bit safe primes (i.e. $p = 2p'+1, q = 2q'+1$ for two primes p', q'). Let X, L, H be the subgroups defined in Lemma 5, and let $\mathsf{G}(L)$ denote the set of generators of L. Assuming the decisional composite residuosity (DCR) assumption, we have*

$$\begin{aligned} &\left\{g, g^a, g^b, g^{ab} \,\middle|\, g \leftarrow \mathsf{G}(L); a, b \leftarrow \lfloor N/2 \rfloor \right\} \\ \approx^c &\left\{g, g^a, g^b, g^{ab}h \,\middle|\, g \leftarrow \mathsf{G}(L); h \leftarrow H; a, b \leftarrow \lfloor N/2 \rfloor \right\}. \end{aligned}$$

Proof. Consider the following hybrids: For notational convenience, let $\mathsf{od}(L) = 2p'q'$ denote the order of the group L.

H1. $\left\{g, g^a, g^b, g^{ab} \,\middle|\, g \leftarrow \mathsf{G}(L); a, b \leftarrow \lfloor N/2 \rfloor \right\}$.
H2. $\left\{g, g^a, g^b, g^{ab} \,\middle|\, g \leftarrow \mathsf{G}(L); a, b \leftarrow [\mathsf{od}(L)] \right\}$.
We have $\mathsf{H}_1 \approx \mathsf{H}_0$ because for $a, b \leftarrow \lfloor N/2 \rfloor$, $a, b \mod \mathsf{od}(L)$ are close to uniform over $\mathbb{Z}_{\mathsf{od}}$.
H3. $\left\{g, g^a, g^b, g^{ab} \,\middle|\, g \leftarrow \mathsf{G}(L); a \leftarrow [\mathsf{od}(L)]; b \leftarrow [N^r \mathsf{od}(L)] \right\}$.
We have $\mathsf{H}_2 \equiv \mathsf{H}_1$.
H4. $\left\{g, x, g^b, x^b \,\middle|\, g \leftarrow \mathsf{G}(L); x \leftarrow X; b \leftarrow [N^r \mathsf{od}(L)] \right\}$.
Assuming DCR, we have $\mathsf{H}_3 \approx_c \mathsf{H}_2$ by a direct application of Lemma 5.
H5. $\left\{g, vh, g^{b_1}, v^{b_1} h^{b_2} \,\middle|\, g \leftarrow \mathsf{G}(L); v \leftarrow L; h \leftarrow H; b \leftarrow [N^r \mathsf{od}(L)] \right\}$, where $b_1 = b \mod \mathsf{od}(L)$, and $b_2 = b \mod N^r$.
We have $\mathsf{H}_4 \equiv \mathsf{H}_3$ because $X = H \cdot L$ (Lemma 5).

H6. $\{g, vh, g^b, v^b h^d \,|$
$g \leftarrow \mathsf{G}(L); v \leftarrow L; h \leftarrow H; b \leftarrow [\mathsf{od}(L)], d \leftarrow [N^r]\}$.
We have $\mathsf{H}_4 \equiv \mathsf{H}_3$ by Chinese remainder theorem.

H7. $\{g, x, g^b, x^b h \,|\, g \leftarrow \mathsf{G}(L); x \leftarrow X; h \leftarrow H; b \leftarrow [\mathsf{od}(L)]\}$.
We have $\mathsf{H}_5 \equiv \mathsf{H}_4$, again because $X = H \cdot L$.

H8. $\{g, g^a, g^b, g^{ab} h \,|\, g \leftarrow \mathsf{G}(L); h \leftarrow H; a, b \leftarrow [\mathsf{od}(L)]\}$.
Assuming DCR, have $\mathsf{H}_6 \approx_c \mathsf{H}_5$, symmetrical to $\mathsf{H}_3 \approx_c \mathsf{H}_2$.

H9. $\{g, g^a, g^b, g^{ab} h \,|\, g \leftarrow \mathsf{G}(L); h \leftarrow H; a, b \leftarrow \lfloor N/2 \rfloor\}$.
We have $\mathsf{H}_7 \approx \mathsf{H}_6$, symmetrical to $\mathsf{H}_1 \approx \mathsf{H}_0$.

□

Lemma 7. *For any modulus of the form N^r where N is a product of two distinct λ bit safe primes, and $r = r(\lambda)$ is any polynomial, Construction 3 is a (exact) key-homomorphic masking scheme for messages in $\mathbb{Z}_{N^r}$ under the DCR assumption.*

Proof. The correctness and key-homomorphism of the above construction are clear. We show that the above construction is indeed secure per Definition 4. For simplicity, we again show the 1-dimensional case ($\ell = 1$). We define a series of hybrids that starts with exactly the experiment $\mathsf{Exp}^{\mathcal{A},1}_{\mathsf{Mask}}(1^\lambda)$ and ends with the experiment $\mathsf{Exp}^{\mathcal{A},0}_{\mathsf{Mask}}(1^\lambda)$. Let A_i be the probability that $\mathcal{A}$ outputs 1 in H_i.

H1. This hybrid is exactly the experiment $\mathsf{Exp}^{\mathcal{A},1}_{\mathsf{Mask}}$. The challenger first samples a masking key $k \leftarrow \lfloor N/2 \rfloor$, and then repeat the following until $\mathcal{A}$ outputs a bit b' and stops.
   (a) The challenger samples $u_i \leftarrow \mathbb{Z}^*_{N^{r+1}}$, and computes either $v_i = u_i^{N^r}$ or $v_i = (a u_i)^{N^r}$ as the tag for iteration i, depending on the Jacobi symbol $(u_i|N)$. It sends v_i, u_i to $\mathcal{A}$. (The randomness for the sampling is u_i.)
   (b) The challenger receives a message m_i from $\mathcal{A}$, computes the masked message $c_i = (1+N)^{m_i} v_i^k$, and sends c_i to $\mathcal{A}$.

H2. In this hybrid, the challenger proceeds identically as before, except that it samples the tag v_i differently. In the beginning of the experiment, it samples a random generator g for L.
Next, during each iteration i, it samples the random value $\tilde{u}_i$ differently. It samples random exponents $a_i \leftarrow \mathbb{Z}_{N^{r+1}}$ and $b_i \leftarrow \lfloor N/2 \rfloor$. With probability 1/2, it either computes $\tilde{u}_i = (1+N)^{a_i} g^{b_i}$ or $\tilde{u}_i = (1+N)^{a_i} g^{b_i}/a$. (Recall a is a fixed constant with Jacobi symbol -1). Finally, it computes the tag $\tilde{v}_i$ from $\tilde{u}_i$ as in the previous hybrid:

$$\tilde{v}_i = \left((1+N)^{a_i} g^{b_i}\right)^{N^r} = g^{b_i N^r}.$$

In this hybrid, the masked message c_i is therefore computed as

$$\tilde{c}_i = (1+N)^{m_i} \tilde{v}_i^k = (1+N)^{m_i} g^{b_i k N^r}.$$

By the properties of listed in Lemma 5, this hybrid is statistically close to the previous one.

H3. In this hybrid, the challenger computes $\tilde{c}_i$ differently. It samples $d_i \leftarrow \mathbb{Z}_{N^r}$, and computes

$$\begin{aligned}\tilde{c}_i &= (1+N)^{m_i} g^{b_i k 2N^r} (1+N)^{d_i} \\ &= (1+N)^{m_i+d_i} g^{k b_i 2N^r}.\end{aligned}$$

By Lemma 6, this hybrid is computationally close to the previous one.

H4. In this hybrid, the challenger computes $\tilde{c}_i$ differently. It samples $d_i \leftarrow \mathbb{Z}_{N^r}$, and computes

$$\tilde{c}_i = (1+N)^{d_i} g^{b_i k 2N^r}.$$

Since d_i is randomly sampled from $\mathbb{Z}_{N^r}$, this hybrid is identical to the previous one. Hence $|A_2 - A_3| = 0$.

H5. H_4: In this hybrid, the challenger computes $\tilde{c}_i$ differently:

$$\tilde{c}_i = g^{k b_i 2N^r}.$$

By Lemma 6 again, this hybrid is computationally close to the previous one, hence $|A_3 - A_4| < \mathrm{negl}(\lambda)$.

H6. In this hybrid, the challenger follows the experiment $\mathsf{Exp}^{\mathcal{A},0}_{\mathsf{Mask}}$. In particular, compared to the previous hybrid, it samples $\tilde{u}_i$ differently as $\tilde{u}_i \leftarrow \mathbb{Z}^*_{N^{r+1}}$. In this hybrid, we have

$$\tilde{v}_i = \tilde{u}_i^{2N^r} \text{or } (a\tilde{u}_i)^{2N^r}, \quad \tilde{c}_i = \tilde{v}_i^k.$$

By the properties listed in Lemma 5, this hybrid is statistically close to the previous one, hence $|A_4 - A_5| < \mathrm{negl}(\lambda)$.

□

Extension to Class Groups. We note that the above construction using DCR can be straightforwardly generalized to using the group theoretic framework introduced in [1]. In particular, when the framework is instantiated under class groups, the above construction has a public coin Setup algorithm, which avoids the need for a trusted setup as under the DCR assumption.

A.2 Instantiating LERNA Under DCR

When instantiating LERNA with the DCR-based exact homomorphic masking scheme (Construction 3), we need to address two issues.

First, the Setup algorithm cannot be a public coin because the scheme is only secure when the factoring of N is hidden. To use the scheme, LERNA needs to assume a trusted setup that computes the public parameters $\mathsf{HM.pp}$. As noted at the end of Appendix A.1, this issue is avoided by extending the construction to use class groups.

Second, the key space of the scheme $\mathcal{K} = \mathbb{Z}_{2p'q'}$ has a modulus $2p'q'$, which is not known to the clients or the server. To secret share masking keys in $\mathcal{K}$, we

instantiate the flat secret sharing scheme with the secret space $\mathcal{M} = \mathbb{Z}$, with bounded range $\lfloor N/2 \rfloor$.

With the above issues addressed, we now set the message modulus p_m for the masking scheme. Since the scheme has *exact* key-homomorphism, there is no need for smudging noise in the LERNA protocol at all, i.e., $B_e = 0$. Consequently, we can simply set the message scaling factor $\Delta = 1$, and $B_{\text{msg}} = \Delta M B_x = M B_x$. As specified in Setup, the message modulus of the masking scheme has bit length at most $2\lambda + \log M + \log B$, and the mask modulus has bit length at most $4\lambda + \log M + \log B$. Here λ is the computational security parameter for Paillier's encryption scheme (e.g., $\lambda = 4096$ bits).

B Secure Aggregation in The UC Framework

The UC framework [15] captures the correctness and security goals of a protocol in an ideal functionality $\mathcal{F}$, such as our secure aggregation functionality $\mathcal{F}_{\text{SecAgg}}$ in Figure Fig. 8. In the ideal protocol execution, an environment $\mathcal{Z}$ provides inputs to and reads outputs from a set of dummy parties, who simply forward messages between the environment $\mathcal{Z}$ and the functionality $\mathcal{F}$. The ideal adversary/simulator $\mathcal{S}$ does not directly corrupt the dummy parties. Instead, it sends corruption commands to the functionality $\mathcal{F}$. How much information is leaked to $\mathcal{S}$, as well as what adversarial influences are allowed from $\mathcal{S}$, are completely specified by the behavior of the functionality $\mathcal{F}$.

In contrast, in the real protocol execution, the environment $\mathcal{Z}$ provides inputs to and reads outputs from the actual protocol parties. An adversary $\mathcal{A}$ directly issues corruption commands to each party. When the adversary $\mathcal{A}$ is allowed to issue such commands and the effect of those commands are specific to different types of adversaries.

In both the ideal and the real protocol executions, the environment $\mathcal{Z}$ communicates with the adversary ($\mathcal{S}$ and $\mathcal{A}$, respectively) freely. A protocol Π is said to *UC-realize* a functionality $\mathcal{F}$ if for all efficient adversary $\mathcal{A}$, there exists an efficient ideal adversary / simulator $\mathcal{S}$ such that no efficient environment $\mathcal{Z}$ can distinguish a real protocol Π against an ideal protocol with the functionality $\mathcal{F}$.

Next, we provide more details on our secure aggregation protocol and the ideal aggregation functionality $\mathcal{F}_{\text{SecAgg}}$.

Protocol Execution of Secure Aggregation. In this work, we construct a protocol between M clients $\{P_i\}$ and a server S. The protocol has a single setup phase, and multiple online phases (also referred to as aggregation sessions). In the beginning of the protocol execution, each client obtains (from the environment $\mathcal{Z}$) as its input a list of public keys of other clients. They then execute the setup phase of the protocol. In the beginning of each aggregation session, each client obtains a new aggregation input, and executes the online phase of the protocol. The server outputs the aggregation result at the end of each aggregation session.

In our construction, we describe a round-based protocol, assuming every party has access to a synchronized communication channel. Formally, this is

modeled by a functionality $\mathcal{F}_{\text{Sync}}$ (as described in the [15]). The functionality $\mathcal{F}_{\text{Sync}}$ keeps a round counter, and only advances the counter, (signaling a round is complete,) after all honest online parties' messages of the current round are received. For ease of presentation, in the main body (Section 4) we describe the round-based protocol without explicitly mentioning the functionality $\mathcal{F}_{\text{Sync}}$.

Note that although our real protocol is round-based, our ideal functionality is defined in an (asynchronized) event-driven way. This does not trivially allow the environment $\mathcal{Z}$ to distinguish the real from the ideal protocol, because the simulator $\mathcal{S}$ in the ideal protocol execution can simulate the round based communication pattern internally when communicating with the environment $\mathcal{Z}$.

Corruption in Secure Aggregation. In this work, we consider two types of corruptions. They respectively model colluding parties and dropout clients. In the following, we describe when the adversary $\mathcal{A}$ is allowed to issue the two types of corruption commands, and how a corrupted party behaves.

The first type of corruption models a statically chosen set of colluding parties, which may include the server S. To corrupt a party of the first type, the adversary $\mathcal{A}$ (in the real protocol execution) sends a command $(\mathsf{corrupt}, 1)$ to the target party (which could be a client or the server). If $\mathcal{A}$ is a semi-honest adversary, then the corrupted party sends its state to $\mathcal{A}$, and forwards all later messages to $\mathcal{A}$. If $\mathcal{A}$ is a malicious adversary, then the corrupted party sends its state to $\mathcal{A}$, and starts to execute arbitrary code as demanded by $\mathcal{A}$. We only consider *static* corruption of the first type, where the adversary $\mathcal{A}$ is only allowed to issue the $(\mathsf{corrupt}, 1)$ command at the very beginning of the protocol execution.

The second type of corruption models a statically chosen set of dropout clients during each aggregation session. To corrupt a party of the second type, the adversary sends a command $(\mathsf{corrupt}, 2)$ to the target party (which can only be a client). The corrupted party ignores any following messages in the protocol until the next aggregation session starts. The start of the next aggregation section is signaled by a new aggregation input from the environment $\mathcal{Z}$. We require the adversary $\mathcal{A}$ to decide a potential set $D_t \subseteq [M]$ of dropout clients for every aggregation session t, all at the beginning of the protocol execution. Then, the adversary $\mathcal{A}$ is allowed to issue the $(\mathsf{corrupt}, 2)$ command at any time during each aggregation session t to any client $P_i \in D_t$ from the potential dropout set in the current session.

Note that in the above modeling, each aggregation session starts with the full set of parties, so the dropout sets D_{t_1}, D_{t_2} for different sessions can overlap in arbitrary ways.

The $\mathcal{F}_{\text{SecAgg}}$ Functionality. The functionality $\mathcal{F}_{\text{SecAgg}}^{\delta,M,T,\ell,\mathsf{SampR}}$ modeling T sessions of secure aggregation among M parties with dropout threshold $0 < \delta < 1$ is described in Fig. 8.

For technical reasons, the functionality is parameterized by a modulus sampling function SampR, instead of a fixed input modulus R. In the beginning, before the first iteration, the functionality samples a modulus $R \leftarrow \mathsf{SampR}()$. In each iteration, parties hold new input vectors $\mathbf{x}_i \in \mathbb{Z}_R^\ell$ of dimension ℓ. Ag-

gregation is performed component-wise and the server recovers $\mathbf{x}_U = \sum_{i \in U} \mathbf{x}_i \mod R$ also of dimension ℓ.

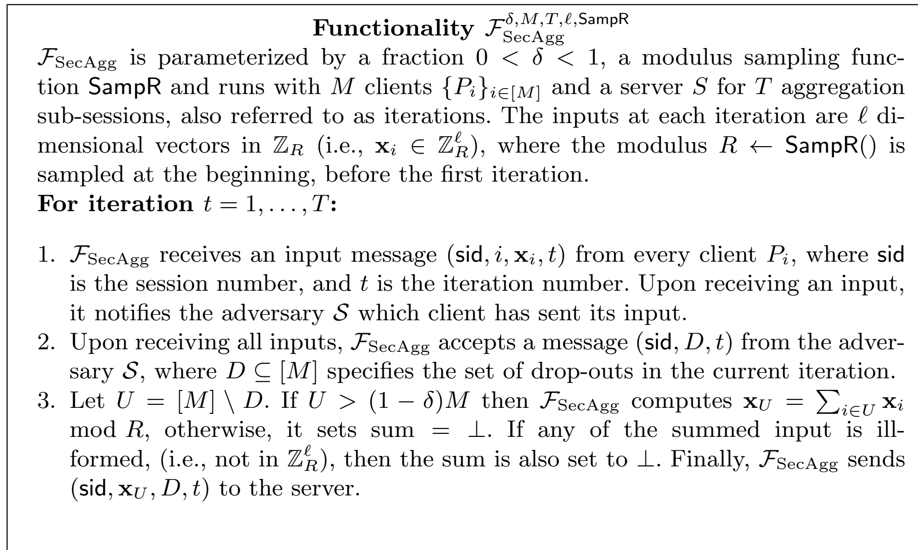

Functionality $\mathcal{F}_{\text{SecAgg}}^{\delta, M, T, \ell, \mathsf{SampR}}$

$\mathcal{F}_{\text{SecAgg}}$ is parameterized by a fraction $0 < \delta < 1$, a modulus sampling function SampR and runs with M clients $\{P_i\}_{i \in [M]}$ and a server S for T aggregation sub-sessions, also referred to as iterations. The inputs at each iteration are ℓ dimensional vectors in $\mathbb{Z}_R$ (i.e., $\mathbf{x}_i \in \mathbb{Z}_R^\ell$), where the modulus $R \leftarrow \mathsf{SampR}()$ is sampled at the beginning, before the first iteration.

For iteration $t = 1, \ldots, T$**:**

1. $\mathcal{F}_{\text{SecAgg}}$ receives an input message $(\mathsf{sid}, i, \mathbf{x}_i, t)$ from every client P_i, where sid is the session number, and t is the iteration number. Upon receiving an input, it notifies the adversary $\mathcal{S}$ which client has sent its input.
2. Upon receiving all inputs, $\mathcal{F}_{\text{SecAgg}}$ accepts a message (sid, D, t) from the adversary $\mathcal{S}$, where $D \subseteq [M]$ specifies the set of drop-outs in the current iteration.
3. Let $U = [M] \setminus D$. If $U > (1-\delta)M$ then $\mathcal{F}_{\text{SecAgg}}$ computes $\mathbf{x}_U = \sum_{i \in U} \mathbf{x}_i \mod R$, otherwise, it sets sum $= \perp$. If any of the summed input is ill-formed, (i.e., not in $\mathbb{Z}_R^\ell$), then the sum is also set to $\perp$. Finally, $\mathcal{F}_{\text{SecAgg}}$ sends $(\mathsf{sid}, \mathbf{x}_U, D, t)$ to the server.

Fig. 8: The ideal functionality for secure aggregation.

The fact that $\mathcal{F}_{\text{SecAgg}}$ does not necessarily computes aggregation over a fixed modulus is not an issue for our application. Since the goal of the LERNA framework is to compute aggregation over the integers $\mathbb{Z}^\ell$, any sampled modulus R that's large enough to avoid wrap-around of the results suffices. Our two instantiations of the framework (Section 4.4, Appendix A) both achieves the functionality $\mathcal{F}_{\text{SecAgg}}$ with such modulus sampling functions.

Another technical detail about the functionality $\mathcal{F}_{\text{SecAgg}}$ is that, in step 2, it accepts a message (sid, D, t) from the ideal adversary $\mathcal{S}$ to specify a set D of dropout clients. This may seem redundant since in the ideal protocol the functionality $\mathcal{F}_{\text{SecAgg}}$ receives any corruption command directly from the adversary $\mathcal{S}$. In particular, it learns directly which dummy clients are dropped out. (See the "Corruption in Secure Aggregation" section above for more details.) Why not just define $\mathcal{F}_{\text{SecAgg}}$ to aggregate over the remaining online clients? The answer is that in the malicious setting, the adversary $\mathcal{A}$ in the real protocol can make a corrupted client (of type 1) under its control to also behave like it has dropped out. To model this additional adversarial influence, we allow $\mathcal{S}$ to specify the "effectively dropout" set D as a separate message to the functionality.

Note on correctness. Since $\mathcal{F}_{\text{SecAgg}}$ allows the ideal adversary $\mathcal{S}$ to choose arbitrary dropout sets D, $\mathcal{S}$ may make the functionality always output $\perp$ by setting D to be larger than a δ fraction. In this case, correctness is vacuous. A trivial protocol, which lets the server S always output $(\mathsf{sid}, \perp, [M], t)$ for all t,

UC-realizes the functionality $\mathcal{F}_{\text{SecAgg}}$. To rule out this trivial protocol, we say that a semi-honest secure aggregation protocol is *non-trivial* if for every iteration $t \in [T]$ the server outputs $(\mathsf{sid}, *, D, t)$ where $D \subseteq D_t$. The malicious variant is defined analogously with the requirement $D \subseteq D_t \cup C$, where C is the set of maliciously corrupted clients.

C Security Proof of LERNA in UC

Theorem 2. *Let λ, κ be the computational and statistical security parameters. For all input range bound B_x with bit length bounded by* $\text{poly}(\lambda)$*, input dimension ℓ, number of clients M, that are polynomials in λ, number of online sessions T that is polynomial in λ, κ, and corruption and dropout thresholds γ, δ such that $\delta + \gamma < 1$, the semi-honest (malicious resp.) protocol in Section 4 is a non-trivial protocol that UC-realizes the secure aggregation functionality $\mathcal{F}_{\text{SecAgg}}^{\delta, M, T, \ell, \mathsf{SampR}}$ in the presence of semi-honest (malicious resp.) adversaries that statically corrupt less than γM clients and the server.*

The modulus sampling function SampR *depends on the concrete instantiations of the key-homomorphic masking scheme, but always samples a modulus $R \geq M \cdot B_x$ to avoid wrap-around in the aggregation result.*

The modulus sampling function SampR in the LWR instantiation (Section 4.4) outputs exactly the desired modulus $R = M \cdot B_x$. In the DCR instantiation (Appendix A), SampR outputs a modulus of the form $R = N^r \geq M \cdot B_x$, where N is an RSA modulus, and r is the smallest integer to satisfy the requirement.

In the following, we sketch a proof for the malicious variant. The semi-honest variant follows analogously.

Proof (of Theorem 2, malicious ver.). We first check non-triviality, which requires the dropout set D output by an honest server at any iteration t does not exceed the adversarially chosen dropout set D_t plus the corruption set C. By inspection, our sever indeed always outputs the set D of clients who don't send a message during the online step 2, which is either actually dropped out, or corrupted.

Next, we describe an ideal adversary/simulator $\mathcal{S}$ that externally interacts with the functionality $\mathcal{F}_{\text{SecAgg}}$ and the environment $\mathcal{Z}$, while internally simulates a protocol execution with an instance of the adversary $\mathcal{A}$.

When interacting with $\mathcal{Z}$, $\mathcal{S}$ simply forwards all communication between $\mathcal{A}$ and $\mathcal{Z}$. During the simulated protocol execution, $\mathcal{S}$ simulates messages from honest parties to interact internally with $\mathcal{A}$. Importantly, $\mathcal{S}$ needs to simulate these messages without knowing honest parties' inputs while ensuring the simulated view of $\mathcal{A}$ is indistinguishable from the real-world execution. During the simulated execution, when $\mathcal{A}$ issues any corruption command, $\mathcal{S}$ forwards them to $\mathcal{F}_{\text{SecAgg}}$.

We show that the view of $\mathcal{Z}$, consisting of

- the inputs supplied by $\mathcal{Z}$ to honest parties,
- the view of $\mathcal{A}$, and

– the output of honest parties, which only exists when the server S is honest,

in the ideal execution and in a real execution are computationally close. Hence $\mathcal{Z}$ cannot distinguish the two executions. In the following, let C denotes the set of corrupted parties during the simulation, and H denote the set of (online) honest parties.

The Case of an Honest Server. When the server S is not corrupted, $\mathcal{A}$ in the simulated execution only sees the following messages:

– Setup: shares of honest clients' masking keys $\{k_j^i\}_{j\in C}^{i\in H}$;
– Online: an online set $U \subseteq [M]$, and honest signatures $\sigma_i(h_U)_{i\in H}$.

Since none of these messages depend on honest clients' inputs, $\mathcal{S}$ can simulate the honest clients and the server by simply following the protocol, without knowing their inputs. In particular, $\mathcal{S}$ simulates the honest clients with input vectors $\mathbf{0}$.

$\mathcal{S}$ also needs to make sure $\mathcal{F}_{\text{SecAgg}}$ outputs (to the dummy server S) consistent aggregation results with the real execution. In particular, through interacting with $\mathcal{F}_{\text{SecAgg}}$, $\mathcal{S}$ needs to emulate the adversarial influences that $\mathcal{A}$ would make on the results.

First consider the case where the honest server outputs $(\mathsf{sid}, \perp, D, t)$ in the simulation. If $|D| \geq \delta M$, then clearly this is caused by too many clients dropping out during the simulation. To emulate this effect, $\mathcal{S}$ simply sends D to $\mathcal{F}_{\text{SecAgg}}$ (step 2). According to the functionality, $\mathcal{F}_{\text{SecAgg}}$ outputs $(\mathsf{sid}, \perp, D, t)$ to the dummy server, as desired.

If $|D| < \delta M$, then by correctness (Lemma 3), we conclude that there exists at least one corrupted client (which cause the server to output $\perp$, through adversarial messages). To emulate this effect, $\mathcal{S}$ sends D to $\mathcal{F}_{\text{SecAgg}}$, and sets one of the corrupted client's inputs to $\perp$. According to the functionality, $\mathcal{F}_{\text{SecAgg}}$ outputs $(\mathsf{sid}, \perp, D, t)$ to the dummy server, due to an ill-formed summand.

Next, consider the case where the server outputs $(\mathsf{sid}, \widetilde{\mathbf{x}}_U, D, t)$ during the simulation. Recall that $\mathcal{S}$ simulates honest clients by setting their inputs to $\mathbf{0}$. We claim the aggregation result $\widetilde{\mathbf{x}}_U$ is exactly the effective sum of corrupted parties, which follows from the key-homomorphic property (Definition 2, 3) of our masking scheme. To emulate this effect, $\mathcal{S}$ sends D to $\mathcal{F}_{\text{SecAgg}}$, and sets one (if there is any) of the corrupted client's input to $\widetilde{\mathbf{x}}_U$, and the rest to $\mathbf{0}$.

The Case of a Corrupted Server. When the server S is corrupted, the remaining honest clients don't have any output in both the ideal and the real execution. $\mathcal{S}$ only needs to make sure the view of $\mathcal{A}$ in the simulated execution is consistent with the honest clients' input. The view of $\mathcal{A}$ in the simulated execution consists of the following:

– Setup: shares of honest parties' masking keys $\{k_j^i\}_{j\in C}^{i\in H}$;
– Online: The masked input vectors

$$\{\mathbf{z}_i \leftarrow \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \Delta\mathbf{x}_i)\}_{i\in H}$$

from honest clients, the honest signatures $\{\sigma_i(h_{U_i})\}_{i\in H}$ where $\{U_i\}$ are possibly different online sets sent to honest client $\{P_i\}_H$, and the reconstruction

vectors

$$\{\mathbf{w}_j = \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i \in U_j} k_j^i, \tau, \mathbf{0}) + \mathbf{e}_j\}_{i \in H \cap Q}$$

from honest committee members.

We first note that in the protocol, an honest committee member only sends a reconstruction message $\mathbf{w}_j$ (online step 5) if it sees at least $(1+\gamma)M/2$ valid signatures on U_j. We claim that there exists at most one such set, denoted U^*, with enough valid signatures.

Claim 4. *For each online session t, there exists at most one set U^* among all online sets $\{U_j\}_{H \cap Q}$ received by honest committee members, that have at least $(1+\gamma)M/2$ valid signatures.*

Proof. Since each such set requires at least $(1+\gamma)M/2$ valid signatures, more than $(1+\gamma)M/2 - \gamma M = (1-\gamma)M/2$ are from honest clients. Two such sets would require more than $(1-\gamma)M$ honest signatures, which is more than the number of honest clients. This is impossible, as each honest client creates at most one such signature. □

When such a set U^* exists, and when enough honest committee members send reconstruction vectors for U^* to the corrupted server, the view of $\mathcal{A}$ allows it to recover the sum of honest inputs $\mathbf{x}_{H \cap U^*}$ included in U^*. In order to simulate such a view, $\mathcal{S}$ clearly needs to learn the sum $\mathbf{x}_{H \cap U^*}$ from $\mathcal{F}_{\text{SecAgg}}$, which is possible by setting $D = [M] \setminus U^*$, and sending (sid, D, t) to $\mathcal{F}_{\text{SecAgg}}$ in step 2.

However, there seems to be a timing issue. $\mathcal{S}$ only learns the set U^* and hence the sum $\mathbf{x}_{H \cap U^*}$ after online step 4 of the simulated protocol, where the corrupted server sends signatures on U_j to committee members P_j. This means $\mathcal{S}$ needs to simulate the masked input vectors $\{\mathbf{z}_i\}_{i \in H}$, during online step 1, without knowing the sum $\mathbf{x}_{H \cap U^*}$.

To solve this issue, the strategy is to set the simulated masked input vectors to empty masks of $\mathbf{0}$: $\tilde{z}_i \leftarrow \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \mathbf{0})$. Later during online step 5 $\mathcal{S}$, already knowing the sum $\mathbf{x}_{H \cap U^*}$, program it into the simulated reconstruction vectors $\{\widetilde{\mathbf{w}}_i\}_{H \cap Q}$, such that they reconstructs to the correct result. This is possible due to the key-homomorphic property of the masking scheme. We describe the simulation strategy in more detail below.

– **Setup phase.** In Step 1, honest clients are supposed to sample their masking keys $\{k_i\}_{i \in H}$ and secret shares them with the committee. $\mathcal{S}$ simulates honest parties simply following the protocol. It obtains key shares $\{k_j^i\}_{j \in Q \cap C} = \mathsf{Share}(\mathsf{SS.pp}, k_i)$, encrypts them, and sends them to the server.
  In Step 2, $\mathcal{S}$ receives encrypted shares from corrupted clients. It decrypts them and obtains $\{k_j^i\}_{j \in H \cap Q}^{i \in C}$. In case of decryption failures, it sets corresponding shares to some default value.
– **Online phase.** In Step 1, honest clients are supposed to compute masked inputs vectors: $\mathbf{z}_i = \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \Delta\mathbf{x}_i)$, with honest input vector $\mathbf{x}_i$. $\mathcal{S}$ simulates those by computing (empty) masks of $\mathbf{0}$: $\tilde{\mathbf{z}}_i \leftarrow \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \mathbf{0})$. We will rely on the security of secret sharing to argue the masking keys k_i are

hidden from $\mathcal{A}$, and then invoke the security of the masking scheme to argue $\{\mathbf{z}_i\}_H$ and $\{\tilde{\mathbf{z}}_i\}_H$ are indistinguishable.

In Steps 2, 3 and 4, the (malicious) server sends h_{U_i} to clients P_i for their signatures, and sends collected signatures on U_j to committee members P_j. $\mathcal{S}$ simulates the honest clients and committee members following the protocol. As argued earlier, there is a unique $U^* = U_j$, received by some honest client P_j that has enough valid signatures. (The case where no such U^* exists is analogous but only simpler, hence is omitted.)

The simulator now pauses to interact with $\mathcal{F}_{\text{SecAgg}}$ at Step 2.

– **Interacting with $\mathcal{F}_{\text{SecAgg}}$.** In Step 1, $\mathcal{F}_{\text{SecAgg}}$ accepts inputs for corrupted clients from $\mathcal{S}$, which are set to $\mathbf{0}$. $\mathcal{F}_{\text{SecAgg}}$ also notifies $\mathcal{S}$ whenever an input is received from an honest dummy client. $\mathcal{S}$ proceeds in the above internally simulated protocol during online step 1 according to these notifications.

  In Step 2, $\mathcal{S}$ sets $\widetilde{D}$ to be the set of honest parties in U^*, i.e., $\tilde{D} = H \cap U^*$, and sends $(\mathsf{sid}, \tilde{D}, t)$ to $\mathcal{F}_{\text{SecAgg}}$. The functionality replies with the sum $\mathbf{x}_{H\cap U^*}$.
– **Online phase cont.** In Step 5, honest committee members who receive U^* with enough valid signatures are supposed to compute the reconstruction vectors $\mathbf{w}_j = \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i\in U^*} k_j^i, \tau, \mathbf{0}) + \mathbf{e}_j$. To simulate $\mathbf{w}_i$, first note that by almost key homomorphism of HM, we have

$$\begin{aligned}\mathbf{w}_j \approx \mathsf{Mask}(\mathsf{HM.pp}, &\sum_{i\in C\cap U^*} k_j^i, \tau, \mathbf{0})\\ &+ \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i\in H\cap U^*} k_j^i, \tau, \mathbf{0}) + \mathbf{e}_j,\end{aligned}$$

  where we abuse notation to homomorphically "add" (+) two masks by implicitly running the Eval algorithm. We use this notation in the following also for other linear functions.

  $\mathcal{S}$ can simulate the first term $\mathsf{Mask}(\mathsf{HM.pp}, \sum_{i\in C\cap U^*} k_j^i, \tau, \mathbf{0})$ exactly using the received shares during the setup phase. The goal is to simulate the second term – which depends on key shares k_j^i of an honest masking key k_i to an honest committee member – without leaking additional information about k_i beyond what's already known to $\mathcal{A}$. In particular, these include

  - The simulated message $\tilde{\mathbf{z}}_i$, which is an empty mask computed under k_i.
  - The corrupted shares of k_i, which are $\{k_{j'}^i\}_{j'\in C\cap Q}$.

  To this end, we first invoke γ-simulation privacy (Definition 7) to simulate the honest share k_j^i according to the corrupted shares $\{k_{j'}^i\}_{j'\in C\cap Q}$, and the masking key itself k_i. We will next invoke flatness (Definition 8) and key-homomorphism to replace the use of k_i with an empty mask under k_i, i.e. the message $\widetilde{\mathbf{z}}_i$.

  In more detail, by simulation privacy and flatness, there exists a deterministic linear function Ext_j that computes $k_j^i = \mathsf{Ext}_j(k_i, \{k_{j'}^i\}_{j'\in C\cap Q})$ with integer coefficients bounded by $O(\log\kappa)$, where κ is the statistical security parameter. We further separate the coefficient before the secret k_i out, and write $k_j^i =$

$c_j \cdot k_i + \mathsf{Ext}'_j(\{k^i_{j'}\}_{j' \in C \cap Q})$. We have

$$\sum_{i \in H \cap U^*} k^i_j = c_j \big(\sum_{i \in H \cap U^*} k_i \big) + \underbrace{\big(\sum_{i \in H \cap U^*} \mathsf{Ext}'_j(\{k^i_{j'}\}_{j' \in C \cap Q}) \big)}_{f_i},$$

where we call the right term f_j for short, which $\mathcal{S}$ can simulate exactly. Finally, by approximate key-homomorhpism, we can simulate the second term as

$$\begin{aligned} &\mathsf{Mask}(\mathsf{HM.pp}, \sum_{i \in H \cap U^*} k^i_j, \tau, \mathbf{0}) + \mathbf{e}_j \\ &\approx c_j \cdot \mathsf{Mask}(\mathsf{HM.pp}, \sum_{i \in H \cap U^*} k_i, \tau, \mathbf{0}) \\ &\quad + \mathsf{Mask}(\mathsf{HM.pp}, f_j, \tau, \mathbf{0}) + \mathbf{e}_j \\ &\approx c_j \sum_{i \in H \cap U^*} \tilde{\mathbf{z}}_i + \mathsf{Mask}(\mathsf{HM.pp}, f_{\mathsf{CShare}}, \tau, \mathbf{0}) + \mathbf{e}_j \end{aligned}$$

The last thing to do is to program the sum $\mathbf{x}_{H \cap U^*}$ in to the simulated $\widetilde{\mathbf{w}}_j$ vectors. Combining the above, $\mathcal{S}$ simulate $\widetilde{\mathbf{w}}_j$ as

$$\begin{aligned} \widetilde{\mathbf{w}}_j = &\mathsf{Mask}(\mathsf{HM.pp}, \sum_{i \in C \cap U^*} k^i_j, \tau, \mathbf{0}) \\ &+ c_j \sum_{i \in H \cap U^*} \tilde{\mathbf{z}}_i + \mathsf{Mask}(\mathsf{HM.pp}, f_j, \tau, \mathbf{0}) \\ &+ \mathsf{Mask}(\mathsf{HM.pp}, 0, \tau, -\mathbf{x}_{H \cap U^*}) + \mathbf{e}_j. \end{aligned}$$

In the above arguments using ε-approximate key-homomorhpism, we assume that the noise $\mathbf{e}_j$ is sampled uniformly from a sufficiently large range $[B_e]$ to smudge the distribution differences introduced by homomorphic evaluation. Inspecting the coefficients shows that $B_e \geq O(\varepsilon \cdot \log \kappa \cdot M \cdot 2^\kappa)$ suffices.

A similar series of hybrid experiments to those in Section 4.3 concludes the proof. □

D Experimental Evaluation: Further Results

We briefly summarize the major computation tasks of LERNA's clients and the server. In the setup phase, the client tasks are the following.

- ShareAndEnc: Each client secret shares its masking key, and encrypt each share under the target share holder's public key.
- DecShares: Each committee member receives encrypted key shares and decrypts them.
- SumShares: Each committee member pre-computes a sum over received key shares. Later during each online phase, when clients drop out, the committee member subtracts their shares from the pre-computed sum.

In the online phase, the client tasks are the following.

– MaskInputs: Each client masks its input vector.
– SubShares: Each committee member learns the dropout clients, and subtracts their key shares from the pre-computed sum.
– MaskShares: Each committee member computes an empty mask using its aggregated key share.

In the setup phase, the server has essentially no computation task beyond forwarding messages. In the online phase, the server tasks are the following.

– SumMasks: Homomorphically sum masked input vectors.
– Recon: Homomorphically evaluate the Recon algorithm over received empty masks.

Table 4 records computation times of a committee member client, broken down by major computation tasks, when the dropout rate is fixed at 10%. We omit the table for a non-member client, since it essentially contains a subset of the rows, namely ShareAndEnc in the setup phase, and MaskInputs in the online phase.

Tasks	M=20K (s)	M=40K (s)	M=80K (s)
ShareAndEnc	3.75	3.76	3.75
DecShares	3.22	6.71	13.6
SumShares	1.13	2.29	4.57
(Setup) Total	**8.10**	**12.76**	**22.00**
MaskInputs	0.06	0.06	0.06
SubShares	0.10	0.18	0.40
MaskShares	0.05	0.04	0.05
(Online) Total	**0.21**	**0.28**	**0.50**

Table 4: Committee member client computation times (s).

Table 5 records computation times of the server, similar to Table 4. Since that the major computation tasks for our server all happen in the online phase, we omit its Setup phase.

Tasks	M=20K (s)	M=40K (s)	M=80K (s)
SumMasks	4.67	9.75	18.92
Recon	0.12	0.13	0.13
(Online) Total	**4.79**	**9.88**	**19.05**

Table 5: Server computation times (s).

E LERNA with Small Number of Clients

When running LERNA with a small number of clients, e.g. $M \leq 80$ it becomes suitable to use Shamir's secret sharing scheme as an alternative to the committee based flat secret sharing constructed in Section 3.2.

1. There is no need for running a committee based scheme when the number of clients M is small.
2. Even though Shamir cannot achieve "flatness" per Definition 8, combined with the standard trick of "clearing the denominator", Shamir's reconstruction coefficients become bounded by $(M!)^2$, wich is tolerable for small M.

We note that since there is no longer a fixed committee throughout multiple iterations in this setting, this variant with small number of clients tolerates adaptive corruption.

In the following, we first introduce two variants of Shamir, to be used with the LWR and DCR based masking scheme respectively. We then highlight the technical changes to the protocol and parameter settings when instantiated with Shamir's secret sharing.

E.1 Two Variants of Shamir

We briefly describe two known variants of Shamir suitable for using with the LWR and DCR based masking scheme respectively.

Shamir with Bounded Recon Coefficients. As described in the beginning of Section 3.2, since the LWR based masking scheme has only approximate key-homomorphism, evaluating Recon homomorphically creates an additive noise, which grows with the magnitude of the coefficients of Recon.

The usual Shamir's secret sharing works over a field and relies on Lagrange interpolation for reconstruction. Due to the use of division for computing the Lagrange coefficients, they can be arbitrarily large (in the field). To avoid the divisions, a standard trick is to multiply a factor $\alpha = (M!)$ to each coefficient to "clear the denominator". The multiplied coefficients are all bounded by $(M!)^2$. As a result, this variant does not reconstruct the original secret, but α times the secret. We refer the reader to [12] for more details of this trick.

Construction 4 ([12])**.** Let M be the number of share holders, q be the modulus of the secret space $\mathbb{Z}_q$, and t be the privacy threshold.

- $\mathsf{Share}(s)$ samples a degree t random polynomial f such that $f(0) = s$, and with other coefficients in $\mathbb{Z}_q$. Output $\{s_j\}_{j \in [M]}$, where $s_j = f(j)$.
- $\mathsf{Recon}_\alpha(W, \{s_j\}_{j \in W})$ outputs $\perp$ if $|W| \leq t$. Otherwise, pick the first $t+1$ indices in W as W', let $\{l_j\}_{W'}$ be the Lagrange coefficients for interpolating those points to position 0, and outputs $s' = \sum_{j \in W'} (\alpha l_j) \cdot s_j \mod q$.

Note that the security proof of LERNA relies on extending shares of 0 for corrupted parties into a full set of shares consistent with any secret s (formally

defined through the Ext algorithm in Definition 7). The usual Shamir's scheme satisfies this requirement: given any subset of shares below the reconstruction threshold, Ext can interpolate them into a polynomial f such that $f(0) = s$, for any s, and then compute the rest of the shares using f. However, this Ext algorithm again has arbitrarily large coefficients, due to the use of division in polynomial interpolation.

We need to use the trick of "clearing the denominator" again with Ext to make its coefficients bounded by $(M!)^2$. We call this scaled version Ext_α. Ext_α can only extend to a set of shares scaled by $\alpha = M!$.

Shamir over Bounded Integers. As mentioned in the overview, since the DCR based masking scheme has an unknown modulus for its secret space, it needs a secret sharing scheme over bounded integers.

First, as there is no division over the integers this variant requires the same trick of "clearing the denominator" to avoid computing division for reconstruction coefficients. Second, Shamir's scheme relies on sampling a random polynomial to share a secret. But how should one sample a random polynomial over the integers? [13] shows that for sharing secrets of bounded magnitude by B, it suffices to sample a polynomial with random coefficients from the range $2B \cdot M \cdot (M!)^2 \cdot 2^\kappa$, where κ is the statistical security parameter.

Construction 5 ([13])**.** Let M be the number of share holders, B be the upperbound on the integer secret values, and t be the privacy threshold.

- $\mathsf{Share}(s)$ samples a random degree t polynomial f such that $f(0) = s$, and other coefficients from $[0, \ldots, 2BM(M!)^2 2^\kappa]$. Output $\{s_j\}_{j\in[M]}$ where $s_j = f(j)$.
- $\mathsf{Recon}_\alpha(W, \{s_j\}_{j\in W})$ outputs $\perp$ if $|W| \le t$. Otherwise, pick the first $t+1$ indices in W as W', let $\{l_j\}_{W'}$ be the Lagrange coefficients (in $\mathbb{Q}$) for interpolating those points to position 0. Note that $\alpha l_j \in \mathbb{Z}$. Output $s' = \sum_{j\in W'}(\alpha l_j) \cdot s_j$ over $\mathbb{Z}$.

Similar to the first variant, there exists an algorithm Ext_α, satisfying an analogous security definition to Definition 7, that extends shares of 0 for corrupted parties into a full set of scaled-by-α shares consistent with any secret s.

E.2 LERNA Protocol with Shamir

Due to the fact that the two variants of Shamir are no longer committee based, and both have a Recon_α algorithm that recovers not the origional secret s, but $\alpha \cdot s$, we need to make some minor changes to the LERNA protocol.

<u>*Setup Phase.*</u> The only change is that the secret sharing scheme doesn't have a $\mathsf{SS.Setup}$ algorithm to sample a committee Q anymore. Each client P_i in step 1 directly secret shares its masking key k_i to all other clients $\{k^i_j\}_{j\in[N]} \leftarrow \mathsf{Share}(k_i)$, and then proceeds as in Figure 1.

<u>*Online Phase.*</u>

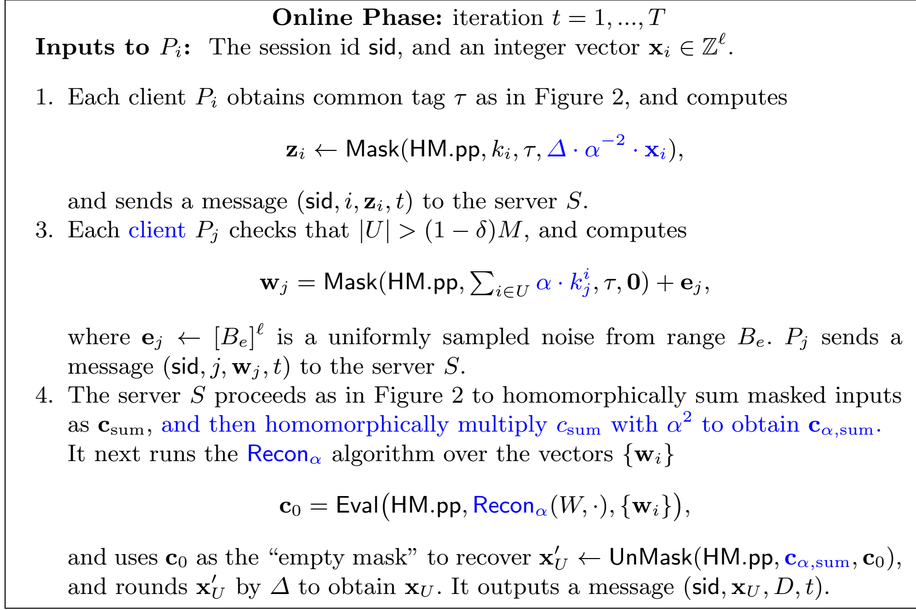

Online Phase: iteration $t = 1, ..., T$

Inputs to P_i: The session id sid, and an integer vector $\mathbf{x}_i \in \mathbb{Z}^\ell$.

1. Each client P_i obtains common tag τ as in Figure 2, and computes

$$\mathbf{z}_i \leftarrow \mathsf{Mask}(\mathsf{HM.pp}, k_i, \tau, \Delta \cdot \alpha^{-2} \cdot \mathbf{x}_i),$$

and sends a message $(\mathsf{sid}, i, \mathbf{z}_i, t)$ to the server S.

3. Each client P_j checks that $|U| > (1-\delta)M$, and computes

$$\mathbf{w}_j = \mathsf{Mask}(\mathsf{HM.pp}, \textstyle\sum_{i \in U} \alpha \cdot k_j^i, \tau, \mathbf{0}) + \mathbf{e}_j,$$

where $\mathbf{e}_j \leftarrow [B_e]^\ell$ is a uniformly sampled noise from range B_e. P_j sends a message $(\mathsf{sid}, j, \mathbf{w}_j, t)$ to the server S.

4. The server S proceeds as in Figure 2 to homomorphically sum masked inputs as $\mathbf{c}_{\text{sum}}$, and then homomorphically multiply c_{sum} with α^2 to obtain $\mathbf{c}_{\alpha,\text{sum}}$. It next runs the Recon_α algorithm over the vectors $\{\mathbf{w}_i\}$

$$\mathbf{c}_0 = \mathsf{Eval}\big(\mathsf{HM.pp}, \mathsf{Recon}_\alpha(W, \cdot), \{\mathbf{w}_i\}\big),$$

and uses $\mathbf{c}_0$ as the "empty mask" to recover $\mathbf{x}'_U \leftarrow \mathsf{UnMask}(\mathsf{HM.pp}, \mathbf{c}_{\alpha,\text{sum}}, \mathbf{c}_0)$, and rounds $\mathbf{x}'_U$ by Δ to obtain $\mathbf{x}_U$. It outputs a message $(\mathsf{sid}, \mathbf{x}_U, D, t)$.

Fig. 9: LERNA protocol for the online phase with Shamir.

1. As explained in the previous section, the two variants of Shamir have a Ext_α algorithm, that can only be used for simulating *scaled* shares for the honest clients. Hence in step 3 of the online phase, each honest client creates its empty mask $\mathbf{w}_j$ under scaled shares $\sum_{i \in U} \alpha \cdot k_j^i$ for the proof to go through.
2. The two variants have a Recon_α algorithm that reconstructs the secret s scaled by α. In step 4, the server runs Recon_α homomorphically over reconstruction vectors $\mathbf{w}_i$, which contain shares that are already scaled by α as required above. The resulting empty mask $\mathbf{c}_0$ is effectively under the key $\alpha^2 k_U$. To match this additional factor of α^2 in $\mathbf{c}_0$, the server needs to homomorphically scale $\mathbf{c}_{\text{sum}}$ also by α^2. And to cancel the effect of scaling $\mathbf{c}_{\text{sum}}$, each client needs to multiplies its input by $\alpha^{-2} \mod p_m$ in step 1 before masking it.

The above changes are highlighted by blue in Figure 9.

Parameter Settings for LWR Instantiation. There are two parameters we need to set, the noise magnitude B_e and the message scaling factor Δ. B_e needs to be sufficiently large to smudge the homomorphic noise created by homomorphically evaluating Ext_α during simulation. Since Ext_α has coefficients bounded by $(M!)^2$, it suffices to set $B_e = O((M!)^2 \cdot M \cdot 2^\kappa)$. We refer the reader to Appendix 4.3 and Appendix C for details on the simulation.

Δ needs to be sufficiently large to remove the homomorphic noise created by homomorphically evaluating Recon_α over the already noisy vectors $\mathbf{w}_j$. The final error has magnitude bounded by $M(M!)^2 B_e = O(M^2(M!)^4 2^\kappa)$, so we set $\Delta = O(M^2(M!)^4 2^\kappa)$.

Under these settings, we can setup the LWR-based masking scheme with message modulus $p_m > \Delta M B_x$, which has bit length $\log p_m \leq (4M+3)\log M + \kappa + \log B_x$. An additional requirement on p_m comes from the highlighted online phase Step 1. The protocol requires multiplying $\alpha^{-2} \mod p_m$ to each client's input. It suffices to choose p_m as a prime, so that α^{-1} always exists.

Parameter Settings for DCR Instantiation. Since the DCR-based masking scheme has exact key-homomorphism, there is no homomorphic noise to smudge during simulation. Therefore, we can set $B_e = 0$, and consequently $\Delta = 1$.

Under these settings, we setup the DCR-based scheme with message modulus $p_m > MB_x$, which has bit-length at most $2\lambda + \log M + \log B$. Here λ is the computational security parameter for Paillier's encryption scheme (e.g., $\lambda = 4096$ bits).