

---

# OPA: One-shot Private Aggregation with Single Client Interaction and its Applications to Federated Learning

---

**Harish Karthikeyan, Antigoni Polychroniadou**

J.P. Morgan AI Research, J.P. Morgan AlgoCRYPT Center of Excellence  
New York, NY, USA

Email: `firstname.lastname@jpmchase.com`

## Abstract

Our work aims to minimize interaction in secure computation due to the high cost and challenges associated with communication rounds, particularly in scenarios with many clients. In this work, we revisit the problem of secure aggregation in the single-server setting where a single evaluation server can securely aggregate client-held individual inputs. Our key contribution is the introduction of One-shot Private Aggregation (OPA) where clients speak only once (or even choose not to speak) per aggregation evaluation. Since each client communicates only once per aggregation, this simplifies managing dropouts and dynamic participation, contrasting with multi-round protocols and aligning with plaintext secure aggregation, where clients interact only once.

We construct OPA based on LWR, LWE, class groups, DCR and demonstrate applications to privacy-preserving Federated Learning (FL) where clients *speak once*. This is a sharp departure from prior multi-round FL protocols whose study was initiated by Bonawitz et al. (CCS, 2017). Moreover, unlike the YOSO (You Only Speak Once) model for general secure computation, OPA eliminates complex committee selection protocols to achieve adaptive security. Beyond asymptotic improvements, OPA is practical, outperforming state-of-the-art solutions. We benchmark logistic regression classifiers for two datasets, while also building an MLP classifier to train on MNIST, CIFAR-10, and CIFAR-100 datasets.

We build two flavors of OPA (1) from (threshold) key homomorphic PRF and (2) from seed homomorphic PRG and secret sharing. The threshold Key homomorphic PRF addresses shortcomings observed in previous works that relied on DDH and LWR in the work of Boneh *et al.* (CRYPTO, 2013), marking it as an independent contribution to our work. Moreover, we also present new threshold key homomorphic PRFs based on class groups or DCR or the LWR assumption.

## 1 Introduction

Minimizing interaction in Multiparty Computation (MPC) stands as a highly sought-after objective in the field of secure computation. This is primarily because each communication round is costly, and ensuring the liveness of participants, particularly in scenarios involving a large number of parties, poses significant challenges. Unlike throughput, latency is now primarily constrained by physical limitations, making it exceedingly difficult to reduce the time required for a communication round substantially. Furthermore, non-interactive primitives offer increased versatility and are better suited as foundational building blocks. However, any non-interactive protocol, which operates with a single communication round, becomes susceptible to a vulnerability referred to as the “residual attack” [54] where the server can collude with some clients and evaluate the function on as many inputs as they wish revealing the inputs of the honest parties.

In this work, we explore a natural “hybrid” model that sits between the 2-round and 1-round settings. Specifically, our model allows for private aggregation, aided by a committee of members, where the client only speaks once. This approach brings us closer to achieving non-interactive protocols while preserving traditional security guarantees. Our specific focus is within the domain of secure aggregation protocols, where a group of  $n$  clients  $P_i$  for  $i \in [n]$  hold a private value  $x_i$ , wish to learn the sum  $\sum_i x_i$  without leaking any information about the individual  $x_i$ . In this model, clients release encoded versions of their confidential inputs  $x_i$  to a designated committee of ephemeral members and they go offline, they only speak once. Later, any subset of the ephemeral members can compute these encodings by simply transmitting a single public message to an unchanging, stateless evaluator or server. This message conveys solely the outcome of the secure aggregation and nothing else. Of significant note, the ephemeral members are stateless, speak only once, and can change (or not) per aggregation session. With that in mind, the committee members can be regarded as another subset of clients who abstain from contributing input when they are selected to serve on the committee during a current aggregation session. Each client/committee member communicates just once per aggregation, eliminating the complexity of handling dropouts commonly encountered in multi-round secure aggregation protocols. The security guarantee is that an adversary corrupting a subset of clients and the committee members learn no information about the private inputs of honest clients, beyond the sum of the outputs of these honest clients. We present a standard simulation-based security proof against both a malicious server and a semi-honest server. We present extensions that offer stronger security guarantees. See Section A for a comparison to other models.

Our main application is Federated Learning (FL) in which a server trains a model using data from multiple clients. This process unfolds in iterations where a randomly chosen subset of clients (or a set of clients based on the history of their availability) receives the model’s current weights. These clients update the model using their local data and send back the updated weights. The server then computes the average of these weights, repeating this cycle until model convergence is achieved. This approach enhances client privacy by only accessing aggregated results, rather than raw data.

## 2 Our Contributions

We introduce OPA designed to achieve maximal flexibility by granting honest parties the choice to speak once or remain silent, fostering dynamic participation from one aggregation to the next. This diverges from prior approaches [17, 13, 62, 67, 52], which necessitate multiple interaction rounds and the management of dropout parties to handle communication delays or lost connections in federated learning.

**Cryptographic Assumptions:** We construct OPA protocols providing a suite of six distinct versions based on a diverse spectrum of assumptions:

- Learning With Rounding (LWR) Assumption
- Learning with Errors (LWE) Assumption
- Hidden Subgroup Membership ( $HSM_M$ ) assumption where  $M$  is a prime integer.
- $HSM_M$  assumption where  $M = p^k$  for some prime  $p$  and integer  $k$ .
- $HSM_M$  assumption where  $M = 2^k$
- $HSM_M$  assumption where  $M = N$  where  $N$  is an RSA modulus (i.e., the DCR assumption)

**Threat Model:** OPA does not require any trusted setup for keys and for  $M$  being either a prime or an exponent of prime, or the LWR assumption, we do not require any trusted setup of parameters either. We allow the server to be maliciously corrupted, maliciously corrupt clients, and corrupt up to a certain threshold  $t$  of the committee members where  $t$  is the corruption threshold for secret reconstruction. Based on the CL framework we additionally strengthen the security to allow for the compromise of all the committee members (while ensuring the server is not corrupt).

**Variants:** In Figure 1, we summarize our variants. We offer two variants<sup>1</sup>:

---

<sup>1</sup>Our construction based on seed-homomorphic PRG can also be built from a (length extended) key-homomorphic PRF (KHPRF). However, such length-extended efficient KHPRFs, unlike seed-homomorphic PRG, are only known from the random oracle model. See Section G for the instantiations. Note that the resulting construction’s asymptotic performance will match the ones from seed-homomorphic PRG.

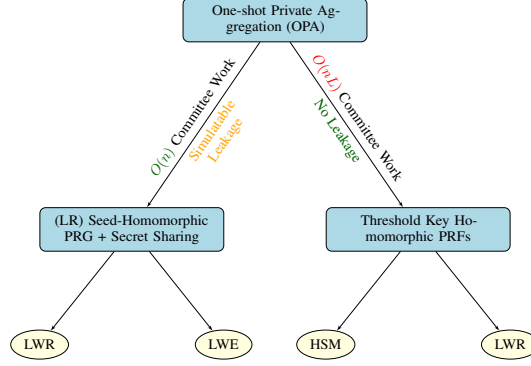


Figure 1: Our Key Contributions. Here,  $n$  is the number of clients, and  $L$  is the length of the input vector. Finally, “LR” refers to leakage-resilient. Note that the HSM assumption also subsumes the DCR assumption.

- OPA: A variant that is built from seed-homomorphic PRG that is leakage resilient and any secret sharing scheme. The resulting construction has the benefit of the committee doing work that is independent of the vector length. However, the server receives the sum of the keys of the honest parties, as leakage which requires this leakage to be simulatable.
- OPA’: We also present a variant built from threshold/distributed key-homomorphic PRF where there is no leakage. In this case, the committee’s work is dependent on the length of the input vector.

**Other contributions:**

- We build the first Key Homomorphic PRF based on the  $HSM_M$  assumption. We later extend it to the distributed key homomorphic setting by using a modified Shamir’s Secret Sharing scheme but over integers.
- We extend the almost Key Homomorphic PRF based on the LWR assumption [19, 48] using Shamir Secret Sharing over prime-order fields. In doing so, we fix gaps in the prior Distributed Key Homomorphic PRF based on LWR, as proposed by Boneh *et al.* [19]
- We also extend Shamir Secret Sharing over Integers to a packed version which enables packing more secrets in one succinct representation.
- We also present a malicious variant of OPA that empowers a server to detect a malicious client that has inconsistent shares. Meanwhile, we extend OPA’ to the setting where all the committee is corrupted but the server is honest. We want to ensure that the honest client’s input is still hidden, even if the committee member gets the ciphertext.
- Of independent interest, we also build a seed homomorphic PRG from the  $HSM_M$  assumption for the most efficient secure aggregation protocol we present in the context of federated learning.

**Applications:** Our motivating application is for Federated Learning (FL). The dynamic participation feature of OPA, crucial for federated learning, facilitates secure federated learning, where participants speak only once, streamlining the process significantly. Learning in the clear involves the client receiving the global update from the server and responding with *one* message corresponding to the new updates generated after the client has trained the model with its local data. In contrast, prior works[17, 13] involve 8 rounds, and the work of [67] requires 7 rounds in total, including the setup. Moreover, our protocols offer adaptive security. Our advantages extend beyond just round complexity. See Section 2.1 and Tables 1 and 2 for a detailed comparison of asymptotics.

**Implementation and Benchmarks:** Our contributions also extend beyond the theoretical domain. We implement OPA as a secure aggregation protocol and benchmark with several state-of-the-art solutions [17, 13, 52, 67]. Specifically, we first implement OPA based on the CL framework with  $M = p$ . Importantly, our server computation time scales the best with a larger number of inputs where our server takes  $< 1s$  for computation even for larger number of clients ( $n = 1000$ ). Meanwhile, our client running time is competitive for a small number of clients but offers significant gains for a larger number of clients.

Table 1: Total asymptotic computation cost for all rounds per aggregation with semi-honest security.  $n$  denotes the total number of users, with committee size  $m$  and  $L$  is the length of the input vector. The “Rounds” column indicates the number of rounds in the setup phase (on the left, if applicable) and in each aggregation iteration (on the right). A “round of communication” refers to a discrete step within a protocol during which a participant or a group of participants send messages to another participant or group of participants, and participants from the latter group must receive these messages before they can send their own messages in a subsequent round. “fwd” means that the server only forwards the messages from the users. The second column in the User Aggregation phase refers to the cost of the committee members.

Protocol	Rounds		Computation Cost					
			Server			User		
	Setup	Agg.	Setup	Agg.	Setup	Agg.		
BIK+17[17]	-	8	-	$O(n^2L)$	-	$O(n^2 + nL)$		
BBG+20[13]	-	8	-	$O(n \log^2 n + nL \log n)$	-	$O(\log^2 n + L \log n)$		
Flamingo[67]	4	3	fwd	$O(nL + n \log^2 n)$	$O(\log^2 n)$	$O(L + n \log n)$	$O(m^2 + n)$	
LERNA [62]	1	1	1	fwd	$O((\kappa + n)L + \kappa^2)$	$O(\kappa^2)$	$O(L)$	$O(L + n)$
SASH[64]	-	10	-	$O(L + n^2)$	-	$O(L + n^2)$		
OPA	-	1	1	-	$O(nL + m \log m)$	-	$O(L + m)$	$O(n)$
OPA'	-	1	1	-	$O(nL + L(m \log m))$	-	$O(L + mL)$	$O(nL)$

Then we benchmark OPA based on the LWR assumption.  $OPA_{LWR}$ , which is based on the seed homomorphic PRG, offers competitive server performance (with significant gains for larger  $L$ ) when compared to  $OPA_{CL}$ . Meanwhile,  $OPA_{LWR}$  outperforms the client performance of even  $OPA_{CL}$ . We also compare the performance of our LWR construction with a naive solution where each client secret-shares their inputs with the committee members. Our experiments show that  $OPA_{LWR}$  offers the best numbers for the committee members, for large  $L$ .

To further demonstrate the feasibility of our protocol, we train a binary classification model using logistic regression, in a federated manner, for two datasets. Our protocol carefully handles floating point values (using two different methods of quantization - multiplying by a global multiplier vs representing floating point numbers as a vector of integer values) and the resulting model is shown to offer performance close to that of simply learning in the clear. We also train a neural-network-based MLP classifier that was trained over popular machine learning datasets including MNIST, CIFAR-10, and CIFAR-100. More details can be found in Section 4.

## 2.1 Detailed Contributions in Federated Learning

Next, we compare our protocol with all efficient summation protocols listed in Table 3, with a specific focus on those that accommodate dynamic participation, a key feature shared by all federated learning methodologies. Asymptotically, we look at the performance of OPA. In Tables 1 and 2, we list the communication complexity, computational complexity, and round complexity per participant. Notably, our protocols are setup-free, eliminating any need for elaborate initialization procedures. Furthermore, they are characterized by a streamlined communication process, demanding just a single round of interaction from the participants.

**Asymptotic Comparison.** More concretely, based on Table 1, our approach stands out by significantly reducing the round complexity, ensuring that each participant’s involvement is limited to a single communication round i.e. each participant speaks only once. That is, users speak once and committee members speak once too. On the contrary previous works[17, 13]<sup>2</sup> require 8 rounds and the work of [67] requires 7 rounds in total, including the setup. This reduction in round complexity serves as a significant efficiency advantage.

Despite our advantage in the round complexity, our advantages extend beyond just round complexity (see Table 1). Notably, as the number of participants ( $n$ ) grows larger, our protocol excels in terms of computational complexity. While previous solutions exhibit complexities that are quadratic [17, 64] or linearithmic [67] in  $n$ , our approach maintains a logarithmic complexity for the users which is noteworthy when considering our protocol’s concurrent reduction in the number of communication

<sup>2</sup>[13] offer a weaker security definition from the other works: for some parameter  $\alpha$  between  $[0, 1]$ , honest inputs are guaranteed to be aggregated at most once with at least  $\alpha$  fraction of other inputs from honest users.

Table 2: Total received and sent asymptotic communication cost for all rounds per aggregation with semi-honest security.  $n$  denotes the total number of users,  $m$  is the size of the committee,  $k$  the security parameter,  $L$  the length of the input vector, and  $\ell$  the bit length of each element.

Protocol	Communication Cost				
	Server		User		
	Setup	Agg.	Setup	Agg.	
BIK+17[17]	-	$O(nL\ell + n^2k)$	-	$O(L\ell + nk)$	
BBG+20[13]	-	$O(nL\ell + nk \log n)$	-	$O(L\ell + k \log n)$	
Flamingo[67]	$O(k \log n)$	$O(nL\ell + nk \log^2 n)$	$O(k \log n)$	$O(L\ell + nk \log n)$	$O(m^2 + n \log n)$
LERNA [62]	-	$O(Lnk + m^2 \cdot k)$	$O(\kappa^2 k)$	$O(\kappa L + L \log n + k)$	$O(L\kappa + L \log n + k)$
SASH [64]	-	$O(nL\ell + \kappa^2 k)$	-	$O(L\ell + nk)$	
OPA	-	$O(Ln + m)$	-	$O(kL + m)$	$O(n)$
OPA'	-	$O(kL(n + m))$	-	$O(k(L + mL))$	$O(k(nL + L))$

rounds. Furthermore, our committee framework demonstrates a linear relationship with  $n$  for the committee members, a notable improvement compared to the linearithmic complexity and setup requirement in the case of [67] which considers a stateful set of decryptors (committee), as opposed to our stateless committee.<sup>3</sup>

When it comes to user communication and message sizes, previous solutions entail user complexities that either scale linearly [17, 64] or linearithmically [67] with the number of participants ( $n$ ) according to Table 2. However, in our case, user communication complexity is reduced to a logarithmic level. Our communication among committee members exhibits a linear relationship with  $n$ . Furthermore, as the number of users  $n$  increases, the communication load placed on the server is also effectively reduced in comparison to other existing protocols. That said, the above advantages underline the scalability and efficiency of our protocols in the federated learning context which typically requires a very large number of  $n$  and  $L$ .

The works of [17, 13, 67] address an active adversary that can provide false information regarding which users are online on behalf of the server, by requiring users to verify the signatures of the claimed online set. This approach introduces an additional two rounds into each protocol, resulting in 10 rounds in [17, 13] and 5 rounds in [67] with 10 rounds of setup. The setup communication complexity of [67] also increases to  $O(k \log^2 n)$ . In our work, we solve this issue without signatures. Note that prior work often required a two-thirds honest majority for the malicious setting. One can leverage this to instead introduce a gap between the reconstruction threshold  $r$  and corruption threshold  $t$ . Specifically, if  $r > (m + t)/2$ , then a malicious server can only reconstruct with respect to a unique set and any other information is purely random information, unhelpful to the server.

**Comparison with LERNA [62].** LERNA requires a fixed, stateful committee (like Flamingo) to secret share client keys, whereas we support smaller, dynamic stateless committees that can change in every round. Concretely, LERNA works by having *each* client (in the entire universe of clients, not just for that iteration) secret-share the keys with the committee. Consequently, LERNA's committee needs to be much larger ( $2^{14}$  members for  $\kappa = 40$  due to the number of shares they receive) and tolerate fewer dropouts, compared to our approach. Furthermore, LERNA's benchmarks assume 20K+ clients, while real-world deployments have 50-5000 clients per iteration. When the client count is low, the committee has to do significantly more work to handle and store the required large number of shares. That said, LERNA is not suitable for traditional FL applications. In the table, we use the same notations, as for LERNA to refer to committee size by utilizing  $\kappa$  in the committee calculations. LERNA could work for less than 16K parties but then the computation of committee members increases significantly as the number of parties decreases. Even concrete costs require the client to send 2GB of data during the setup phase (with 20K clients and  $L = 50,000$ ). Per iteration, the cost is 0.91 MB. For the same parameters, OPA<sub>CL</sub> (the construction based on distributed key-homomorphic PRF in CL framework), requires the client to send 5.6 MB (across both server and committee, per iteration) As a result, LERNA only becomes cost-effective after more than 400 iterations, during which it requires a fixed, stateful, and a large committee to stay alive.

<sup>3</sup>Flamingo [67] employed *decryptors* which were a random subset of clients chosen by the server to interact with it to remove masks from masked data that were sent by the larger set of clients. [67] lacks security guarantees in the event of collusion among all decryptors.

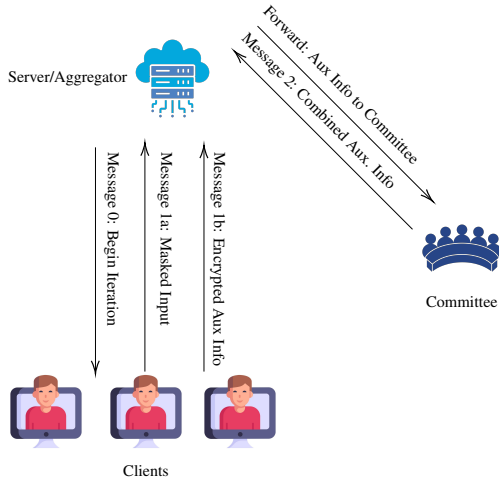


Figure 2: The OPA system model operates in iterations. Each iteration begins with the server sending a message to initiate the process (Message 0). In response, clients train the model on their local data, obtain updates, and mask the input. (Message 1): masked input is sent to the server, while auxiliary information is transmitted (via encryption) to the committee, via the server. Upon receiving the forwarded information, the committee members combine these into a single value. Finally, this consolidated data is sent to the server (Message 2), concluding the iteration. Section 3 presents a succinct overview of OPA.

**Comparison with Willow [14].** Willow, a concurrent and independent work, requires a static, stateful committee of clients. This is similar to LERNA but they can allow for much smaller sizes. Each committee member is expected to sample a key pair and provide the public key to the clients. Clients, with knowledge of this public key, can encrypt information to these committee members, via the server. The server coordinates with the committee members to later aggregate. Note that this is also similar to our setting where there exists a PKI, or some mechanism, to ensure that the client can encrypt to the committee member and then route the message through the server. Unlike OPA, Willow requires a one-round setup ceremony among the committee members to agree on the public key. Specifically, they need to agree on the public key for threshold encryption and each committee member holds a share of the decryption key. This setup phase implies either a static committee or performing the setup phase, at each iteration, to support a dynamic committee. Furthermore, to achieve security against a malicious server, they require another party (or a set of parties) called verifier(s). In essence, the verifier certifies that the server has not behaved maliciously. Furthermore, it appears that a verifier also needs to be stateful, across iterations, to ensure that the server does not replay an old ciphertext.

### 3 Preliminaries and Overview

**Notations.** For a distribution  $X$ , we use  $x \leftarrow_s X$  to denote that  $x$  is a random sample drawn from the distribution  $X$ . We denote by  $\mathbf{u}$  a vector and by  $\mathbf{A}$  a matrix. For a set  $S$  we use  $x \leftarrow_s S$  to denote that  $x$  is chosen uniformly at random from the set  $S$ . By  $[n]$  for some integer  $n$ , we denote the set  $\{1, \dots, n\}$ .

**Cryptographic Preliminaries.** This is deferred to Section C. In Section C.1 we discuss seed-homomorphic PRG and Section C.2, we discuss the definition and construction of secret-sharing schemes over a field and over integers. In Section C.3, we present lattice-based assumptions. This is followed by an exposition on the syntax and security definitions of various types of pseudorandom functions through Sections C.4 and C.5. Finally, we introduce the CL Framework in Section C.6.

#### 3.1 Overview of OPA

While we defer a rigorous exposition on the technical overview of our construction to Section B, we now present the intuition behind our protocol. The communication flow is described in Figure 2.






Figure 3: Client and Server Computation Time as a function of client count across different algorithms.

Every client needs to ensure the privacy of their input, therefore a client has to mask their input. In iteration  $\ell$ , if client  $i$  has input  $\mathbf{x}_{i,\ell}$ , then it chooses a mask of the same length to “add” to the inputs. Let the mask be  $\mathbf{mask}_{i,\ell}$  and the ciphertext is defined as  $\mathbf{ct}_{i,\ell} = \mathbf{x}_{i,\ell} + \mathbf{mask}_{i,\ell}$ . To ensure privacy, we need the mask to be chosen, uniformly at random, from a large distribution. Furthermore, by performing the addition with respect to a modulus  $M$ , we get the property that for a random  $\mathbf{mask}_{i,\ell}$ ,  $\mathbf{ct}_{i,\ell}$  is identically distributed to a random element from  $\mathbb{Z}_M$ . The client  $i$  sends  $\mathbf{ct}_{i,\ell}$  to the server. This is Message 1a in Figure 2.

The server, upon receiving the ciphertexts, can simply add up the ciphertexts. This leaves it with  $\sum_i \mathbf{ct}_{i,\ell} = \sum_i \mathbf{x}_{i,\ell} + \sum_i \mathbf{mask}_{i,\ell}$ . The goal of the server is to recover  $\sum_i \mathbf{x}_{i,\ell}$ . Therefore, it requires  $\sum_i \mathbf{mask}_{i,\ell}$  to complete the computation. In works on Private Stream Aggregation [81, 15, 56, 48], the assumption made is that  $\sum_i \mathbf{mask}_{i,\ell} = 0$ . However, this requires all the clients to participate which is a difficult requirement in federated learning. Instead, we enlist the help of the committee to provide the server with  $\sum_i \mathbf{mask}_{i,\ell}$ , for only the participating clients.

**Working with the Committee.** Each client  $i$ , therefore, has to communicate information about its respective  $\mathbf{mask}_{i,\ell}$  to the committee members. This is what we refer to as “Aux info”, and we route it through the server, to the committee member, and the information is encrypted under the public key of the respective committee member. This is Message 1b in Figure 2. This ensures that the server cannot recover the auxiliary information. Eventually, each committee member “combines” the aux info it has received to the server (Message 2 in Figure 2), with the guarantee that this is sufficient to reconstruct  $\sum_i \mathbf{mask}_{i,\ell}$ . Our key cryptographic technique here is known as secret-sharing which ensures that a secret  $s$  can be shared with a committee of  $m$  such that as long as  $r$  number of them participate, the server can learn the secret  $s$ . The security of the secret is guaranteed even if the server colludes with  $t$  committee members. We require  $t \leq m - 1$ . The additional property is that if a committee member receives shares of two secrets  $s_1, s_2$ . Then, adding up these shares will help reconstruct the secret  $s_1 + s_2$ .

**Optimizing Committee Performance.** The solution laid out above requires the clients to “secret-share”  $\mathbf{mask}_{i,\ell}$ . However, note that  $\mathbf{mask}_{i,\ell}$  is as long as  $\mathbf{x}_{i,\ell}$ . Therefore, each committee member will receive communication  $O(nL)$  where  $n$  is the number of clients and  $L$  is the length of the vector. It also has to perform computation proportional to  $O(nL)$ . OPA reduces the burden of the committee by introducing a succinct communication that is independent of  $L$  to the committee. We do so by relying on a structured pseudorandom generator (PRG) which is called as “seed-homomorphic PRG”. Recall that a PRG guarantees that one can send a small  $sd$  which can be later expanded to generate a large vector by computing  $\text{PRG}(sd)$  with the guarantee that a random  $sd$  implies that the output also looks random. A seed-homomorphic PRG ensures that  $\text{PRG}(sd_1 + sd_2) = \text{PRG}(sd_1) + \text{PRG}(sd_2)$ .

## 4 Experiments

We benchmark  $\text{OPA}_{\text{CL}}$  (Construction 12) and  $\text{OPA}_{\text{LWR}}$  (Construction 1). Recall that the former is based on threshold key-homomorphic PRF and where there is no leakage simulation needed, while




Figure 4: MLP Accuracy for different datasets: MNIST, CIFAR-10, and CIFAR-100




Figure 5: Measure of MCC as a function of number of iterations and number of clients, comparing  $\text{OPA}_{\text{CL}}$  and plaintext learning.

the latter is based on seed-homomorphic PRG. We run our experiments on an Apple M1 Pro CPU with 16 GB of unified memory, without any multi-threading or related parallelization. We use the ABIDES simulation [25] to simulate real-world network connections. ABIDES supports a latency model which is represented as a base delay and a jitter which controls the number of messages arriving within a specified time. Our base delay is set to the “global” setting in ABIDES’s default parameters (the range is 21 microseconds to 53 milliseconds), and use the default parameters for the jitter. This framework was used to measure the performance of other prior work including [67, 52].

- **Parameter Choices for  $\text{OPA}_{\text{LWR}}$ :**  $\text{OPA}_{\text{LWR}}$  is parametrized by  $\rho, q, p$ . We use the LWE-estimator [5] to estimate the security level. We follow the parameters similar to [48]. The value of  $1/p$  is the error rate  $\alpha$  in an LWE instance. Using the LWE estimator, we set  $\rho := 1024$ . We set  $q$  to match the field used for Shamir’s Secret Sharing, which is a 128-bit prime and set  $p = 2^{85}$ . The hardness estimated is  $2^{129}$ . We do use Packed Secret Sharing, for benchmarking the server and client computation cost. We pack it to 64, i.e., each committee member receives 64 shares. For reconstruction to hold we set  $m = 50, r = 34$ . Assuming  $\eta = \delta = 0.01$ , we can estimate that setting  $m = 50$ , then the  $\Pr[\eta_C + \delta_C > 1/3] \leq 5 \cdot 10^{-5}$ . This also satisfies the requirements of packed secret sharing which will require that  $m \geq 3/2 \cdot 1024/64$ . Recall that, for packed secret sharing, we can set the corruption threshold to be  $r - \rho$  where  $\rho$  is the number of secrets being packed. Here  $\rho$  is 16. Note that this construction achieves committee performance independent of the vector length and is the preferred one.
- **Parameter Choices for  $\text{OPA}_{\text{CL}}$ :** We rely on the BICYCL library [20] and use pybind to convert the C++ code to Python. Our implementation will assume that the plaintext space is  $\mathbb{Z}_p$  for a prime  $p$ . Our experiments will assume that  $m = \log n$ .

**Microbenchmarking Secure Aggregation.** Our first series of experiments is to run  $\text{OPA}_{\text{CL}}$ ,  $\text{OPA}_{\text{LWR}}$  to build a secure aggregation protocol for  $L = 1$ . We also compare with existing work including [17, 13, 52, 67]. We vary the offline rates, and the ability to group clients, along with increasing the number of clients to study the performance of related work. Recall that the offline rate (denoted by  $\eta$ ) controls the number of clients who do not participate, despite being selected. Meanwhile, we denote by  $g$  the size of the neighborhood or group. For [52], we set the input size to be bounded by  $10^4$ . Recall that [52] does not have efficient aggregate recovery and requires input bounding. Also, [52] incurs a setup/offline client computation time of nearly 30ms, even for 100 clients. Our implementation sets  $m = \lceil \log(n) \rceil$ . As can be seen from Figure 3. The key takeaway from our experiments is that our client and server performance outperforms all prior works.



**Performance of  $\text{OPA}_{\text{CL}}$ .**  $\text{OPA}_{\text{CL}}$ 's server running time is less than 1 second - owing to a single round protocol with support of efficient recovery of the aggregate. This outperforms all existing protocols. Each committee member computes  $< 1ms$ .  $\text{OPA}_{\text{CL}}$  assumes concurrent client communication to committee and server. Note that  $\text{OPA}_{\text{CL}}$  scales linearly with the size of the input.

**Performance of  $\text{OPA}_{\text{LWR}}$ .** It is to be noted that the server performance matches that of  $\text{OPA}_{\text{CL}}$ 's server performance. However, when aggregating longer vectors, which is the use-case for OPA,  $\text{OPA}_{\text{LWR}}$  significantly outperforms  $\text{OPA}_{\text{CL}}$ , with the key savings coming from having to perform no exponentiations. Meanwhile, the client's performance in  $\text{OPA}_{\text{LWR}}$  is much better than  $\text{OPA}_{\text{CL}}$ , again due to the lack of exponentiations. The communication cost for  $\text{OPA}_{\text{LWR}}$  are as follows:

- Client  $i$  to Server:  $L$  field elements
- Client  $i$  to Committee:  $k \cdot m$  field elements where  $k$  is the number of elements being shared. With naive secret sharing, this would be 1024. Instead, when we pack into 64 different polynomials,  $k = 64$ .
- Committee  $j$  to Server:  $k$  field elements
- Total Sent/Received per Client:  $L + k \cdot m$  field elements
- Total Sent/Received by Server:  $nL + mk$  field elements
- Total Sent/Received per Committee Member:  $n \cdot k + k$  field elements

In Section D, we describe the naive secret-sharing-based construction and also study the computation cost of  $\text{OPA}_{\text{LWR}}$  with respect to this naive construction by micro-benchmarking the running times. Our experiments indicate that the running time of the committee member is the best with  $\text{OPA}_{\text{LWR}}$ , when compared to the naive solution. This is in addition to the improvement in the communication cost. Recall that the motivating use case for avoiding the naive solution is improved performance of the committee members, who are often other clients.

**Benchmarking FL Models.** To demonstrate  $\text{OPA}_{\text{CL}}$ 's viability for federated learning, we train a logistic regression model on two skewed datasets. We show that  $\text{OPA}_{\text{CL}}$  performs very close to learning in the clear, indicating feasibility for machine learning. As our goal was to show feasibility, our experiments use one committee member for all  $n$  clients. We vary  $n$  and the number of iterations for model convergence, measuring accuracy and Matthew's Correlation Coefficient (MCC) [68] which better evaluates binary classification with unbalanced classes.

- **Adult Census Dataset:** We first run experiments on the adult census income dataset from [26, 55] to predict if an individual earns over \$50,000 per year. The preprocessed dataset has 105 features and 45,222 records with a 25% positive class. We randomly split into training and testing, with further splitting by the clients. First, we train in the clear with weights sent to the server to aggregate. With 100 clients and 50 iterations, we achieve 82.85% accuracy and 0.51 MCC. We repeat with  $\text{OPA}_{\text{CL}}$ , one committee member, and 100 clients. With 10 iterations, we achieve 82.38% accuracy and 0.48 MCC. With 20 iterations, we achieve 82% accuracy and 0.51 MCC. Our quantization technique divides weights into integer and decimal parts (2 integer and 8 decimal values per weight). Training with 50 clients takes under 1 minute per client per iteration with no accuracy loss. This quantization yields a vector size of 1050 (10 per feature).
- We use the Kaggle Credit Card Fraud dataset [74], comprising 26 transformed principal components and amount and time features. We omit time and use the raw amount, adding an intercept. The goal is to predict if a transaction was indeed fraudulent or not. There are 30 features and 284,807 rows, with  $< 0.2\%$  fraudulent. Weights are multiplied by 10,000 and rounded to an integer, accounted for in aggregation. Figure 5 shows  $\text{OPA}_{\text{CL}}$ 's MCC versus clear learning for varying clients and iterations. With the accuracy multiplier,  $\text{OPA}_{\text{CL}}$ 's MCC is very close to clear learning and even outperforms sometimes. The highly unbalanced dataset demonstrates  $\text{OPA}_{\text{CL}}$  can achieve strong performance even in challenging real-world scenarios.
- We then train a vanilla multi-layer perceptron (MLP) classifier on three datasets: MNIST, CIFAR-10, CIFAR-100. We quantize the weights by multiplying with  $2^{16}$ . The MLP accuracy, as a function of the iteration count, is plotted in Figure 4. Our experiments demonstrate that  $\text{OPA}_{\text{CL}}$  preserves accuracy while ensuring the privacy of client data. Note that vanilla MLP classifiers do not typically offer good performance for CIFAR datasets, but note that the goal of our experiments was to show that  $\text{OPA}_{\text{CL}}$  does not impact accuracy.

## 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, *Advances in Cryptology – CRYPTO 2022, Part IV*, volume 13510 of *Lecture Notes in Computer Science*, pages 421–452. Springer, Cham, August 2022.
- [2] Anasuya Acharya, Carmit Hazay, Vladimir Kolesnikov, and Manoj Prabhakaran. SCALES - MPC with small clients and larger ephemeral servers. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part II*, volume 13748 of *Lecture Notes in Computer Science*, pages 502–531. Springer, 2022.
- [3] Surya Addanki, Kevin Garbe, Eli Jaffe, Rafail Ostrovsky, and Antigoni Polychroniadou. Prio+: Privacy preserving aggregate statistics via boolean shares. In Clemente Galdi and Stanislaw Jarecki, editors, *Security and Cryptography for Networks - 13th International Conference, SCN 2022, Amalfi, Italy, September 12-14, 2022, Proceedings*, volume 13409 of *Lecture Notes in Computer Science*, pages 516–539. Springer, 2022.
- [4] Shweta Agrawal, Xavier Boyen, Vinod Vaikuntanathan, Panagiotis Voulgaris, and Hoeteck Wee. Functional encryption for threshold functions (or fuzzy ibe) from lattices. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012: 15th International Conference on Theory and Practice of Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 280–297. Springer, Berlin, Heidelberg, May 2012.
- [5] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Math. Cryptol.*, 9(3):169–203, 2015.
- [6] Apple and Google. Exposure notification privacy-preserving analytics (ENPA), 2021.
- [7] Arasu Arun, Chaya Ganesh, Satya V. Lokam, Tushar Mopuri, and Sriram Sridhar. Dew: A transparent constant-sized polynomial commitment scheme. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023: 26th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 13941 of *Lecture Notes in Computer Science*, pages 542–571. Springer, Cham, May 2023.
- [8] Thomas Attema, Ignacio Cascudo, Ronald Cramer, Ivan Bjerre Damgård, and Daniel Escudero. Vector commitments over rings and compressed  $\sigma$ -protocols. Cryptology ePrint Archive, Report 2022/181, 2022.
- [9] Abhishek Banerjee, Chris Peikert, and Alon Rosen. Pseudorandom functions and lattices. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EURO-CRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 719–737. Springer, Berlin, Heidelberg, April 2012.
- [10] Laasya Bangalore, Mohammad Hossein Faghihi Sereshgi, Carmit Hazay, and Muthuramakrishnan Venkatasubramanian. Flag: A framework for lightweight robust secure aggregation. In Joseph K. Liu, Yang Xiang, Surya Nepal, and Gene Tsudik, editors, *ASIACCS 23: 18th ACM Symposium on Information, Computer and Communications Security*, pages 14–28. ACM Press, July 2023.
- [11] Daniela Becker, Jorge Guajardo, and Karl-Heinz Zimmermann. Revisiting private stream aggregation: Lattice-based PSA. In *ISOC Network and Distributed System Security Symposium – NDSS 2018*. The Internet Society, February 2018.
- [12] James Bell, Adrià Gascón, Tancrede Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and Cathie Yun. ACORN: Input validation for secure aggregation. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 4805–4822, Anaheim, CA, August 2023. USENIX Association.
- [13] James Henry Bell, Kallista A Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 1253–1269, 2020.

- [14] James Bell-Clark, Adrià Gascón, Baiyu Li, Mariyana Raykova, Phillipp Schoppmann, and Pierre Tholoniati. Willow: Secure aggregation with one-shot clients. Theory and Practice of Multi-Party Computation Workshop, 2024.
- [15] Fabrice Benhamouda, Marc Joye, and Benoît Libert. A new framework for privacy-preserving aggregation of time-series data. *ACM Trans. Inf. Syst. Secur.*, 18(3), mar 2016.
- [16] Alexander Bienstock, Daniel Escudero, and Antigoni Polychroniadou. On linear communication complexity for (maximally) fluid MPC. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 263–294. Springer, Cham, August 2023.
- [17] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages 1175–1191. ACM, 2017.
- [18] Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 561–586. Springer, Cham, August 2019.
- [19] Dan Boneh, Kevin Lewi, Hart William Montgomery, and Ananth Raghunathan. Key homomorphic PRFs and their applications. In Ran Canetti and Juan A. Garay, editors, *Advances in Cryptology – CRYPTO 2013, Part I*, volume 8042 of *Lecture Notes in Computer Science*, pages 410–428. Springer, Berlin, Heidelberg, August 2013.
- [20] Cyril Bouvier, Guilhem Castagnos, Laurent Imbert, and Fabien Laguillaumie. I want to ride my BICYCL : BICYCL implements cryptography in class groups. *J. Cryptol.*, 36(3):17, 2023.
- [21] Lennart Braun, Ivan Damgård, and Claudio Orlandi. Secure multiparty computation from threshold encryption based on class groups. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 613–645. Springer, Cham, August 2023.
- [22] Joakim Brorsson and Martin Gunnarsson. Dipsauce: Efficient private stream aggregation without trusted parties. To Appear in NordSec23, 2023.
- [23] Johannes Buchmann and Hugh C. Williams. A key-exchange system based on imaginary quadratic fields. *Journal of Cryptology*, 1(2):107–118, June 1988.
- [24] Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in Computer Science*, pages 677–706. Springer, Cham, May 2020.
- [25] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. Abides: Towards high-fidelity multi-agent market simulation. In *Proceedings of the 2020 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, SIGSIM-PADS '20, page 11–22, New York, NY, USA, 2020. Association for Computing Machinery.
- [26] David Byrd, Vaikkunth Mugunthan, Antigoni Polychroniadou, and Tucker Balch. Collusion resistant federated learning with oblivious distributed differential privacy. In *Proceedings of the Third ACM International Conference on AI in Finance, ICAIF '22*, page 114–122, New York, NY, USA, 2022. Association for Computing Machinery.
- [27] Ignacio Cascudo and Bernardo David. SCRAPE: Scalable randomness attested by public entities. In Dieter Gollmann, Atsuko Miyaji, and Hiroaki Kikuchi, editors, *ACNS 17: 15th International Conference on Applied Cryptography and Network Security*, volume 10355 of *Lecture Notes in Computer Science*, pages 537–556. Springer, Cham, July 2017.

- [28] Ignacio Cascudo and Bernardo David. Publicly verifiable secret sharing over class groups and applications to DKG and YOSO. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part V*, volume 14655 of *Lecture Notes in Computer Science*, pages 216–248. Springer, Cham, May 2024.
- [29] Guilhem Castagnos, Dario Catalano, Fabien Laguillaumie, Federico Savasta, and Ida Tucker. Two-party ECDSA from hash proof systems and efficient instantiations. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 191–221. Springer, Cham, August 2019.
- [30] Guilhem Castagnos, Dario Catalano, Fabien Laguillaumie, Federico Savasta, and Ida Tucker. Bandwidth-efficient threshold EC-DSA. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020: 23rd International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 12111 of *Lecture Notes in Computer Science*, pages 266–296. Springer, Cham, May 2020.
- [31] Guilhem Castagnos, Laurent Imbert, and Fabien Laguillaumie. Encryption switching protocols revisited: Switching modulo  $p$ . In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 255–287. Springer, Cham, August 2017.
- [32] Guilhem Castagnos and Fabien Laguillaumie. Linearly homomorphic encryption from DDH. In Kaisa Nyberg, editor, *Topics in Cryptology – CT-RSA 2015*, volume 9048 of *Lecture Notes in Computer Science*, pages 487–505. Springer, Cham, April 2015.
- [33] Guilhem Castagnos, Fabien Laguillaumie, and Ida Tucker. Practical fully secure unrestricted inner product functional encryption modulo  $p$ . In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 733–764. Springer, Cham, December 2018.
- [34] Pyrros Chaidos and Geoffroy Couteau. Efficient designated-verifier non-interactive zero-knowledge proofs of knowledge. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*, volume 10822 of *Lecture Notes in Computer Science*, pages 193–221. Springer, Cham, April / May 2018.
- [35] Jung Hee Cheon, Dongwoo Kim, Duhyeong Kim, Joohee Lee, Junbum Shin, and Yongsoo Song. Lattice-based secure biometric authentication for hamming distance. In Joonsang Baek and Sushmita Ruj, editors, *ACISP 21: 26th Australasian Conference on Information Security and Privacy*, volume 13083 of *Lecture Notes in Computer Science*, pages 653–672. Springer, Cham, December 2021.
- [36] Jihoon Cho, Jincheol Ha, Seongkwang Kim, ByeongHak Lee, Joohee Lee, Jooyoung Lee, Dukjae Moon, and Hyojin Yoon. Transciphering framework for approximate homomorphic encryption. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part III*, volume 13092 of *Lecture Notes in Computer Science*, pages 640–669. Springer, Cham, December 2021.
- [37] Kevin Choi, Aathira Manoj, and Joseph Bonneau. SoK: Distributed randomness beacons. In *2023 IEEE Symposium on Security and Privacy*, pages 75–92. IEEE Computer Society Press, May 2023.
- [38] Arka Rai Choudhuri, Aarushi Goel, Matthew Green, Abhishek Jain, and Gabriel Kaptchuk. Fluid MPC: Secure multiparty computation with dynamic participants. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part II*, volume 12826 of *Lecture Notes in Computer Science*, pages 94–123, Virtual Event, August 2021. Springer, Cham.
- [39] Jean-Sébastien Coron. On the exact security of full domain hash. In Mihir Bellare, editor, *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer, Berlin, Heidelberg, August 2000.
- [40] Henry Corrigan-Gibbs. Privacy-preserving firefox telemetry with prio, 2020.

- [41] 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.
- [42] Geoffroy Couteau, Dahmun Goudarzi, Michael Klooß, and Michael Reichle. Sharp: Short relaxed range proofs. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 609–622. ACM Press, November 2022.
- [43] Geoffroy Couteau, Michael Klooß, Huang Lin, and Michael Reichle. Efficient range proofs with transparent setup from bounded integer commitments. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 247–277. Springer, Cham, October 2021.
- [44] 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: 9th International Conference on Theory and Practice of Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 75–90. Springer, Berlin, Heidelberg, April 2006.
- [45] Yi Deng, Shunli Ma, Xinxuan Zhang, Hailong Wang, Xuyang Song, and Xiang Xie. Promise  $\Sigma$ -protocol: How to construct efficient threshold ECDSA from encryptions based on class groups. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part IV*, volume 13093 of *Lecture Notes in Computer Science*, pages 557–586. Springer, Cham, December 2021.
- [46] Nico Döttling, Dimitris Kolonelos, Russell W. F. Lai, Chuanwei Lin, Giulio Malavolta, and Ahmadreza Rahimi. Efficient laconic cryptography from learning with errors. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part III*, volume 14006 of *Lecture Notes in Computer Science*, pages 417–446. Springer, Cham, April 2023.
- [47] Drand. Drand/drاند: a distributed randomness beacon daemon - go implementation.
- [48] Johannes Ernst and Alexander Koch. Private stream aggregation with labels in the standard model. *Proc. Priv. Enhancing Technol.*, 2021(4):117–138, 2021.
- [49] Matthew K. Franklin and Moti Yung. Communication complexity of secure computation (extended abstract). In *24th Annual ACM Symposium on Theory of Computing*, pages 699–710. ACM Press, May 1992.
- [50] Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 640–658. Springer, 2014.
- [51] Noemi Glaeser, Matteo Maffei, Giulio Malavolta, Pedro Moreno-Sanchez, Erkan Tairi, and Sri AravindaKrishnan Thyagarajan. Foundations of coin mixing services. *Cryptology ePrint Archive*, Report 2022/942, 2022.
- [52] Yue Guo, Antigoni Polychroniadou, Elaine Shi, David Byrd, and Tucker Balch. MicroFedML: Privacy preserving federated learning for small weights. *Cryptology ePrint Archive*, Report 2022/714, 2022.
- [53] Shai Halevi, Yuval Ishai, Eyal Kushilevitz, and Tal Rabin. Additive randomized encodings and their applications. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology - CRYPTO 2023 - 43rd Annual International Cryptology Conference, CRYPTO 2023, Santa Barbara, CA, USA, August 20-24, 2023, Proceedings, Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 203–235. Springer, 2023.
- [54] Shai Halevi, Yehuda Lindell, and Benny Pinkas. Secure computation on the web: Computing without simultaneous interaction. In Phillip Rogaway, editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011*.

- Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 132–150. Springer, 2011.
- [55] Bargav Jayaraman, Lingxiao Wang, David Evans, and Quanquan Gu. Distributed learning without distrust: privacy-preserving empirical risk minimization. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 6346–6357, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [56] Marc Joye and Benoît Libert. A scalable scheme for privacy-preserving aggregation of time-series data. In Ahmad-Reza Sadeghi, editor, *FC 2013: 17th International Conference on Financial Cryptography and Data Security*, volume 7859 of *Lecture Notes in Computer Science*, pages 111–125. Springer, Berlin, Heidelberg, April 2013.
- [57] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista A. Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaïd Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrede Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Hang Qi, Daniel Ramage, Ramesh Raskar, Mariana Raykova, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *Found. Trends Mach. Learn.*, 14(1-2):1–210, 2021.
- [58] Aniket Kate, Easwar Vivek Mangipudi, Pratyay Mukherjee, Hamza Saleem, and Sri Aravinda Krishnan Thyagarajan. Non-interactive VSS using class groups and application to DKG. Cryptology ePrint Archive, Report 2023/451, 2023.
- [59] Duhyeong Kim, Dongwon Lee, Jinyeong Seo, and Yongsoo Song. Toward practical lattice-based proof of knowledge from hint-MLWE. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 549–580. Springer, Cham, August 2023.
- [60] Russell W. F. Lai and Giulio Malavolta. Subvector commitments with application to succinct arguments. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 530–560. Springer, Cham, August 2019.
- [61] Iraklis Leontiadis, Kaoutar Elkhiyaoui, and Refik Molva. Private and dynamic time-series data aggregation with trust relaxation. In Dimitris Gritzalis, Aggelos Kiayias, and Ioannis G. Askoxylakis, editors, *CANS 14: 13th International Conference on Cryptology and Network Security*, volume 8813 of *Lecture Notes in Computer Science*, pages 305–320. Springer, Cham, October 2014.
- [62] Hanjun Li, Huijia Lin, Antigoni Polychroniadou, and Stefano Tessaro. LERNA: Secure single-server aggregation via key-homomorphic masking. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part I*, volume 14438 of *Lecture Notes in Computer Science*, pages 302–334. Springer, Singapore, December 2023.
- [63] Helger Lipmaa. Secure accumulators from euclidean rings without trusted setup. In Feng Bao, Pierangela Samarati, and Jianying Zhou, editors, *ACNS 12: 10th International Conference on Applied Cryptography and Network Security*, volume 7341 of *Lecture Notes in Computer Science*, pages 224–240. Springer, Berlin, Heidelberg, June 2012.
- [64] 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.



- [65] Hidde Lycklama, Lukas Burkhalter, Alexander Viand, Nicolas Küchler, and Anwar Hithnawi. RoFL: Robustness of secure federated learning. In *2023 IEEE Symposium on Security and Privacy*, pages 453–476. IEEE Computer Society Press, May 2023.
- [66] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010.
- [67] 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 *2023 IEEE Symposium on Security and Privacy*, pages 477–496. IEEE Computer Society Press, May 2023.
- [68] B.W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2):442–451, 1975.
- [69] K. MCCURLEY. Cryptographic key distribution and computation in class groups. *Proceedings of NATO ASI Number Theory and applications*, pages 459–479, 1989.
- [70] Moni Naor, Benny Pinkas, and Omer Reingold. Distributed pseudo-random functions and KDCs. In Jacques Stern, editor, *Advances in Cryptology – EUROCRYPT’99*, volume 1592 of *Lecture Notes in Computer Science*, pages 327–346. Springer, Berlin, Heidelberg, May 1999.
- [71] Dinh Duy Nguyen, Duong Hieu Phan, and David Pointcheval. Verifiable decentralized multi-client functional encryption for inner product. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part V*, volume 14442 of *Lecture Notes in Computer Science*, pages 33–65. Springer, Singapore, December 2023.
- [72] Dario Pasquini, Danilo Francati, and Giuseppe Ateniese. Eluding secure aggregation in federated learning via model inconsistency. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 2429–2443. ACM Press, November 2022.
- [73] 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.
- [74] Andrea Dal Pozzolo, Olivier Caelen, Reid A. Johnson, and Gianluca Bontempi. Calibrating probability with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence*, pages 159–166, 2015.
- [75] Mayank Raikwar and Danilo Gligoroski. SoK: Decentralized randomness beacon protocols. In Khoa Nguyen, Guomin Yang, Fuchun Guo, and Willy Susilo, editors, *ACISP 22: 27th Australasian Conference on Information Security and Privacy*, volume 13494 of *Lecture Notes in Computer Science*, pages 420–446. Springer, Cham, November 2022.
- [76] Mayank Rathee, Conghao Shen, Sameer Wagh, and Raluca Ada Popa. Elsa: Secure aggregation for federated learning with malicious actors. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 1961–1979, 2023.
- [77] Mayank Rathee, Conghao Shen, Sameer Wagh, and Raluca Ada Popa. ELSA: Secure aggregation for federated learning with malicious actors. In *2023 IEEE Symposium on Security and Privacy*, pages 1961–1979. IEEE Computer Society Press, May 2023.
- [78] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), sep 2009.
- [79] István András Seres, Péter Burcsi, and Péter Kutas. How (not) to hash into class groups of imaginary quadratic fields? *Cryptology ePrint Archive*, Report 2024/034, 2024.
- [80] Adi Shamir. How to share a secret. *Communications of the Association for Computing Machinery*, 22(11):612–613, November 1979.

- [81] Elaine Shi, T.-H. Hubert Chan, Eleanor G. Rieffel, Richard Chow, and Dawn Song. Privacy-preserving aggregation of time-series data. In *ISOC Network and Distributed System Security Symposium – NDSS 2011*. The Internet Society, February 2011.
- [82] Victor Shoup. Practical threshold signatures. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 207–220. Springer, Berlin, Heidelberg, May 2000.
- [83] Jonathan Takeshita, Zachariah Carmichael, Ryan Karl, and Taeho Jung. TERSE: tiny encryptions and really speedy execution for post-quantum private stream aggregation. In Fengjun Li, Kaitai Liang, Zhiqiang Lin, and Sokratis K. Katsikas, editors, *Security and Privacy in Communication Networks - 18th EAI International Conference, SecureComm 2022, Virtual Event, October 2022, Proceedings*, volume 462 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 331–352. Springer, 2022.
- [84] Jonathan Takeshita, Ryan Karl, Ting Gong, and Taeho Jung. SLAP: Simple lattice-based private stream aggregation protocol. Cryptology ePrint Archive, Report 2020/1611, 2020.
- [85] Jonathan Takeshita, Ryan Karl, Ting Gong, and Taeho Jung. SLAP: Simpler, improved private stream aggregation from ring learning with errors. *Journal of Cryptology*, 36(2):8, April 2023.
- [86] Sri Aravinda Krishnan Thyagarajan, Guilhem Castagnos, Fabien Laguillaumie, and Giulio Malavolta. Efficient CCA timed commitments in class groups. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2663–2684. ACM Press, November 2021.
- [87] Ida Tucker. *Functional encryption and distributed signatures based on projective hash functions, the benefit of class groups*. Theses, Université de Lyon, October 2020.
- [88] Hendrik Waldner, Tilen Marc, Miha Stopar, and Michel Abdalla. Private stream aggregation from labeled secret sharing schemes. Cryptology ePrint Archive, Report 2021/081, 2021.
- [89] Benjamin Wesolowski. Efficient verifiable delay functions. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages 379–407. Springer, Cham, May 2019.
- [90] Benjamin Wesolowski. Efficient verifiable delay functions. *Journal of Cryptology*, 33(4):2113–2147, October 2020.
- [91] Tsz Hon Yuen, Handong Cui, and Xiang Xie. Compact zero-knowledge proofs for threshold ECDSA with trustless setup. In Juan Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 12710 of *Lecture Notes in Computer Science*, pages 481–511. Springer, Cham, May 2021.
- [92] Yulin Zhao, Hualin Zhou, and Zhiguo Wan. SuperFL: Privacy-preserving federated learning with efficiency and robustness. Cryptology ePrint Archive, Report 2024/081, 2024.

**Disclaimer.** This paper was prepared for informational purposes by the Artificial Intelligence Research group of JPMorgan Chase & Co and its affiliates (“J.P. Morgan”) and is not a product of the Research Department of J.P. Morgan. J.P. 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.

## A Related Work

**Multi-Round Private Summation.** We begin by revisiting the concept outlined in [54]. The first multi-round secure aggregation protocol, designed to enable a single server to learn the sum of inputs  $x_1, \dots, x_n$  while hiding each input  $x_i$  is based on the following idea. Each user  $i$  adds a mask  $r_i$  to their private input  $x_i$ . This mask remains hidden from both the server and all other users it exhibits the property of canceling out when combined with all the other masks, i.e.,  $\sum_{i \in [n]} r_i = 0$ . Subsequently, each user forwards  $X_i = x_i + r_i$  to the server. By aggregating all the  $X_i$  values, the server is then able to determine the sum of all  $x_i$ . More specifically, to generate these masks, a common key  $k_{ij} = k_{ji} = \text{PRG}(g^{s_i s_j})$  is established by every pair of clients  $i, j$ . Here,  $g^{s_i}$  serves as an ephemeral “public key” associated with each client  $i \in [n]$ . This public key is shared with all other clients during an initial round, facilitated through the server. Importantly, the value  $s_i$  remains secret by each client  $i$ . Then, each client  $i \in [n]$  computes the mask  $r_i = \sum_{j < i} k_{ij} - \sum_{j > i} k_{ij}$  and due to the cancellation property the server outputs  $\sum_i X_i = \sum_i x_i$ . In this protocol, users are required to engage in multiple rounds of communication, where each user communicates more than once. Moreover, the protocol does not permit users to drop out from an aggregation iteration.

**Non-Interactive Private Summation with trusted setup.** If we were to require users to communicate only once in a protocol iteration, we encounter the challenge of mitigating residual attacks. In a prior study conducted by [81], a solution based on DDH was proposed to mitigate residual attacks by involving a trusted setup that assumes the generation of the common keys  $k_{ij} = k_{ji}$  into the protocol. However, it is important to highlight that this particular setup lacked the necessary mechanisms to accommodate dropouts and facilitate dynamic participation for multiple aggregation iterations. An additional limitation of this construction is the necessity of establishing a trusted setup that can be utilized across multiple iterations. Furthermore, to ensure that the server is unable to recover the masking key given a client’s masked inputs, the work relies on the DDH Assumption. An unfortunate consequence is that the server has to compute the discrete logarithm to recover the aggregate, a computationally expensive operation, particularly when dealing with large exponents. Numerous other works within this framework have emerged, each relying on distinct assumptions, effectively sidestepping the requirement for laborious discrete logarithm calculations. These include works based on the DCR assumption [56, 15], and lattice-based cryptography [11, 48, 84, 83, 88].

**Multi-Round Private Summation without Trusted Setup.** A separate line of research endeavors to eliminate the necessity for a trusted setup by introducing multi-round decentralized reusable setups designed to generate masks while adhering to the crucial cancellation property. However, akin to the previously mentioned approaches, these protocols come with a caveat—they do not accommodate scenarios involving dropouts or dynamic user participation across multiple iterations. Dipsauce [22] is the first to formally introduce a definition for a distributed setup PSA with a security model based on  $k$ -regular graph, non-interactive key exchange protocols, and a distributed randomness beacon [37, 47, 75] to build their distributed setup PSA. Meanwhile, the work of Nguyen *et al.* [71], assuming a PKI (or a bulletin board where all the public keys are listed), computed the required Diffie-Hellman keys on the fly to then build a one-time decentralized sum protocol which allowed the server to sum up the inputs *one-time*, with their construction relying on class group-based cryptography. To facilitate multiple iterations of such an aggregation, they combined their one-time decentralized sum protocol with Multi-client Functional Encryption (MCFE) to build a privacy-preservation summation protocol that can work over multiple rounds, without requiring a trusted setup and merely requiring a PKI. Unfortunately, per iteration, the clients need to be informed of the set of users participating in that round and unfortunately, they cannot drop out once chosen.

**Non-Interactive Private Summation with a collector.** To circumvent the need for a trusted setup and multi-round decentralized arrangements, an approach is presented in the work of [61] which introduces an additional server known as the “collector”. The fundamental premise here is to ensure that the collector and the evaluation server do not collude, thus effectively mitigating the risks associated with residual attacks. This protocol does allow dynamic participation and dropouts per iteration.

**Multi-Round Private Summation with Dynamic Participation (aka Secure Aggregation).** Secure aggregation of users’ private data with the aid of a server has been well-studied in the context of federated learning. Given the iterative nature of federated learning, dynamic participation is crucial. It enables seamless integration of new parties and those chosen to participate in various learning iterations, while also addressing the challenge of accommodating parties that may drop out during the aggregation phase due to communication failures or delays. Furthermore, an important problem in federated learning with user-constrained computation and wireless network resources is the computation and communication overhead which wastes bandwidth, increases training time, and can even impact the model accuracy if many users drop out. Seminal contributions by Bonawitz *et al.* [17] and Bell *et al.* [13] have successfully proposed secure aggregation protocols designed to cater to a large number of users while addressing the dropout challenge in a federated learning setting. However, it’s important to note that these protocols come with a notable drawback—substantial round complexity and overhead are incurred during each training iteration. Even in the extensive corpus of research based on more complex cryptographic machinery (see [57] for a plethora of previous works) such as threshold additive homomorphic encryption etc., these persistent drawbacks continue to pose challenges. Notably, all follow up works [13, 62, 52, 67, 12, 64, 65] of [17] require multiple rounds of interaction based on distributed setups. Secure aggregation protocols, with their adaptable nature, hold relevance across a wide array of domains. They are applicable in various scenarios, including ensuring the security of voting processes, safeguarding privacy in browser telemetry as illustrated in [40], and facilitating data analytics for digital contact tracing, as seen in [6] besides enabling secure federated learning.

It is also important to note that some of these works - ACORN [12] and RoFL [65] build on top of the works of [17, 13] to tackle the problem of “input validation” using efficient zero-knowledge proof systems. The goal is for the clients to prove that the inputs being encrypted are “well-formed” to prevent poisoning attacks. RoFL allows for detection when a malicious client misbehaves, while ACORN presents a *non-constant* round protocol to identify and remove misbehaving clients. We leave it as a direction for future research on how to augment our protocol to also support input validation. The above discussion is also summarized in Table 3, by looking at four properties (a) whether the aggregate can be efficiently recovered, (b) whether it allows dynamic participation, (c) whether it requires trusted setup or multi-round distributed setup, and (d) the security assumptions.

For completeness, we also compare with other communication models that bear similarities with OPA in Section A.1.

### A.1 OPA vs Other Communication/Computation Models

**Shuffle Model.** Note that our model bears similarities to the shuffle model, in which clients dispatch input encodings to a shuffler or a committee of servers responsible for securely shuffling before the data reaches the server for aggregation as in the recent work of Halevi *et al.* [53]. Nonetheless, it is important to note that such protocols typically entail multiple rounds among the committee servers to facilitate an efficient and secure shuffle protocol.

**Multi-Server Secure Aggregation Protocols.** It’s worth emphasizing that multi-server protocols, as documented in [50, 41, 3, 76, 92], have progressed to a point where their potential standardization by the IETF, as mentioned in [73], is indeed noteworthy. In the multi-server scenario, parties can directly share their inputs securely among a set of servers, which then collaborate to achieve secure aggregation. Some of the works in this domain include two-server solutions Elsa [77] and SuperFL [92] or the generic multi-server solution Flag [10]. Unfortunately, in the case of federated learning, which involves handling exceptionally long inputs, the secret-sharing approach becomes impractical due to the increase in communication complexity associated with each input. Furthermore, these servers are required to have heavy computation power and be stateful (retaining data/state from iteration to iteration). Jumping ahead, in our protocol the ephemeral parties are neither stateful nor require heavy computation.

Table 3: Comparison of Various Private Summation Protocols. **TD** stands for trusted dealer/trusted setup, **DS** stands for multi-round distributed setup. Note that **DS** implies several rounds of interaction while our protocol does not require any interaction. Here, efficient aggregate recovery refers to whether the server can recover the aggregate efficiently. For example, [81, 15, 52] require some restrictions on input sizes to recover the aggregate due to the discrete logarithm bottleneck.

	Efficient Aggregate	Dynamic Participation	TD vs DS	Assumptions
[81]	✗	✗	TD	DDH
[56]	✓	✗	TD	DCR
[15]	✗/✓	✗	TD	DDH/DCR
[11]	✓	✗	TD	LWE/R-LWE
[85]	✓	✗	TD	R-LWE
[88]	✓	✗	TD	AES
[83]	✓	✗	TD	RLWE
[61]	✗	✓	TD	DCR
[48]	✓	✗	TD	LWR
[22]	✗	✗	DS	LWR
[17, 13]	✓	✓	DS*	DDH
[52]	✓	✓	DS	DDH
[67]	✓	✓	DS	DDH
[62]	✓	✓	DS	DDH
Our Work	✓	✓	NA	HSM, LWR, (R)LWE

**Committee-Based MPC.** Committee-based MPC is widely used for handling scenarios involving a large number of parties. However, it faces a security vulnerability known as adaptive security, where an adversary can corrupt parties after committee selection. The YOSO model, introduced by Gentry *et al.* [36] proposes a model that offers adaptive security. In YOSO, committees speak once and are dynamically changed in each communication round, preventing adversaries from corrupting parties effectively. The key feature of YOSO is that the identity of the next committee is known only to its members, who communicate only once and then become obsolete to adversaries. YOSO runs generic secure computation calculations and aggregation can be one of them. However, its efficiency is prohibitive for secure aggregation. In particular, the communication complexity of YOSO in the computational setting scales quadratic with the number of parties  $n$  (or linear in  $n$  if the cost is amortized over  $n$  gates for large circuits). Additionally, to select the committees, an expensive role assignment protocol is applied. Like LERNA, in YOSO also specific sizes for the committee need to be fulfilled to run a protocol execution. Last but not least, our protocol does not rely on any secure role assignment protocol to choose the committees since even if all committee members are corrupted, privacy is still preserved. Fluid MPC [38, 16] also considers committee-based general secure computation. However, like YOSO, it is not practical. Unlike YOSO, it lacks support for adaptive security.

Moreover, SCALES [2] considers ephemeral servers a la YOSO responsible for generic MPC computations where the clients only provide their inputs. This approach is of theoretical interest as it is based on heavy machinery such as garbling and oblivious transfer if they were to be considered for the task of secure aggregation. Moreover, SCALES needs extra effort to hide the identities of the servers which we do not require.

## B Technical Overview

In this work, we focus on building a primitive, One-shot Private Aggregation (OPA), that enables privacy-preserving aggregation of multiple inputs, across several aggregation iterations whereby a client only speaks once on his will, per iteration.

**Seed-Homomorphic PRG (SHPRG).** A PRG  $G : \mathcal{K} \rightarrow \mathcal{Y}$  is said to be seed-homomorphic if  $G(s_1 \oplus s_2) = G(s_1) \otimes G(s_2)$  where  $(\oplus, \mathcal{K})$  and  $(\otimes, \mathcal{Y})$  are groups. We show how to build SHPRG from LWR and LWE Assumption. Note that the former was known but the latter was not. For ease of exposition,

we will focus on the LWR construction. For a randomly chosen  $\mathbf{A} \leftarrow_s \mathbb{Z}_q^{L \times n}$ ,  $\text{PRG}_{\text{LWR}, \mathbf{A}}(\mathbf{s}) = \lfloor \mathbf{A}\mathbf{s} \rfloor_p$  where  $p < q$ . Unfortunately, this construction is only almost seed-homomorphic in that there is an induced error. This is formally defined in Construction 4.

**Secret Sharing over Finite Fields.** In standard Shamir secret sharing [80], one picks a secret  $s$  and generate a polynomial  $f(X) = \sum_{i=0}^{r-1} a_i \cdot X^i$  where  $a_0 = s$  and  $a_1, \dots, a_{r-1}$  are randomly chosen from the field. Assuming there are  $m$  parties, the share for party  $i \in [m]$  is  $f(i)$ , and any subset of at least  $r$  parties can reconstruct  $s$  and any subset of  $r - 1$  shares are independently random. We can lower the corruption threshold from  $r - 1$  to  $t$  to obtain additional properties from the scheme. In packed secret sharing [49], one can hide multiple secrets using a single polynomial.

**One-shot Private Aggregation (based on SHPRG).** With the needed background in place, we are ready to build the new primitive called One-shot Private Aggregation. Critically, we want to empower a client to speak once, per iteration, and help the server successfully aggregate long vectors (of length  $L$ , say). To this end, we have the client communicate at the same time with the server and to a set of committee members, via the server. We will assume that there is a PKI or a mechanism to retrieve the public key of these committee members so that the communication to the committee can be encrypted. These are  $m$  ephemeral chosen clients who are tasked with helping the server aggregate, for that iteration. Flamingo [67, Figure 1, Lines 2-7] presents an approach on how to sample these committee members using a randomness beacon such as [47]. Observe that the one-shot communication can be leveraged to successfully avoid the complex setup procedures that were reminiscent of prior work. Instead, the client  $i$  samples a seed from the seed space of the PRG, i.e.,  $\text{sd}_i \leftarrow_s \text{PRF}.\mathcal{K}$ . Then, the PRG is expanded under this seed  $\text{sd}_i$  and effectively serves as a mask for input  $\mathbf{x}_{i,\ell}$ . Here  $\ell$  is the current iteration number. In other words, it computes  $\mathbf{ct}_{i,\ell} = \mathbf{x}_{i,\ell} + \text{PRG}.\text{Expand}(\text{sd}_i)$ . Intuitively, the PRF security implies that the evaluation is pseudorandom and is an effective mask for the input. Then, the client secret shares using standard Shamir’s Secret Sharing,  $\text{sd}_i$  to get shares  $\left\{ \text{sd}_i^{(j)} \right\}_{j \in [m]}$ .

Share  $\text{sd}_i^{(j)}$  is sent to committee member  $j$ , via the server through an appropriate encryption algorithm. Each committee member simply adds up the shares, using the linear homomorphism property of Shamir’s secret sharing. After receiving from the clients for that round, the committee member  $j$  sends the added shares, which corresponds to  $\sum_{i=1}^n \text{sd}_i^{(j)}$ . The server can successfully reconstruct from the information sent by the committee to get  $\sum_{i=1}^n \text{sd}_i$ . Note that adding up the ciphertext from the clients results in  $\sum_{i=1}^n \mathbf{ct}_{i,\ell} = \sum_{i=1}^n \mathbf{x}_{i,\ell} + \sum_{i=1}^n \text{PRG}.\text{Expand}(\text{sd}_i)$ . Key Homomorphism of the PRF implies that  $\sum_{i=1}^n \text{PRG}.\text{Expand}(\text{sd}_i) = \text{PRG}.\text{Expand}(\sum_{i=1}^n \text{sd}_i)$ . Note that the server, with the reconstructed information, can compute  $\text{PRG}.\text{Expand}(\sum_{i=1}^n \text{sd}_i)$  and subtract it from  $\sum_{i=1}^n \mathbf{ct}_{i,\ell}$  to recover the intended sum. This is the core idea behind our construction. However, there are a few caveats. First, the server recovers the sum of the keys  $\sum_{i=1}^n \text{sd}_i$ , which constitutes a leakage on the seed of the PRG. In other words, we require a leakage-resilient, seed homomorphic PRG. Second, LWR and the LWE based construction of seed homomorphic PRG are only almost seed homomorphic. Thus, care needs to be taken to encode and decode to ensure that the correct sum is recovered. We show that the LWR and LWE based schemes are key-homomorphic and leakage-resilient while describing the necessary encoding and decoding algorithms for the input. A key benefit of this construction is that the key shared with the committee is usually independent of the vector length. Our proof of security is in the simulation-based paradigm against both a semi-honest and an active server.

**One-shot Private Aggregation against Malicious Clients.** Prior work to ensure the detection of client misbehavior such as ACORN [12] requires the usage of verifiable secret sharing. Unfortunately, this is an expensive process as the committee member is required to do at  $r$  exponentiations, per client. Instead, we take an alternative route by designing a publicly verifiable secret-sharing scheme. This employs SCRAPE [27] and combines it with a simple  $\Sigma$ -protocol. As a result, the server can verify that given the commitment to shares, the shares are indeed a valid Shamir Secret Sharing. Meanwhile, the committee member simply needs to verify if the commitment to the share matches its encrypted share with the committee member raising an alarm should the check fail. This is a much leaner verification process for the committee members.

**CL Framework.** We use the generalized version of the framework, as presented by Bouvier *et al.* [20]. Broadly, there exists a group  $\hat{\mathbb{G}}$  whose order is  $M \cdot \hat{s}$  where  $\text{gcd}(M, \hat{s}) = 1$  and  $\hat{s}$  is unknown while  $M$  is a parameter of the scheme. Then,  $\hat{\mathbb{G}}$  admits a cyclic group  $\mathbb{F}$ , generated by  $f$  whose order is  $M$ . Consider the cyclic subgroup  $\mathbb{H}$  which is generated by  $h = x^M$ , for a random  $x \in \hat{\mathbb{G}}$ . Then, one can consider the cyclic subgroup  $\mathbb{G}$  generated by  $g = f \cdot h$  with  $\mathbb{G}$  factoring as  $\mathbb{F} \cdot \mathbb{H}$ . The order of



$\mathbb{G}$  is also unknown. The  $\text{HSM}_M$  assumption states that an adversary cannot distinguish between an element in  $\mathbb{G}$  and  $\mathbb{H}$ , while a discrete logarithm is easy in  $\mathbb{F}$ . Meanwhile,  $\bar{s}$ , an upper bound for  $\hat{s}$  is provided as input. Note that for  $M = N$  where  $N$  is an RSA modulus, the  $\text{HSM}_M$  assumption reduces to the DCR assumption. Therefore, the  $\text{HSM}_M$  assumption can be viewed as a generalization of the DCR assumption.

**Secret Sharing over Integers.** Braun *et al.* [21] was the first to identify how to suitably modify Shamir’s secret sharing protocol to ensure that the operations can work over a group of unknown order, such as the ones we use on the CL framework. This stems from two reasons. The first is leakage in that a share  $f(i)$  corresponding to some sharing polynomial  $f$  always leaks information about the secret  $s$ ,  $\text{mod } i$  when the operation is over the set of integers. Meanwhile, the standard approach to reconstruct the polynomial requires the computation of the Lagrange coefficients which involves dividing by an integer, which again needs to be “reinterpreted” to work over the set of integers. The solution to remedy both these problems is to multiply with an offset  $\Delta = m!$  where  $m$  is the total number of shares.

**(Almost) Key Homomorphic Pseudorandom Functions.** The concept of key homomorphic PRFs (KH-PRFs) was introduced by [70], who demonstrated that  $H(x)^k$  is a secure KH-PRF under the DDH assumption in the Random Oracle Model, where  $H$  is a hash function,  $k$  is the key, and  $x$  is the input. [19] later constructed an almost KH-PRF under the Learning with Rounding assumption [9], which was formally proven secure by [48]. Our work leverages almost KH-PRF constructions from both LWR and LWE assumptions in the Random Oracle Model. Additionally, we present a novel KH-PRF construction in the CL framework, yielding new constructions under the HSM assumption (including DCR-based constructions). We adapt the DDH-based construction to the CL framework and prove that  $F(k, x) = H(x)^k$ , where  $k \leftarrow_s \mathcal{K}$  and  $H : \{0, 1\}^* \rightarrow \mathbb{H}$  is modeled as a random oracle, is a secure KH-PRF under the HSM assumption. This adaptation requires careful consideration of appropriate groups, input spaces, and key spaces.

**Distributed Key Homomorphic Pseudorandom Functions (KHPRF).** [19] presented generic constructions of Distributed PRFs from any KH-PRF using secret sharing techniques. However, the CL framework’s use of groups with unknown order necessitates working over integer spaces. While Linear Integer Secret Sharing [44] exists, it can be computationally expensive. Instead, we utilize Shamir Secret Sharing over Integer Space as described by [21], refining it with appropriate offsets to construct Distributed PRFs from our HSM-based construction. In other words, rather than reconstructing the secret, the solution reconstructs a deterministic function of the secret which is accounted for in the PRF evaluation. This induces complications in reducing the security of our distributed key homomorphic PRF to that of the HSM-based KH-PRF. Finally, we demonstrate that our construction maintains the key homomorphism property, allowing combinations of partial evaluations of secret key shares to match the evaluation of the sum of keys at the same input point.

**Building Distributed Key Homomorphic PRFs.** [19] introduced a generic construction of distributed KHPRFs from almost key homomorphic PRFs. They proposed  $F_{\text{LWR}}(\mathbf{k}, x) := \lfloor \langle H(x), \mathbf{k} \rangle \rfloor_p$  as an almost key homomorphic PRF, where  $q > p$  are primes,  $\mathbf{k} \leftarrow_s \mathbb{Z}_q^r$ , and  $H$  is a suitably defined hash function. Their reduction to build a distributed PRF utilizes standard Shamir’s Secret Sharing over fields, simplifying the process compared to integer secret sharing due to the prime nature of  $q$  and  $p$ . However, their proposed construction contained shortcomings affecting both correctness and security proofs. We will now provide a brief overview of these issues. An almost KH-PRF satisfies  $F(k_1 + k_2, x) = F(k_1, x) + F(k_2, x) - e$  for some error  $e$ . For the LWR construction,  $e \in 0, 1$ . However, this implies  $F(T \cdot k, x) = T \cdot F(k, x) - e_T$  where  $e_T \in 0, \dots, T - 1$  for any integer  $T$ , leading to error growth and affecting Lagrange interpolation. The authors proposed multiplying by an offset  $\Delta = m!$  to bound the error and use rounding to mitigate its impact by ensuring that the error terms are “eaten” up. However, their security reduction faces challenges when simulating partial evaluations for unknown  $i^*$ . Lagrange interpolation with the “clearing out the denominator” technique causes further error growth, necessitating additional rounding. Consequently, the challenger can only provide  $\left\lfloor \frac{\Delta F(k^{(i)}, x)}{u} \right\rfloor$ . Thus, their definition of partial evaluation function needs to be updated to be consistent with what is simulatable. We identified further issues in their rounding choices. Specifically, partial evaluations should be rounded down to  $u$  where  $\lfloor p/u \rfloor > (\Delta + 1) \cdot r \cdot \Delta$  ( $r$  being the reconstruction threshold). Moreover, their framework only addressed single-key PRFs, not vector-key cases like LWR. We address these issues in our construction of a distributed, almost key

homomorphic PRF based on LWR. We formally prove that  $F_{\text{LWR}}(\mathbf{k}, x) = \lfloor \Delta \lfloor \Delta \lfloor \langle \mathbf{H}(x), \mathbf{k} \rangle_p \rfloor_u \rfloor_v$  for appropriate choices of  $u$  and  $v$  is a secure, distributed, almost key homomorphic, PRF.

**One-shot Private Aggregation without Leakage Simulation.** The construction of OPA' is similar to the earlier ones based on seed-homomorphic PRG. There exist the following differences:

- A client  $i$  has to sample  $L$  different keys. Each key  $in$  is used to evaluate the PRF at a point  $\ell$  and is used to mask the input  $x_i^{(in)}$ .
- It then secret shares each of the  $L$  keys by running the DPRF.Share, the algorithm to generate the shares of the DPRF key.
- Each share for committee member  $j$  is evaluated at  $\ell$ . This evaluation is sent to the committee member.
- Finally, the server runs DPRF.Combine to combine the information from the committee members to get the PRF evaluation at  $\ell$  under the sum of the keys, for each index  $in$ . Combine is the algorithm that helps reconstruct the evaluation from partial evaluations.

**Stronger Security Definitions and Construction.** We also present a stronger security guarantee, which was not provided by the committees of Lerna and Flamingo, whereby the committee members can all collude and observe all encrypted ciphertexts and all auxiliary information, and cannot mount an IND-CPA-style attack. Unfortunately, our current construction where the inputs are solely blinded by the PRF evaluation, which is also provided to the committee members in shares, can be unblinded by the committee leaking information about the inputs. We modify the syntax where each label/iteration begins with the server, which has its secret key, advertising a “public key” for that iteration (the keys for all iterations can also be published beforehand). The auxiliary information sent by a client to the committee is a function of this public key while the actual ciphertext is independent of this public key. Intuitively, this guarantees that the committee member’s information is “blinded” by the secret key of the server and cannot be used to unmask the information sent by the client. We now describe our updated construction. For each label  $\ell$ , the server publishes  $F(k_0, \ell)$  where  $k_0$  is its public key. Then, the auxiliary information sent by the client is of the form  $F(k_i^{(j)}, \ell)$  where  $k_i^{(j)}$  is the  $j$ -th share of the  $i$ -th clients key. Client  $i$  masks its input by doing  $f^{x_i} \cdot F(k_0, \ell)^{k_i}$ . Committee member  $j$  combines the results to then send  $F(\sum_{i=1}^n k_i^{(j)}, \ell)$  to the server. The server uses Lagrange interpolation and its own key  $k_0$  to compute:  $F(k_0 \sum_{i=1}^n k_i, \ell)$ . Meanwhile, the server, upon multiplying the ciphertexts gets  $X_\ell = f^{\sum_{i=1}^n x_i \cdot \ell} \cdot F(k_0 \sum_{i=1}^n k_i, \ell)$ . The recovery is straightforward after this point.

## B.1 Our Construction of One-shot Private Aggregation Scheme

The key idea behind our construction is that a seed-homomorphic PRG allows us to have a single seed, much shorter than  $L$ , to mask all the  $L$  inputs. Then, one can simply secret-share the seed, which reduces computation and communication. Our seed-homomorphic PRGs are in the standard model. However, it is important to note that for the intended application of Federated Learning one might have to rely on the random oracle model to thwart attacks such as those pointed out by Pasquini *et al.* [72]. Though the underpinning idea of OPA is the combination of seed-homomorphic PRG and an appropriate secret-sharing scheme, there are technical issues with presenting a generic construction. We begin by presenting the construction based on the Learning with Rounding Assumption.

### B.1.1 Construction of OPA<sub>LWR</sub>

We now present OPA<sub>LWR</sub> and prove its correctness and security. We rely on the seed-homomorphic PRG (Construction 4) and combine it with Shamir’s Secret Sharing Scheme over  $\mathbb{F}_q$  (Construction 6).

**Construction 1.** We present our construction in Figure 6.

**Correctness.** First, recall that Construction 4 is only almost seed homomorphic. In other words,

$$\text{PRG}(sd_1 + sd_2) = \text{PRG}(sd_1) + \text{PRG}(sd_2) + e$$

where  $e \in \{0, 1\}$ .

For ease of presentation, our correctness proof is for  $L = 1$ , but it extends to any arbitrary  $L$ . Therefore, while the correctness of the Shamir’s Secret Sharing scheme guarantees that  $sd_\ell$ , computed

### Protocol Construction of OPA<sub>LWR</sub>

#### One-Time System Parameters Generation

Run  $pp \leftarrow \text{SS.Setup}(1^\kappa, 1^t, 1^r, 1^m)$   
 Run  $pp_{\text{PRG}} \leftarrow \text{PRG.Gen}(1^{\kappa c}, 1^{\kappa s})$   
 Set  $pp = (pp_{\text{PRG}}, t, r, m)$   
**return** Committee of size  $m$  and  $pp$ .

#### Data Encryption Phase by Client $i$ in iteration $\ell$

Sample  $sd_{i,\ell} \leftarrow \text{PRG.C}, \text{dig}_{i,\ell} \leftarrow \{0, 1\}^{\log q}$   
 Compute  $(x_1, \dots, x_L) \leftarrow \text{Encode}(\mathbf{x}_i)$   
 Compute  $\text{mask}_{i,\ell} = \text{PRG.Expand}(sd_{i,\ell}), \text{mask}'_{i,\ell} := \text{H}(\text{dig}_{i,\ell})$   
 Compute  $\text{ct}_{i,\ell} = (x_1, \dots, x_L) + \text{mask}_{i,\ell} + \text{mask}'_{i,\ell}$   
 $(sd_i^{(j)})_{j \in [m]} \leftarrow \text{SS.Share}(sd_{i,\ell}, t, r, m)$   
 $(\text{dig}_{i,\ell}^{(j)})_{j \in [m]} \leftarrow \text{SS.Share}(\text{dig}_{i,\ell}, t, r, m)$   
**for**  $j = 1, \dots, m$  **do**  
    $\text{aux}_{i,\ell}^{(j)} \leftarrow sd_i^{(j)}, \text{dig}_{i,\ell}^{(j)}$   
 Send  $\text{ct}_{i,\ell}$  to the Server  
 Send  $c_{i,\ell}^{(j)} = \text{Enc}_{pk_j}(\ell, i, \text{aux}_{i,\ell}^{(j)})$  to committee member  $j$  for each  $j \in [m]$ , via server.

#### Set Intersection Phase by Server in iteration $\ell$

**For**  $j \in [m]$ , let  $\mathcal{C}^{(j)} := \{i : c_{i,\ell}^{(j)} \text{ was received by server}\}$   
 Let  $\mathcal{C}^{(0)} := \{i : \text{ct}_{i,\ell} \text{ was received by the server}\}$   
 Compute  $\mathcal{C} := \bigcap_{j \in \mathcal{S} \cup \{0\}} \mathcal{C}^{(j)}$  // This is bit-wise AND operation of the  $r+1$  bit strings.  
**assert**  $|\mathcal{C}| \geq (1-\delta)n$   
 Send  $\mathcal{C}, \{c_{i,\ell}^{(j)}\}_{i \in \mathcal{C}}$  for every  $j$  in  $[m]$

#### Data Combination Phase by Committee Member $j$ in iteration $\ell$

Recover  $(i, \ell, \{\text{aux}_{i,\ell}^{(j)} = (sd_i^{(j)}, \text{dig}_{i,\ell}^{(j)})\})$  from  $\{c_{i,\ell}^{(j)}\}_{i \in \mathcal{C}}$   
 Verify  $i \in \mathcal{C}$  and  $\ell$  is the current iteration, else abort.  
 Compute  $\text{AUX}^{(j)} \leftarrow \sum_{i \in \mathcal{C}} sd_{i,\ell}^{(j)}$   
 Send  $\text{AUX}^{(j)}, \{\text{dig}_{i,\ell}^{(j)}\}_{i \in \mathcal{C}}$  to server

#### Data Aggregation Phase by Server in iteration $\ell$

Let  $\{\text{AUX}_\ell^{(j)}\}_{j \in \mathcal{S}}, \{\text{ct}_{i,\ell}\}_{i \in \mathcal{C}}$  be the inputs received by the server with  $|\mathcal{S}| \geq r$ .  
 Run  $sd_\ell \leftarrow \text{SS.Reconstruct}(\{\text{AUX}_\ell^{(j)}\}_{j \in \mathcal{S}})$   
 $\text{AUX}_\ell = \text{PRG.Expand}(sd_\ell)$   
**for**  $i \in \mathcal{C}$  **do**  
    $\text{dig}_{i,\ell} \leftarrow \text{SS.Reconstruct}(\{\text{dig}_{i,\ell}^{(j)}\}_{j \in \mathcal{S}})$   
 Compute  $\text{CT}_\ell \leftarrow \sum_{i \in \mathcal{C}} \text{ct}_{i,\ell}$   
 Compute  $(X_1, \dots, X_L) = \text{CT}_\ell - \text{AUX}_\ell - \sum_{i \in \mathcal{C}} \text{H}(\text{dig}_{i,\ell})$   
 Compute  $\mathbf{X}_\ell \leftarrow \text{Decode}(pp, (X_1, \dots, X_L))$   
**return**  $\mathbf{X}_\ell$

Figure 6: Our Construction of OPA built from LWR-based Seed Homomorphic PRG (PRG) and Shamir's secret sharing scheme SS. Here,  $\text{Encode}(\mathbf{x}_{i,\ell}) := n \cdot \mathbf{x}_{i,\ell} + 1$  and  $\text{Decode}(X_i) := \lceil X_i/n \rceil - 1$ . The [lines](#) are for security against an active server. The use of the second mask  $\text{mask}'_{i,\ell}$ , as the output of a H is for simulation proof, for an active server. We will model H as a programmable random oracle.

in S-Combine, is indeed  $\sum_{i=1}^n sd_i \bmod q$ , there is an error growth in  $\text{AUX}_\ell$ . Specifically, we get that:

$$\text{AUX}_\ell := \text{PRG.Expand} \left( sd_\ell = \sum_{i=1}^n sd_{i,\ell} \right) = \sum_{i=1}^n \text{PRG.Expand}(sd_{i,\ell}) + e'$$

where  $e' \in \{0, \dots, n-1\}$ . We know that  $\text{Encode}(x_{i,\ell}) := n \cdot x_{i,\ell} + 1$ .

$$\begin{aligned}
\sum_{i=1}^n \text{ct}_{i,\ell} - \text{AUX}_\ell &= \left( \sum_{i=1}^n (x_{i,\ell} \cdot n + 1) + \text{Expand}(\text{sd}_{i,\ell}) \right) - \text{Expand} \left( \sum_{i=1}^n \text{sd}_{i,\ell} \right) \bmod p \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n + \sum_{i=1}^n \text{Expand}(\text{sd}_{i,\ell}) - \text{Expand} \left( \sum_{i=1}^n \text{sd}_{i,\ell} \right) \bmod p \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n + \text{Expand} \left( \sum_{i=1}^n \text{sd}_{i,\ell} \right) - \text{Expand} \left( \sum_{i=1}^n \text{sd}_{i,\ell} \right) - e' \bmod p \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n - e' \bmod p \\
\bar{X}_\ell &= n \cdot \sum_{i=1}^n x_{i,\ell} + n - e'
\end{aligned}$$

To make the last jump in the proof, we require:

$$0 \leq n \cdot \sum_{i=1}^n x_{i,\ell} + n - e' < p$$

First,  $e' \leq n-1$ . This guarantees that:  $0 \leq n \cdot \sum_{i=1}^n x_{i,\ell} + n - e'$ . Now, if  $\sum_{i=1}^n x_{i,\ell} < (p-n)/n$  then we also get:

$$n \cdot \sum_{i=1}^n x_{i,\ell} + n - e' < p$$

Now, we show the correctness of Decode algorithm to recover  $\sum_{i=1}^n x_{i,\ell}$  from  $\bar{X}_\ell$ .

- $\bar{X}_\ell/n = \sum x_{i,\ell} + (n - e')/n$
- $0 \leq e' \leq n-1 \Rightarrow 1/n \leq (n - e')/n \leq 1$
- Therefore,  $\lceil \bar{X}_\ell/n \rceil = \sum x_{i,\ell} + 1$

**Theorem 1.** Let  $\kappa_s$  and  $\kappa_c$  be the statistical and computational security parameters. Let  $L$  be the input dimension and  $n$  be the number of clients, that are  $\text{poly}(\kappa_c)$ . Let  $\delta$  be the dropout threshold and  $\eta$  be the corruption threshold such that  $\delta + \eta < 1$ . Then, there exists an efficient simulator  $\text{Sim}$  such that for all  $K \subset [n]$  such that  $|K| \leq \eta n$ , inputs  $X = \{\mathbf{x}_{i,\ell}\}_{i \in n \setminus K}$ , and for all adversaries  $\mathcal{A}$  against Construction 1, that controls the server and the set of corrupted clients  $K$ , which behave semi-honestly (resp. maliciously), the output of  $\text{Sim}$  is computationally indistinguishable from the joint view of the server and the corrupted clients.  $\text{Sim}$  is allowed to query  $F_{D,\delta}^\ell(X)$  (defined in Equation ??) once, per iteration.

Proof is deferred to Section H.

**Remark 1** (Malicious Server and Inconsistent Updates). Recent work by Pasquini *et al.* [72] describes an attack where the server can send different models to different client updates with the goal that the model sent to a particular client can negate the training done by other clients on different models. In our case, this attack can be easily remedied with no overhead. Rather than evaluating the PRG with just the public matrix  $\mathbf{A}$ , one can first compute a hash  $\mathbb{H}(\text{model})$ , and multiply it with  $\mathbf{A}$ . This would ensure that the matrix used by the client is tied to the model update sent. If different clients use different  $\mathbf{A}$ , the seed homomorphism fails. This would make it difficult, for a malicious server to send different models to different clients, to ensure that a particular client's contributions are not aggregated and therefore can be recovered. We can also simply switch to the key-homomorphic PRF constructions described in Section G.

In Section I.1, we model security when the entire committee is corrupted and the committee, through some attack, recovers the ciphertext of an honest client. The goal is to ensure that the input of the honest client is still preserved. In this definition, the server is honest.

### B.1.2 Construction of $\text{OPA}_{\text{LWE}}$

**Construction 2.** As alluded to before,  $\text{OPA}_{\text{LWE}}$  is largely similar to  $\text{OPA}_{\text{LWR}}$  with the following differences:

- While the seed of  $\text{PRG}_{\text{LWE}}$  is  $(\text{sd}, \mathbf{e})$ , we will only secret share  $\text{sd}$ . We will argue below that the correctness still holds, for suitable definition of  $\chi$ .
- The plaintext space for  $\text{OPA}_{\text{LWE}}$ , like the one for  $\text{OPA}_{\text{LWR}}$ , is  $\mathbb{Z}_p$ . Meanwhile the seed space for both  $\text{OPA}_{\text{LWE}}$  and  $\text{OPA}_{\text{LWR}}$  will be  $\mathbb{Z}_q$ . Let  $\Delta := \lfloor q/p \rfloor$ .
- We will use Shamir's Secret Sharing over  $q$ , as before, which is the seed space.
- There is a change in S-Combine. To compute  $\text{AUX}_\ell$ , the server uses the reconstructed seed  $\text{sd}_\ell$ , and additionally sets the error component of the PRG seed to be 0.
- $\text{Encode}(\mathbf{x}_{i,\ell}) := \Delta \cdot \mathbf{x}_{i,\ell}$
- $\text{Decode}(X_i) := \lceil X_i / \Delta \rceil - 1$

Due to the similarities, we do not present the construction in its entirety. Meanwhile, we present the proof of correctness and proof of security.

**Correctness.** First, observe that Construction 4 is only almost seed-homomorphic, i.e.

$$\text{PRG}((\text{sd}_1 + \text{sd}_2, \mathbf{e}), \ell) = \text{PRG}((\text{sd}_1, \mathbf{e}_1), \ell) + \text{PRG}((\text{sd}_2, \mathbf{e}_2), \ell) + \mathbf{e}'$$

for some error  $\mathbf{e}'$ . Indeed, assuming the correctness of Shamir's Secret Sharing, we get that the server computes:

$$\text{AUX}_\ell := \text{PRG.Expand} \left( \left( \sum_{i=1}^n \text{sd}_i, 0 \right), \ell \right) := \mathbf{A} \cdot \sum_{i=1}^n \text{sd}_i$$

Meanwhile,

$$\begin{aligned} \sum_{i=1}^n \text{ct}_{i,\ell} &= \sum_{i=1}^n (\mathbf{A} \text{sd}_i + \mathbf{e}_i + \Delta \cdot x_{i,\ell}) \\ &= \mathbf{A} \sum_{i=1}^n \text{sd}_i + \sum_{i=1}^n \mathbf{e}_i + \sum_{i=1}^n \Delta \cdot x_{i,\ell} \\ \text{Let } X_\ell &:= \sum_{i=1}^n \text{ct}_{i,\ell} - \text{AUX}_\ell = \sum_{i=1}^n \mathbf{e}_i + \sum_{i=1}^n \Delta \cdot x_{i,\ell} \\ \frac{X_\ell}{\Delta} &= \frac{\sum_{i=1}^n \mathbf{e}_i}{\Delta} + \sum_{i=1}^n x_{i,\ell} \end{aligned}$$

If  $\sum_{i=1}^n \mathbf{e}_i < \frac{\Delta}{2}$ , then  $\lceil \frac{X_\ell}{\Delta} \rceil = \sum_{i=1}^n x_{i,\ell} + 1$ . This shows the correctness of our algorithm.

**Theorem 2.** Let  $\kappa_s$  and  $\kappa_c$  be the statistical and computational security parameters. Let  $L$  be the input dimension and  $n$  be the number of clients, that are  $\text{poly}(\kappa_c)$ . Let  $\delta$  be the dropout threshold and  $\eta$  be the corruption threshold such that  $\delta + \eta < 1$ . Then, there exists an efficient simulator  $\text{Sim}$  such that for all  $K \subset [n]$  such that  $|K| \leq \eta n$ , inputs  $X = \{\mathbf{x}_{i,\ell}\}_{i \in [n] \setminus K}$ , and for all adversaries  $\mathcal{A}$  against Construction 2, that controls the server and the set of corrupted clients  $K$ , which behave semi-honestly (resp. maliciously), the output of  $\text{Sim}$  is computationally indistinguishable from the joint view of the server and the corrupted clients.  $\text{Sim}$  is allowed to query  $F_{D,\delta}(X)^\ell$  (defined in Equation ??) once, per iteration.

Proof deferred to Section H.

## B.2 Security Against Malicious Clients

In this section, we will focus on a misbehaving client. Observe that the client can send inconsistent shares to the committee. A standard approach would be for the client to rely on verifiable secret sharing which would empower each committee member to verify if the share received by the committee member is consistent with the commitments to the polynomial that was sent by the client. However, this requires each client to perform  $n \cdot r$  exponentiations which can be expensive. Instead, we take the approach of a modified publicly verifiable secret sharing scheme. This approach empowers the server, which has more computation power, to verify if the sharing is consistent.

Our public verifiability will rely on a modification of SCRAPE [27]. SCRAPE test is done to check if  $(\text{sd}_{i,\ell}^{(1)}, \dots, \text{sd}_{i,\ell}^{(m)})$  is a Shamir sharing over  $\mathbb{F}$  of degree  $d = r - 1$  (namely there exists a polynomial  $p$

of degree  $\leq d$  such that  $p(i) = s_i$  for  $i = 1, \dots, n$ ), one can sample  $w^{(1)}, \dots, w^{(m)}$  uniformly from the dual code to the Reed-Solomon code formed by the evaluations of polynomials of degree  $\leq d$ , and check if  $\sum_{i=1}^m w^{(i)} \cdot \text{sd}_{i,\ell}^{(i)} = 0$  in  $\mathbb{F}$ . If the test passes, then  $\text{sd}_{i,\ell}^{(1)}, \dots, \text{sd}_{i,\ell}^{(m)}$  are Shamir Shares, except with probability  $1/|\mathbb{F}|$ .

**Lemma 3** (SCRAPE Test [28]). *Let  $\mathbb{F}$  be a finite field and let  $d = r - 1, m$  be parameters of the Shamir's Secret Sharing scheme such that  $0 \leq d \leq m - 2$ , and inputs  $\text{sd}_{i,\ell}^{(1)}, \dots, \text{sd}_{i,\ell}^{(m)} \in \mathbb{F}$ . Define  $v_i := \prod_{j \in [m] \setminus \{i\}} (i - j)^{-1}$  and let  $m^*(X) := \sum_{i=0}^{m-d-2} m_i \cdot X^i \leftarrow_{\mathbb{F}[X]_{\leq m-d-2}}$  (i.e., a random polynomial over the field of degree at most  $m - d - 2$ ). Now, let  $\mathbf{w} := (v_1 \cdot m^*(1), \dots, v_n \cdot m^*(n))$  and  $\mathbf{s} := (\text{sd}_{i,\ell}^{(1)}, \dots, \text{sd}_{i,\ell}^{(m)})$ . Then,*

- If there exists  $p \in \mathbb{F}[X]_{\leq d}$  such that  $\text{sd}_{i,\ell}^{(i)} = p(i)$  for all  $i \in [n]$ , then  $\langle \mathbf{w}, \mathbf{s} \rangle = 0$ .
- Otherwise,  $\Pr[\langle \mathbf{w}, \mathbf{s} \rangle = 0] = 1/|\mathbb{F}|$ .

Typically, we compute the polynomial  $m^*(X)$  by using the Fiat-Shamir transform over public values. Then, the vector  $\mathbf{w}$  is a public vector. One simply has to hide the vector  $\mathbf{s}$ . As a result, we have proved that the shares do lie on the same polynomial. Note that in standard Shamir's Secret Sharing, we set  $p(0) = \text{sd}_{i,\ell}$ , i.e., the secret. Therefore, we will have to perform inner product over a vector of length  $m + 1$ .

This results in the following additional steps on the part of the client, the committee member, and the server. We detail only these additional steps.

**Construction 3** (Detecting Malicious Client Behavior). Let  $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{F}^{m-d-2}$  where  $d = r - 1$  be a hash function which is modeled as a random oracle. Let  $\mathcal{H}' : \{0, 1\}^* \rightarrow \mathbb{F}$  be the hash function used to generate the challenge. Let  $G$  be a group generated by  $g$  where Discrete Logarithm and DDH is hard, and is of prime order  $q$ , the same as the order of the field for Shamir Secret Sharing.

- Client  $i$  does the following:
  - Commit to  $\text{sd}_{i,\ell}^{(0)} = \text{sd}_{i,\ell}, \text{sd}_{i,\ell}^{(1)}, \dots, \text{sd}_{i,\ell}^{(m)}$  as  $C_s^{(j)} := g^{\text{sd}_{i,\ell}^{(j)}}$ .
  - Generate the coefficients to polynomial  $m - d - 2$  by using Fiat-Shamir transform. In other words, get  $m_0, \dots, m_{m-d-2} \leftarrow \mathcal{H}(C_s^{(0)}, \dots, C_s^{(m)})$
  - Compute  $v_0, \dots, v_m$  as  $v_i := \prod_{j \in \{0, \dots, m\} \setminus \{i\}} (i - j)^{-1}$ .
  - Compute  $\mathbf{w} := (v_0 \cdot m^*(0), \dots, v_m \cdot m^*(m))$
  - Compute  $\mathbf{t} := (t_0, \dots, t_m) \leftarrow_{\mathbb{F}}$
  - Compute  $C_t^{(0)}, \dots, C_t^{(m)}$  as commitments to  $t_0, \dots, t_m$  where  $C_t^{(j)} := g^{t_j}$
  - Set  $r := \langle \mathbf{t}, \mathbf{w} \rangle$
  - Compute  $c := \mathcal{H}'(C_s^{(0)}, \dots, C_s^{(m)}, C_t^{(0)}, \dots, C_t^{(m)}, \mathbf{w}, r)$
  - Compute  $z_0, \dots, z_m$  where  $z_i := t_i + c \cdot \text{sd}_{i,\ell}^{(i)}$
  - Set  $\pi_i := \left( \left\{ C_s^{(j)} \right\}, r, \mathbf{z} = (z_0, \dots, z_m), c \right)$
- Server does the following:
  - Upon receiving  $\pi_i$  from client  $i$ , the server parses  $\pi_i := \left( \left\{ C_s^{(j)} \right\}, r, \mathbf{z} = (z_0, \dots, z_m), c \right)$
  - It computes  $\mathbf{w}$  (similar to how the client does it). It then computes  $\langle \mathbf{w}, \mathbf{z} \rangle$ . It checks to see if this is equal to the value  $r$  sent by the client  $i$ .
  - For each  $j = 0, \dots, m$ , compute  $C_t^{(j)} = g^{z_j} \cdot \left( C_s^{(j)} \right)^{-c}$
  - Compute  $c' = \mathcal{H}'(C_s^{(0)}, \dots, C_s^{(m)}, C_t^{(0)}, \dots, C_t^{(m)}, \mathbf{w}, r)$
  - Accept input from client  $i$  if  $c == c'$ .
  - The server sends  $C_s^{(j)}$  to committee member  $j$ , along with the encrypted shares for committee member  $j$ .
- Committee member  $j$  does the following:



- Decrypt and recover the share  $sd_{i,\ell}^{(j)}$ . Verify that this matches the commitment forwarded by the server.

It is easy to verify that an honest prover will satisfy the proof. This follows from the SCRAPE test (Lemma 3) and also the completeness of the generalized  $\Sigma$ -protocol. Meanwhile, if there exists two accepting transcripts  $\mathbf{z}_1, \mathbf{z}_2$  corresponding to  $c_1, c_2$ , then one can extract a witness for  $sd_{i,\ell}^{(0)}, \dots, sd_{i,\ell}^{(m)}$ . This guarantees soundness. The zero-knowledge property follows by simply sampling a random  $\mathbf{z}$  and then setting the choice of the  $C_t^{(j)}$  by following the verification steps. Further, it can set  $r = \langle \mathbf{z}, \mathbf{w} \rangle$ .

### B.3 One-shot Private Aggregation without Leakage Simulation

It is easy to observe that the server, in the previous constructions of OPA, recovers the sum of the keys. This constitutes leakage of the honest clients' keys. Consequently, our proofs relied on the pseudorandomness, even contingent on this leakage. Furthermore, these constructions allowed the generation of a new key, in every iteration.

We will now present a construction OPA' such that: (a) a client's keys can be reused across multiple iterations, and, (b) the server does not get the sum of the keys but rather a function of pseudorandom values, which can argued as itself being pseudorandom. The core technique we employ in this work is a distributed, key-homomorphic PRF. We formally present constructions from the CL framework in Section E. Specifically, we defer OPA<sub>CL</sub> to the appendix in Section E.2. Similarly, we present LWR based construction in Section F. Meanwhile, we broadly describe the intuition behind our construction.

A distributed-key-homomorphic PRF has two specific algorithms: Eval which allows to evaluation of the PRF with key  $k_i$  at a point. While, P-Eval allows the PRF to be evaluated at a share of the key  $k_i^{(j)}$  to get a partial evaluation. Therefore, in our construction, the clients mask a vector of inputs  $\mathbf{x}_{i,\ell}$  by computing a pseudorandom evaluation of  $\text{DPRF.Eval}(k_{i,1}, \ell, \dots, \text{DPRF.Eval}(k_{i,L}, \ell))$ . Meanwhile, the auxiliary information send to the committee member will be  $\text{P-Eval}(k_{i,k}^{(j)}, \ell)$  for  $k = 1, \dots, L$  and  $j \in [m]$ . Then, the committee members combine by multiplying the auxiliary information. The server then reconstructs on its end. Note that the server only computes  $\text{DPRF.Eval}(\sum_{i=1}^n k_{i,k}, \ell)$  for  $k = 1, \dots, L$ . This is a PRF evaluation and therefore the leakage is pseudorandom and can be easily simulated, replacing it with random.

## C Cryptographic Preliminaries

For completeness, we discuss secret sharing in Section C.2. We discuss pseudorandom functions in Section C.4. We then introduce lattice-based cryptographic assumptions in Section C.3.

### C.1 Seed Homomorphic PRG

#### C.1.1 Syntax and Security

**Definition 1** (Seed Homomorphic PRG (SHPRG)). *Consider an efficiently computable function  $\text{PRG} : \mathcal{K} \rightarrow \mathcal{Y}$  where  $(\mathcal{K}, \oplus), (\mathcal{Y}, \otimes)$  are groups. Then  $(\text{PRG}, \oplus, \otimes)$  is said to be a secure seed homomorphic pseudorandom generator (SHPRG) if:*

- PRG is a secure pseudorandom generator (PRG), i.e., for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:

$$\Pr \left[ b = b' \mid \begin{array}{l} \text{pp}_{\text{PRG}} \leftarrow \text{PRG.Gen}, b \leftarrow \{0, 1\}, \text{sd} \leftarrow \mathcal{K} \\ Y_0 = \text{PRG.Expand}(\text{sd}), Y_1 \leftarrow cY \\ b' \leftarrow \mathcal{A}(Y_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\kappa)$$

- For every  $\text{sd}_1, \text{sd}_2 \in \mathcal{K}$ , we have that  $\text{PRG.Expand}(\text{sd}_1) \otimes \text{PRG.Expand}(\text{sd}_2) = \text{PRG.Expand}(\text{sd}_1 \oplus \text{sd}_2)$

In the construction below, we abuse notation and simply write  $\text{PRG}(\text{sd})$  as a shorthand for  $\text{PRG.Expand}(\text{sd})$

### C.1.2 Construction from LWR Assumption

**Construction 4** (SHPRG from LWR Assumption). Let  $\mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{n_1 \times L}$  be the output of PRG.Gen, then consider the following seed homomorphic PRG  $\text{PRG}_{\text{LWR}} : \{0, 1\}^{n_1} \rightarrow \{0, 1\}^L$  where  $L > n_1$  is defined as  $\text{PRG}_{\text{LWR}, \mathbf{A}}(\text{sd} = \mathbf{s}) = \lceil \mathbf{A}^\top \cdot \mathbf{s} \rceil_p$  where  $q > p$  with  $\lceil x \rceil_p = \lceil x \cdot p/q \rceil$  for  $x \in \mathbb{Z}_q$ .

This is almost seed homomorphic in that:  $\text{PRG}(\mathbf{s}_1 + \mathbf{s}_2) = \text{PRG}(\mathbf{s}_1) + \text{PRG}(\mathbf{s}_2) + e$  where  $e \in \{-1, 0, 1\}^L$

**Theorem 4** (Leakage Resilience of Construction 4). *Let  $\text{PRG}_{\text{LWR}}$  be the PRG defined in Construction 4. Then, it is leakage resilient in the following sense:*

$$\{\text{PRG}_{\text{LWR}}(\text{sd}) \bmod p, \text{sd} + \text{sd}' \bmod q : \text{sd}, \text{sd}' \leftarrow_{\$} \mathbb{Z}_q^{n_1}\} \approx_c \{\mathbf{y}, \text{sd} + \text{sd}' \bmod q : \mathbf{y} \leftarrow_{\$} \mathbb{Z}_p^L, \text{sd}, \text{sd}' \leftarrow_{\$} \mathbb{Z}_q^{n_1}\}$$

Proof deferred to Section H.

### C.1.3 Construction from LWE Assumption

**Construction 5** (SHPRG from LWE Assumption). Let  $\mathbf{A} \leftarrow_{\$} \mathbb{Z}_q^{L \times \lambda}$  be the output of PRG.Gen. Then, consider the following seed homomorphic PRG  $\text{PRG}_{\text{LWE}} : \mathbb{Z}_q^\lambda \times \chi^L \rightarrow \mathbb{Z}_q^L$  is defined as  $\text{PRG}_{\text{LWE}, \mathbf{A}}(\mathbf{s}, \mathbf{e}) = \mathbf{A}\mathbf{s} + \mathbf{e}$

The proof of security is a direct application of the LWE Assumption. However, we now prove the (almost) seed homomorphic property.

$$\begin{aligned} \text{PRG}_{\text{LWE}, \mathbf{A}}(\mathbf{s}_1, \mathbf{e}_1) &= \mathbf{A}\mathbf{s}_1 + \mathbf{e}_1 \\ \text{PRG}_{\text{LWE}, \mathbf{A}}(\mathbf{s}_2, \mathbf{e}_2) &= \mathbf{A}\mathbf{s}_2 + \mathbf{e}_2 \\ \text{PRG}_{\text{LWE}, \mathbf{A}}(\mathbf{s}_1, \mathbf{e}_1) + \text{PRG}_{\text{LWE}, \mathbf{A}}(\mathbf{s}_2, \mathbf{e}_2) &= \mathbf{A}(\mathbf{s}_1 + \mathbf{s}_2) + (\mathbf{e}_1 + \mathbf{e}_2) = \mathbf{A}(\mathbf{s}_1 + \mathbf{s}_2) + \mathbf{e} \end{aligned}$$

where  $\mathbf{e} \leq \mathbf{e}_1 + \mathbf{e}_2 + 1$ .

Looking ahead, when we use this PRG to mask the inputs, we will do the following:  $\text{PRG}_{\text{LWE}, \mathbf{A}}(s_i, e_i) = \mathbf{A}\mathbf{s}_i + \mathbf{e}_i + \lceil q/p \rceil \cdot \mathbf{x}_i$  where  $\mathbf{x}_i \in \mathbb{Z}_p^m$ . Upon adding  $n$  such ciphertexts (mod  $q$ ), we get:

$$\mathbf{A} \sum_{i=1}^n \mathbf{s}_i + \mathbf{e} + \lceil q/p \rceil \cdot \sum_{i=1}^n \mathbf{x}_i$$

where  $\mathbf{e} \leq 1 + \sum_{i=1}^n \mathbf{e}_i$ . Therefore, to eventually recover  $\sum_{i=1}^n \mathbf{x}_i \bmod p$  from just the value of  $\sum_{i=1}^n \mathbf{s}_i$ , we will require  $\|\mathbf{e}\|_\infty < \frac{q}{2p}$ . Looking ahead, we will rely on the Hint-LWE Assumption 6 to show that it is leakage resilient when we use it to build our aggregation tool.

**Remark 2** (Construction based on Ring-LWE). One can also extend the above LWE construction to the Ring-LWE [66] setting.

Recall the R-LWE Assumption. Let  $N$  be a power of two, and let  $m > 0$  be an integer. Let  $R$  be a cyclotomic ring of degree  $N$ , and let  $R_q$  be its residue ring modulo  $q > 0$ . Then, the following holds:

$$\{(\mathbf{a}, \mathbf{a} \cdot k + \mathbf{e}) : \mathbf{a} \leftarrow_{\$} R_q^m, k \leftarrow_{\$} R_q, \mathbf{e} \leftarrow_{\$} \chi^m\} \approx_c \{(\mathbf{a}, \mathbf{u}) : \mathbf{a} \leftarrow_{\$} R_q^m, \mathbf{u} \leftarrow_{\$} R_q^m\}$$

This gives us the following construction:  $\text{PRG}_{\text{R-LWE}}((k, e)) : \mathbf{a}k + \mathbf{e}$

## C.2 Secret Sharing

A key component of threshold cryptography is the ability to compute distributed exponentiation by sharing a secret. More formally, the standard approach is to compute  $g^s$  for some  $g \in \mathbb{G}$  where  $\mathbb{G}$  is a finite group and  $s$  is a secret exponent that has been secret-shared among multiple parties. This problem is much simpler when you assume that the group order is a publicly known prime  $p$  which then requires you to share the secret over the field  $\mathbb{Z}_p$ . This was the observation of Shamir [80] whereby a secret  $s$  can be written as a linear combination of  $\sum_{i \in \mathcal{S}} \alpha_i s_i \bmod p$  where  $\mathcal{S}$  is a set of servers that is sufficiently large and holds shares of the secret  $s_i$  and  $\alpha_i$  is only a function of the indices in  $\mathcal{S}$ . It follows that if each server provides  $g_i = g^{s_i}$ , then one can compute  $g^s = g^{\sum_{i \in \mathcal{S}} \alpha_i s_i} = \prod_{i \in \mathcal{S}} g_i^{\alpha_i}$ . Formally, this is defined below.

**Construction 6** (Shamir’s Secret Sharing over  $\mathbb{F}_q$ ). Consider the following  $(t, r, m)$  Secret Sharing Scheme where  $m$  is the total number of parties,  $t$  is the corruption threshold,  $r$  is the threshold for reconstruction. Then, we have the following scheme:

- **Share**( $s, t, r, m$ ): Sample a random polynomial  $f(X) \in \mathbb{F}_q[X]$  of degree  $r - 1$  such that  $f(0) = s$ . Then, **return**  $\{s^{(j)} := f(j)\}_{j \in [m]}$
- **Coeff**( $\mathcal{S}$ ): On input of a set  $\mathcal{S} = \{i_1, \dots, i_r, \dots\} \subseteq [m]$  of at least  $r$  indices, compute  $\lambda_{i_j} = \prod_{\zeta \in [r] \setminus \{j\}} \frac{i_\zeta}{i_\zeta - i_j}$ . Then, **return**  $\{\lambda_{i_j}\}_{i_j \in \{i_1, \dots, i_r\}}$
- **Reconstruct**( $\{s^{(j)}\}_{j \in \mathcal{S}}$ ): If  $|\mathcal{S}| \geq r$ , then output  $\sum_{j \in \mathcal{S}} \lambda_j \cdot s^{(j)}$  where  $\{\lambda_j\} \leftarrow \text{Coeff}(\mathcal{S})$ .

The correctness of the scheme guarantees that the secret  $s$  is correctly reconstructed.

**Construction 7** (Packed Secret Sharing over  $\mathbb{F}_q$ ). Consider the following  $(t, r, m)$  Secret Sharing Scheme where  $m$  is the total number of parties,  $t$  is the corruption threshold,  $r$  is the threshold for reconstruction. Further, let  $\rho$  be the number of secrets being packed which are to be embedded at points  $\text{pos}_1, \dots, \text{pos}_\rho$  where  $\text{pos}_i = m + i$ . Here,  $t := r - \rho$ . Then, we have the following scheme:

<p><b>PShare</b>(<math>\mathbf{s} = (s_1, \dots, s_\rho), t, r, m</math>)</p> <p><math>(r_0, \dots, r_{r-\rho-1}) \leftarrow \mathbb{F}_q</math></p> <p><math>q(X) := \sum_{i=0}^{r-\rho-1} X^i \cdot r_i</math></p> <p><math>\text{pos}_i = m + i</math> <b>for</b> <math>i = 1, \dots, \rho</math></p> <p><b>for</b> <math>i \in [\rho]</math> <b>do</b></p> <p style="padding-left: 20px;"><math>L_i(X) := \prod_{j \in [\rho] \setminus i} \frac{X - \text{pos}_j}{\text{pos}_i - \text{pos}_j} \cdot \Delta</math></p> <p><math>f(X) := q(X) \prod_{i=1}^{\rho} (X - \text{pos}_i) + \sum_{i=1}^{\rho} s_i \cdot L_i(X)</math></p> <p><b>return</b> <math>\{s^{(i)}\}_{i \in [m]}</math></p>	<p><b>Reconstruct</b>(<math>\{s^{(i)}\}_{i \in \mathcal{S}}</math>)</p> <p><b>if</b> <math> \mathcal{S}  &lt; r</math> <b>return</b> <math>\perp</math></p> <p><b>Parse</b> <math>\mathcal{S} := \{i_1, \dots, i_t, \dots\}</math></p> <p><b>for</b> <math>k \in [\rho]</math></p> <p style="padding-left: 20px;"><b>for</b> <math>j \in [t]</math></p> <p style="padding-left: 40px;"><math>\Lambda_{i_j}(X) := \prod_{\zeta \in [t] \setminus j} \frac{i_\zeta - X}{i_\zeta - i_j}</math></p> <p style="padding-left: 40px;"><math>s'_k := \sum_{j \in [t]} \Lambda_{i_j}(m + k) \cdot s^{(j)}</math></p> <p><b>return</b> <math>\mathbf{s}' := (s'_1, \dots, s'_\rho)</math></p>
---	---

### Parameters.

- For reconstruction, we require  $r < m(1 - \delta_C)$  where  $\delta_C$  is the dropout rate within the committee.
- For security, we require that  $(r - \rho) > m \cdot \eta_C$  where  $\eta_C$  is the corruption rate.

Combining, we get  $m > \rho / (1 - \delta_C - \eta_C)$ . Recall that we need  $\delta_C + \eta_C < 1/3$ , for byzantine fault tolerance. In other words, setting  $m \geq 3\rho/2$  is sufficient.

**Remark 3** (Optimizations for Packed Secret Sharing). Observe that the polynomial  $L_i(X)$  is only dependent on points  $\text{pos}_i$ , which are the points where the secrets are embedded. This can be pre-processed, and indeed, can be a part of the setup algorithm which distributes it to all the clients. Furthermore, rather than naively reconstructing the Lagrange polynomial, one can also rely on FFT techniques to achieve speed up.

Unfortunately, the above protocols do not extend to settings where the order of the group is not prime, not publicly known, or even possibly unknown to everyone. In this setting, the work of Damgård and Thorbek presents a construction to build Linear Integer Secret Sharing (LISS) schemes. In this work, we rely on the simpler scheme that extends Shamir’s secret sharing into the integer setting from the work of Braun *et al.* [21]. We also extend the Packed Secret Sharing scheme to this integer setting in Construction 9.

**Definition 2** (Secret Sharing over  $\mathbb{Z}$ ). A  $(t, r, m)$  Linear Integer Secret Sharing Scheme LISS is a tuple of PPT algorithms  $\text{LISS} := (\text{Share}, \text{GetCoeff}, \text{Reconstruct})$ , with the following public parameters: the statistical security parameter  $\kappa_s$ , the number of parties  $m$ , the corruption threshold  $t$ , and reconstruction threshold  $r$  of secrets needed for reconstruction, the randomness bit length  $\ell_r$ , the bit length of the secret  $\ell_s$  and the offset by which the secret is multiplied, denoted by  $\Delta = m!$ , and the following syntax:

- $(s_1, \dots, s_m) \leftarrow \text{Share}(s, m, r, t)$ : On input of the secret  $s$ , the number of parties  $m$ , and the threshold  $t$ , the share algorithm outputs shares  $s_1, \dots, s_m$  such that party  $i$  receives  $s_i$ .
- $\{\lambda_i\}_{i \in \mathcal{S}} \leftarrow \text{GetCoeff}(\mathcal{S})$ : On input of a set  $\mathcal{S}$  of at least  $r$  indices, the GetCoeff algorithm outputs the set of coefficients for polynomial reconstruction.
- $s' \leftarrow \text{Reconstruct}(\{s_i\}_{i \in \mathcal{S}})$ : On input of a set of secrets of at least  $r$  shares, the reconstruction algorithm outputs the secret  $s'$ .

We further require the following security properties.

- **Correctness:** For any  $m, \kappa_s, t, r, \ell_s, \ell_r \in \mathbb{Z}$  with  $t < r \leq m$ , and any set  $\mathcal{S} \subseteq [m]$  with  $|\mathcal{S}| \geq r$ , for any  $s \in \mathbb{Z}$  such that  $s \in [0, 2^{\ell_s})$  the following holds:

$$\Pr \left[ s' = f(s) \mid \begin{array}{l} (s_1, \dots, s_m) \leftarrow \text{Share}(s, m, r, t) \\ s' \leftarrow \text{Reconstruct}(\{s_i\}_{i \in \mathcal{S}}) \end{array} \right]$$

where  $f$  is some publicly computable function, usually  $f(s) = \Delta^2 \cdot s$ .

- **Statistical Privacy [44]:** We say that a  $(t, r, m)$  linear integer secret sharing scheme LISS is statistically private if for any set of corrupted parties  $\mathcal{C} \subset [m]$  with  $|\mathcal{C}| \leq t$ , and any two secrets  $s, s' \in [0, 2^{\ell_s})$  and for independent random coins  $\rho, \rho'$  such that  $\{s_i\}_{i \in [m]} \leftarrow \text{Share}(s; \rho)$ ,  $\{s'_i\}_{i \in [m]} \leftarrow \text{Share}(s'; \rho')$  we have that the statistical distance between:  $\{s_i | i \in \mathcal{C}\}$  and  $\{s'_i | i \in \mathcal{C}\}$  is negligible in the statistical security parameter  $\kappa_s$ .

**Construction 8 (Shamir's Secret Sharing over  $\mathbb{Z}$ ).** Consider the following  $(t, r, m)$  Integer Secret Sharing scheme where  $m$  is the number of parties,  $t$  is the corruption threshold, and  $r$  is the threshold for reconstruction. Further, let  $\kappa_s$  be a statistical security parameter. Let  $\ell_s$  be the bit length of the secret and let  $\ell_r$  be the bit length of the randomness. Then, we have the following scheme:

Share( $s, t, r, m$ )	GetCoeff( $\mathcal{S}$ )	Reconstruct( $\{s^{(i)}\}_{i \in \mathcal{S}}$ )
$\Delta := m!, \bar{s} := s \cdot \Delta$	if $ \mathcal{S}  \geq r$	if $ \mathcal{S}  \geq r$
$(r_1, \dots, r_{r-1}) \leftarrow [0, 2^{\ell_r + \kappa_s})$	for $i \in \mathcal{S}$ do	$\{\Lambda_i\}_{i \in \mathcal{S}} \leftarrow \text{GetCoeff}(\mathcal{S})$
$f(X) := \bar{s} + \sum_{i=1}^{r-1} r_i \cdot X^i$	$\Lambda_i := \prod_{j \in \mathcal{S} \setminus \{i\}} \frac{x_j}{x_j - x_i} \cdot \Delta$	$s' := \sum_{i \in \mathcal{S}} \Lambda_i \cdot s^{(i)}$
return $\{s^{(i)} = f(i)\}_{i \in [m]}$	return $\{\Lambda_i\}_{i \in \mathcal{S}}$	return $s'$

We omit the proof of correctness as it is similar to the original Shamir's Secret Sharing scheme. However, we highlight the critical differences:

- Unlike Shamir's Secret Sharing over fields, the secret here is already multiplied by the offset  $\Delta$ . Therefore, any attempt to reconstruct can only yield  $s \cdot \Delta$
- However, note that the inverse of  $x_j - x_i$  which was defined over the field  $\mathbb{Z}_q$  might not exist or be efficiently computable in a field of unknown order. Instead, we multiply the Lagrange coefficients by  $\Delta$ . Consequently, the reconstruction yields  $\Delta \cdot \bar{s}$  which equals  $s \cdot \Delta^2$ .

**Theorem 5 ([21]).** Construction 8 is statistically private provided  $\ell_r \geq \ell_s + \lceil \log_2(h_{\max} \cdot (t-1)) \rceil + 1$  where  $h_{\max}$  is an upper bound on the coefficients of the sweeping polynomial.

We refer the readers to the proof in [21, §B.1]. The key idea behind the proof is first to show that there exists a "sweeping polynomial" such that at each of the points that the adversary has a share of, the polynomial evaluates to 0 while at the point where the secret exists, it contains the offset  $\Delta$ . Implicitly, one can add the sweeping polynomial to the original polynomial whereby the sweeping polynomial "sweeps" away the secret information that the adversary has gained knowledge of. Meanwhile, in the later section, we present the proof for the generic construction that uses Shamir's Packed Secret Sharing over the integer space. This again uses the idea of a sweeping polynomial.

**Construction 9 (Shamir's Packed Secret Sharing over  $\mathbb{Z}$ ).** Let  $m$  be the number of parties and  $\rho$  be the number of secrets that are packed in one sharing. Further, let  $t$  denote the threshold for reconstruction (implies that corruption threshold is  $t - \rho$ ). Then, consider the following  $(m, t, \rho)$  Integer Secret Sharing Scheme with system parameters  $\kappa_s$  as the statistical security parameter,  $\ell_s$  is the bit length of the a secret, and let  $\ell_r$  be the bit length of the randomness. Then, we have the following scheme:

$\begin{aligned} & \text{PackedShare}(\mathbf{s} = (s_1, \dots, s_\rho), \mathbf{t}, m) \\ & \Delta := m!, \tilde{\mathbf{s}} := \mathbf{s} \cdot \Delta \\ & (r_0, \dots, r_{t-\rho-1}) \leftarrow \mathfrak{s}[0, 2^{\ell_r + \kappa_s}] \\ & q(X) := \sum_{i=0}^{t-\rho-1} X^i \cdot r_i \\ & \text{pos}_i = m + i \text{ for } i = 1, \dots, \rho \\ & \text{for } i \in [\rho] \text{ do} \\ & \quad L_i(X) := \prod_{j \in [\rho] \setminus \{i\}} \frac{X - \text{pos}_j}{\text{pos}_i - \text{pos}_j} \cdot \Delta \\ & f(X) := q(X) \prod_{i=1}^{\rho} (X - \text{pos}_i) + \sum_{i=1}^{\rho} \tilde{s}_i \cdot L_i(X) \\ & \text{return } \{s^{(i)}\}_{i \in [m]} \end{aligned}$	$\begin{aligned} & \text{Reconstruct}(\{s^{(i)}\}_{i \in \mathcal{S}}) \\ & \text{if }  \mathcal{S}  < t \text{ return } \perp \\ & \text{Parse } \mathcal{S} := \{i_1, \dots, i_t, \dots\} \\ & \text{for } k \in [\rho] \\ & \quad \text{for } j \in [t] \\ & \quad \quad \Lambda_{i_j}(X) := \prod_{\zeta \in [t] \setminus \{j\}} \frac{i_\zeta - X}{i_\zeta - i_j} \cdot (\Delta) \\ & \quad \quad s'_k := \sum_{j \in [t]} \Lambda_{i_j}(m + k) \cdot s^{(j)} \\ & \text{return } \mathbf{s}' := (s'_1, \dots, s'_\rho) \end{aligned}$
--	--

**Correctness.** Observe that for all  $i = 1, \dots, \rho$ , we have the following:

- $L_i(\text{pos}_i) = \Delta$
- $L_j(\text{pos}_i) = \Delta$  for all  $j \in [\rho], j \neq i$
- $f(\text{pos}_i) = \tilde{s}_i \cdot \Delta = s_i \cdot \Delta^2$

Meanwhile, for  $\lambda_{i_j}(X) := \prod_{\zeta \in [t] \setminus \{j\}} \frac{i_\zeta - X}{i_\zeta - i_j}$ , the polynomial we will be able to compute the polynomial  $f(x) = \sum_{j \in [t]} \lambda_{i_j} \cdot s^{(j)}$  by correctness of Lagrange Interpolation. Consequently,  $f(\text{pos}_i)$  would return  $s_i \cdot \Delta^2$ . However, we compute  $\Lambda_{i_j}$  instead, by multiplying with  $\Delta$  to remove need for division. Consequently, the resulting polynomial has  $\Delta$  multiplied throughout yielding a  $\Delta^3$  as the total offset.

**Definition 3** (Vector of Sweeping Polynomials). *Let  $\mathcal{C} \subset [m]$  such that  $|\mathcal{C}| = t - \rho$ . Then, we have a vector of sweeping polynomials, denoted by  $\text{sp}_{\mathcal{C}}(X) = (\text{sp}_{1,\mathcal{C}}, \dots, \text{sp}_{\rho,\mathcal{C}})$  where  $\text{sp}_{i,\mathcal{C}}(X) := \sum_{j=0}^{t-\rho} \text{sp}_{i,j} \cdot X^j \in \mathbb{Z}[X]_{\leq t-1}$  is the unique polynomial whose degree is at most  $t - 1$  such that  $\text{sp}_{i,\mathcal{C}}(m + i) = \Delta^2$ ,  $\text{sp}_{i,\mathcal{C}}(m + j) = 0$  for  $j \in [\rho], j \neq i$ , and  $\text{sp}_{i,\mathcal{C}}(j) = 0$  for all  $j \in \mathcal{C}$ . Further, one can define  $\text{sp}_{\max}$  as the upper bound for the coefficients for the sweeping polynomials, i.e.,  $\text{sp}_{\max} := \{\text{sp}_{i,j} \mid i \in \{1, \dots, \rho\}, j \in \{0, \dots, t - 1\}\}$*

**Lemma 6** (Existence of Sweeping Polynomial). *For any  $\mathcal{C} \subset [m]$  with  $|\mathcal{C}| = t - \rho$ , there exists  $\text{sp}_{\mathcal{C}} \in (\mathbb{Z}[X]_{\leq t-\rho})^\rho$  satisfying Definition 3.*

*Proof.* For any  $i = 1, \dots, \rho$ , we have that  $\text{sp}_{i,\mathcal{C}}(m + i) = \Delta^2$  and  $\text{sp}_{i,\mathcal{C}}(j) = 0$  for  $j \in \mathcal{C}$ . Let  $\mathcal{C} := (i_1, \dots, i_{t-\rho})$ . In other words, we can use these evaluations to construct a polynomial as follows:

$$\text{sp}_{i,\mathcal{C}}(X) := \Delta^2 \cdot \prod_{j=1}^{t-\rho} \frac{(X - i_j)}{(m + i) - i_j} \cdot \prod_{j \in [\rho] \setminus \{i\}} \frac{(X - (m + j))}{(i - j)}$$

Note that  $i_1, \dots, i_j \in [m]$  and are distinct. Therefore,  $\prod_{j=1}^{t-\rho} (m + i) - i_j$  perfectly divides  $\Delta$  and so does  $\prod_{j \in [\rho] \setminus \{i\}} (i - j)$ , which implies that the coefficients are all integers. Further, the degree of this polynomial is at most  $t - 1$ . Thus,  $\text{sp}_{i,\mathcal{C}}(X) \in \mathbb{Z}[X]_{t-1}$ . This defines the resulting vector of sweeping polynomials  $\text{sp}_{\mathcal{C}}$ .  $\square$

**Theorem 7.** *Construction 9 is statistically private provided*

$$\ell_r \geq \ell_s + \lceil \log_2(\text{sp}_{\max} \cdot (t - 1) \cdot \rho) \rceil + 1$$

*Proof.* Let  $\mathbf{s}, \mathbf{s}' \in [0, 2^{\ell_s}]^\rho$  be two vectors of secrets. Then,  $\tilde{\mathbf{s}} := \mathbf{s} \cdot \Delta$  and  $\tilde{\mathbf{s}}' := \mathbf{s}' \cdot \Delta$ . Let  $\mathcal{C}$  denote an arbitrary subset of corrupted parties of size  $|\mathcal{C}| = t - \rho$ . Further, let us assume that  $\tilde{\mathbf{s}}$  is shared using the polynomial  $f(X)$  as defined below:

$$f(X) := q(X) \cdot \prod_{k=1}^{\rho} (X - \text{pos}_k) + \sum_{k=1}^{\rho} \tilde{s}_k \cdot L_k(X)$$

where  $L_k(X) := \prod_{j \in [\rho] \setminus \{k\}} \frac{X - \text{pos}_j}{\text{pos}_k - \text{pos}_j} \cdot \Delta$ . and  $q(X)$  is a random polynomial of degree  $t - \rho - 1$ .

Now observe that the adversary see  $|\mathcal{C}| = t - \rho$  shares corresponding to  $f(i_j)$  for  $i_j \in \mathcal{C}$ . By Lagrange interpolation, this induces a one-to-one map from possible secrets to corresponding sharing polynomials. Specifically, we can use the vector of sweeping polynomials, as defined in Definition 3 to explicitly map any secret vector  $\mathbf{s}^*$  to its sharing polynomial defined by  $f(X) + \langle \mathbf{s}^* - \mathbf{s}, \mathbf{sp}_{\mathcal{C}}(X) \rangle$

In other words, the sharing polynomial to share  $\mathbf{s}^*$  is defined by

$$f^*(X) = f(X) + \sum_{k=1}^{\rho} (s_k^* - s_k) \cdot \text{sp}_{k,\mathcal{C}}(X)$$

One can verify the correctness. For example, to secret share  $s_1^*$ , at position  $m + 1$ , we get:

$$f^*(m + 1) = f(m + 1) + \sum_{k=1}^{\rho} (s_k^* - s_k) \cdot \text{sp}_{k,\mathcal{C}}(X)$$

Now, observe that  $f(m + 1) = s_1 \cdot (\Delta^2)$ . Meanwhile,  $\text{sp}_{1,\mathcal{C}}(m + 1) = \Delta^2$  while  $\text{sp}_{j,\mathcal{C}}(m + 1) = 0$  for  $1 < j \leq \rho$ . This simplifies to:  $f^*(m + 1) = s_1 \cdot \Delta^2 + (s_1^* - s_1) \cdot \Delta^2 = s_1^* \cdot \Delta^2$ . However, while we have an efficient mapping, note that  $f^*(X)$  could have coefficients that are not of the prescribed form, i.e., coefficients do not lie in the range  $[0, 2^{\ell_r + \kappa_s}]$ . We will call the event good if the coefficients lie in the range and bad even if one of the coefficients does not lie in the range.

Let us apply the above mapping to the secret  $\mathbf{s}'$  and we have the resulting polynomial:

$$f'(X) = f(X) + \sum_{k=1}^{\rho} (s'_k - s_k) \cdot \text{sp}_{k,\mathcal{C}}(X)$$

Now, observe that if  $f'(X)$  was a good polynomial, then  $f'(j) = f(j)$  for every  $j \in \mathcal{C}$ . It follows that if  $f'$  was good, then an adversary cannot distinguish whether the secret vector was  $\mathbf{s}$  or  $\mathbf{s}'$ .

We will now upper bound the probability that  $f'$  was bad in at least one of the coefficients. We know that  $|s'_k - s_k| \in [0, 2^{\ell_s})$  for  $k = 1, \dots, \rho$ . Further all coefficients of  $\text{sp}_{k,\mathcal{C}}(X)$  are upper bounded by  $\text{sp}_{\max}$ . Therefore, to any coefficient of  $f(X)$ , the maximum perturbation in value is:  $2^{\ell_s} \cdot \text{sp}_{\max} \cdot \rho$ . Therefore, one requires that the original coefficients of  $f$  be sampled such that they lie in  $[2^{\ell_s} \cdot \text{sp}_{\max} \cdot \rho, 2^{\ell_r + \kappa_s} - 2^{\ell_s} \cdot \text{sp}_{\max} \cdot \rho]$ . In other words, the probability that one coefficient of  $f'$  is bad is:

$$\frac{2 \cdot 2^{\ell_s} \cdot \text{sp}_{\max} \cdot \rho}{2^{\ell_r + \kappa_s}}$$

There are  $t - 1$  such coefficients. This gives us that the probability is  $\leq 2^{-\kappa_s}$  assuming that  $\ell_r \geq \ell_s + \lceil \log_2(\text{sp}_{\max} \cdot (t - 1) \cdot \rho) \rceil + 1$   $\square$

### C.3 Lattice-Based Assumptions

In this section, we will look at constructions based on three different lattice-based assumptions.

#### C.3.1 Learning with Rounding Assumption

We will begin by defining the learning with rounding (LWR) assumption, which can be viewed as a deterministic version of the learning with errors (LWE) assumption [78]. LWR was introduced by Banerjee *et al.* [9].

**Definition 4** (Learning with Rounding). *Let  $\rho, q, p \leftarrow \text{LWRGen}(1^\rho)$  with  $\rho, q, p \in \mathbb{N}$  such that  $q > p$ . Then, the Learning with Rounding assumption states that for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:*

$$\Pr \left[ b = b' \mid \begin{array}{l} \mathbf{s}, \mathbf{a}_0 \leftarrow \mathbb{Z}_q^\rho, \\ Y_0 := \lfloor \langle \mathbf{a}_0, \mathbf{s} \rangle \rfloor_p \\ Y_1 \leftarrow \mathbb{Z}_p, \mathbf{a}_1 \leftarrow \mathbb{Z}_q^\rho \\ b \leftarrow \{0, 1\}, b' \leftarrow \mathcal{A}(\mathbf{a}_b, Y_b) \end{array} \right] = \frac{1}{2} + \text{negl}(\rho)$$

where  $\lfloor x \rfloor_p = i$  where  $i \cdot \lfloor q/p \rfloor$  is the largest multiple of  $\lfloor q/p \rfloor$  that does not exceed  $x$ .



### C.3.2 Learning with Errors Assumption

**Definition 5** (Learning with Errors Assumption (LWE)). Consider integers  $\lambda, L, q$  and a probability distribution  $\chi$  on  $\mathbb{Z}_q$ , typically taken to be a normal distribution that has been discretized. Then, the  $LWE_{\lambda, L, q, \chi}$  assumption states that for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:

$$\Pr \left[ b = b' \mid \begin{array}{l} \mathbf{A} \leftarrow_s \mathbb{Z}_q^{L \times \lambda}, \mathbf{x} \leftarrow_s \mathbb{Z}_q^\lambda, \mathbf{e} \leftarrow_s \chi^L \\ \mathbf{y}_0 := \mathbf{A}\mathbf{x} + \mathbf{e} \\ \mathbf{y}_1 \leftarrow_s \mathbb{Z}_q^L, b \leftarrow_s \{0, 1\}, b' \leftarrow_s \mathcal{A}(\mathbf{A}, \mathbf{y}_b) \end{array} \right] = \frac{1}{2} + \text{negl}(\rho)$$

**Definition 6** (Hint-LWE [35, 46]). Consider integers  $\lambda, L, q$  and a probability distribution  $\chi'$  on  $\mathbb{Z}_q$ , typically taken to be a normal distribution that has been discretized. Then, the Hint-LWE assumption states that for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:

$$\Pr \left[ b = b' \mid \begin{array}{l} \mathbf{A} \leftarrow_s \mathbb{Z}_q^{L \times \lambda}, \mathbf{k} \leftarrow_s \mathbb{Z}_q^\lambda, \mathbf{e} \leftarrow_s \chi'^L \\ \mathbf{r} \leftarrow_s \mathbb{Z}_q^\lambda, \mathbf{f} \leftarrow_s \chi'^L \\ \mathbf{y}_0 := \mathbf{A}\mathbf{k} + \mathbf{e}, \mathbf{y}_1 \leftarrow_s \mathbb{Z}_q^L, b \leftarrow_s \{0, 1\} \\ b' \leftarrow_s \mathcal{A}(\mathbf{A}, (\mathbf{y}_b, \mathbf{k} + \mathbf{r}, \mathbf{e} + \mathbf{f})) \end{array} \right] = \frac{1}{2} + \text{negl}(\kappa)$$

where  $\kappa$  is the security parameter.

Kim *et al.* [59] demonstrates that the Hint-LWE assumption is computationally equivalent to the standard LWE assumption. In essence, this assumption posits that  $\mathbf{y}_0$  maintains its pseudorandom properties from an adversary's perspective, even when provided with certain randomized information about the secret and error vectors. Consider a secure LWE instance defined by parameters  $(\lambda, m, q, \chi)$ , where  $\chi$  represents a discrete Gaussian distribution with standard deviation  $\sigma$ . The corresponding Hint-LWE instance, characterized by  $(\lambda, m, q, \chi')$ , where  $\chi'$  denotes a discrete Gaussian distribution with standard deviation  $\sigma'$ , remains secure under the condition  $\sigma' = \sigma/\sqrt{2}$ . As a result, we can decompose any  $\mathbf{e} \in \chi$  into the sum  $\mathbf{e}_1 + \mathbf{e}_2$ , where both  $\mathbf{e}_1$  and  $\mathbf{e}_2$  are drawn from  $\chi'$ .

### C.4 Pseudorandom Functions

**Definition 7** (Pseudorandom Function (PRF)). A pseudorandom function family is defined by a tuple of PPT algorithms  $\text{PRF} = (\text{Gen}, \text{Eval})$  with the following definitions:

- $\text{pp}_{\text{PRF}} \leftarrow_s \text{Gen}(1^\kappa)$ : On input of the security parameter  $\kappa$ , the generation algorithm outputs the system parameters required to evaluate the function  $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  where  $\mathcal{K}$  is the key space,  $\mathcal{X}$  is the input space, and  $\mathcal{Y}$  is the output space.
- $y \leftarrow \text{Eval}(k, x)$ : On input of  $x \in \mathcal{X}$  and a randomly chosen key  $k \leftarrow_s \mathcal{K}$ , the algorithm outputs  $y \in \mathcal{Y}$  corresponding to the evaluation of  $F(k, x)$ .

We further require the following security property that: for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:

$$\Pr \left[ b = b' \mid \begin{array}{l} b \leftarrow_s \{0, 1\}, k \leftarrow_s \mathcal{K} \\ \mathcal{O}_0(\cdot) := F(k, \cdot), \mathcal{O}_1(\cdot) := U(\mathcal{Y}) \\ b' \leftarrow_s \mathcal{A}^{\mathcal{O}_b(\cdot)} \end{array} \right] \leq \frac{1}{2} + \text{negl}(\kappa)$$

where  $U(\mathcal{Y})$  outputs a randomly sampled element from  $\mathcal{Y}$ .

**Definition 8** ( $(\gamma)$ -Key Homomorphic PRF). Let PRF be a pseudorandom function that realizes an efficiently computable function  $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$  such that  $(\mathcal{K}, \oplus)$  is a group. Then, we say that it is

- key homomorphic if:  $(\mathcal{Y}, \otimes)$  is also a group and for every  $k_1, k_2 \in \mathcal{K}$  and every  $x \in \mathcal{X}$  we get:  $\text{Eval}(k_1, x) \otimes \text{Eval}(k_2, x) = \text{Eval}(k_1 \oplus k_2, x)$ .
- $\gamma = 1$ -almost key homomorphic if:  $\mathcal{Y} = \mathbb{Z}_p$  if for every  $k_1, k_2 \in \mathcal{K}$  and every  $x \in \mathcal{X}$ , there exists an error  $e \in \{0, 1\}$  we get:  $\text{Eval}(k_1, x) \otimes \text{Eval}(k_2, x) = \text{Eval}(k_1 \oplus k_2, x) + e$ .

### C.5 Distributed Key Homomorphic PRF

**Definition 9** (Distributed Key Homomorphic PRF (DPRF)). A  $(t, m)$ -Distributed PRF is a tuple of PPT algorithms  $\text{DPRF} := (\text{Gen}, \text{Share}, \text{Eval}, \text{P-Eval}, \text{Combine})$  with the following syntax:



- $\text{pp}_{\text{PRF}} \leftarrow \text{Gen}(1^\kappa, 1^t, 1^m)$ : On input of the threshold  $t$  and number of parties  $m$ , and security parameter  $\kappa$ , the  $\text{Gen}$  algorithm produces the system parameter  $\text{pp}_{\text{PRF}}$  which is implicitly consumed by all the other algorithms.
- $k^{(1)}, \dots, k^{(m)} \leftarrow \text{Share}(k, t, r, m)$ : On input of the number of parties  $m$ , corruption threshold  $t$ , reconstruction threshold  $r$ , and a key  $k \leftarrow \mathcal{K}$ , the share algorithm produces the key share for each party.
- $Y \leftarrow \text{Eval}(k, x)$ : On input of the PRF key  $k$  and input  $x$ , the algorithm outputs  $y$  corresponding to some pseudorandom function  $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ .
- $y_i \leftarrow \text{P-Eval}(k^{(i)}, x)$ : On input of the PRF key share  $k^{(i)}$ , the partial evaluation algorithm outputs a partial evaluation  $y_i$  on input  $x \in \mathcal{X}$ .
- $Y' \leftarrow \text{Combine}(\{y_i\}_{i \in \mathcal{S}})$ : On input of partial evaluations  $y_i$  corresponding to some subset of shares  $\mathcal{S}$  such that  $|\mathcal{S}| \geq t$ , the algorithm outputs  $Y'$ .

We further require the following properties:

- **Correctness:** We require that the following holds for any  $m, t, r, \kappa \in \mathbb{Z}$  with  $t < r \leq m$  and any set  $\mathcal{S} \subseteq [m]$  with  $|\mathcal{S}| \geq r$ , any input  $x \in \mathcal{X}$ :

$$\Pr \left[ Y = Y' \mid \begin{array}{l} \text{pp}_{\text{PRF}} \leftarrow \text{Gen}(1^\kappa, 1^t, 1^r, 1^m), k \leftarrow \mathcal{K} \\ \{k^{(i)}\}_{i \in [m]} \leftarrow \text{Share}(k), \{y_i \leftarrow \text{P-Eval}(k^{(i)}, x)\}_{i \in \mathcal{S}} \\ Y' \leftarrow \text{Combine}(\{y_i\}_{i \in \mathcal{S}}, Y = \text{Eval}(k, x)) \end{array} \right] = 1$$

- **Pseudorandomness with Static Corruptions:** We require that for any integers  $t, r, m$  with  $t < r \leq m$ , and for all PPT adversary  $\mathcal{A}$ , there exists a negligible function  $\text{negl}$  such that:

$$\Pr \left[ \begin{array}{l} b = b' \\ |\mathbf{K} \cup \{j : (j, x^*) \in \mathbf{E}\}| \leq t \end{array} \mid \begin{array}{l} \text{pp}_{\text{PRF}} \leftarrow \text{Gen}(1^\kappa, 1^t, 1^r, 1^m), k \leftarrow \mathcal{K} \\ (st, \mathbf{K}) \leftarrow \mathcal{A}(\text{pp}_{\text{PRF}}) \\ \{k^{(i)}\}_{i \in [m]} \leftarrow \text{Share}(k, t, r, m, ) \\ (st, x^*) \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Eval}}}(\{k^{(i)}\}_{i \in \mathbf{K}}) \\ b \leftarrow \{0, 1\}, Y_0 \leftarrow \text{Eval}(k, x^*) \\ Y_1 \leftarrow \mathcal{U}(\mathcal{Y}), b' \leftarrow \mathcal{A}^{\mathcal{O}_{\text{Eval}}}(st, Y_b) \end{array} \right] = 1$$

where:

$$\begin{array}{l} \mathcal{O}_{\text{Eval}}(i, x) \\ \mathbf{E} := \mathbf{E} \cup \{(i, x)\} \\ \text{return P-Eval}(k^{(i)}, x) \end{array}$$

- **Key Homomorphic:** We require that if  $(\mathcal{K}, \oplus), (\mathcal{Y}, \otimes)$  are groups such that:
  - $\forall x \in \mathcal{X}, \forall k_1, k_2 \in \mathcal{K}, \text{Eval}(k_1, x) \otimes \text{Eval}(k_2, x) = \text{Eval}(k_1 \oplus k_2, x)$ , and
  - $\forall x \in \mathcal{X}, \forall k_1, k_2 \in \mathcal{K}, \left\{ k_b^{(j)} \leftarrow \text{Share}(k_b) \right\}_{j \in [m], b \in \{1, 2\}}$ ,
  - $\forall j \in [m], y_{1,2}^{(j)} := \left( \text{P-Eval}(k_1^{(j)}, x) \otimes \text{P-Eval}(k_2^{(j)}, x) \right)$ , and  $\forall \mathcal{S} \subseteq [m]$  with  $|\mathcal{S}| \geq r$ ,
  - $\text{Combine} \left( \left\{ y_{1,2}^{(j)} \right\}_{j \in \mathcal{S}} \right) = \text{Eval}(k_1 \oplus k_2, x)$

## C.6 Class Groups Framework

Class group-based cryptography is a cryptographic technique that originated in the late 1980s, with the idea that the class group of ideals of maximal orders of imaginary quadratic fields may be more secure than the multiplicative group of finite fields [23, 69]. The CL framework was first introduced by the work of Castagnos and Laguillaumie [32]. This framework operates on a group where there exists a subgroup with support for efficient discrete logarithm construction. Subsequent works

[33, 29, 30, 87, 21] have refined the original framework. The framework has been used in various applications over the years [33, 29, 30, 91, 45, 51, 86, 21, 58]. Meanwhile, class group cryptography itself has been employed in numerous applications [89, 90, 63, 18, 31, 34, 60, 24, 7, 1, 43, 42, 8].

Broadly, the framework is defined by two functions - CLGen, CLSolve with the former outputting a tuple of public parameters. The elements of this framework are the following:

- *Input Parameters:*  $\kappa_c$  is the computational security parameter,  $\kappa_s$  is a statistical security parameter, a prime  $p$  such that  $p > 2^{\kappa_c}$ , and uniform randomness  $\rho$  that is used by the CLGen algorithm and is made public.
- *Groups:*  $\widehat{\mathbb{G}}$  is a finite multiplicative abelian group,  $\mathbb{G}$  is a cyclic subgroup of  $\widehat{\mathbb{G}}$ ,  $\mathbb{F}$  is a subgroup of  $\mathbb{G}$ ,  $H = \{x^p, x \in \mathbb{G}\}$
- *Orders:*  $\mathbb{F}$  has order  $p$ ,  $\widehat{\mathbb{G}}$  has order  $p \cdot \widehat{s}$ ,  $\mathbb{G}$  has order  $p \cdot s$  such that  $s$  divides  $\widehat{s}$  and  $\gcd(p, \widehat{s}) = 1, \gcd(p, s) = 1$ ,  $H$  has order  $s$  and therefore  $\mathbb{G} = \mathbb{F} \times H$ .
- *Generators:*  $f$  is the generator of  $\mathbb{F}$ ,  $g$  is the generator of  $\mathbb{G}$ , and  $h$  is the generator of  $H$  with the property that  $g = f \cdot h$
- *Upper Bound:* Only an upper bound  $\bar{s}$  of  $\widehat{s}$  (and  $s$ ) is provided.
- *Additional Properties:* Only encodings of  $\widehat{\mathbb{G}}$  can be recognized as valid encodings and  $s, \widehat{s}$  are unknown.
- *Distributions:*  $\mathcal{D}$  (resp.  $\mathcal{D}_p$ ) be a distribution over the set of integers such that the distribution  $\{g^x : x \leftarrow \mathcal{D}\}$  (resp.  $\{g_p^x : x \leftarrow \mathcal{D}_p\}$ ) is at most distance  $2^{-\kappa_s}$  from the uniform distribution over  $\mathbb{G}$  (resp.  $H$ ).
- *Additional Group and its properties:*  $\widehat{\mathbb{G}}^p = \{x^p, x \in \widehat{\mathbb{G}}\}$ ,  $\widehat{\mathbb{G}}$  factors as  $\widehat{\mathbb{G}}^p \times \mathbb{F}$ .<sup>4</sup> Let  $\bar{\omega}$  be the group exponent of  $\widehat{\mathbb{G}}^p$ . Then, the order of any  $x \in \widehat{\mathbb{G}}^p$  divides  $\bar{\omega}$ .<sup>5</sup>

**Remark 4.** The motivations behind these additional distributions are as follows. One can efficiently recognize valid encodings of elements in  $\widehat{\mathbb{G}}$  but not  $\mathbb{G}$ . Therefore, a malicious adversary  $\mathcal{A}$  can run our constructions by inputting elements belonging to  $\widehat{\mathbb{G}}^p$  (rather than in  $H$ ). Unfortunately, this malicious behavior cannot be detected which allows  $\mathcal{A}$  to obtain information on the sampled exponents modulo  $\bar{\omega}$  (the group exponent of  $\widehat{\mathbb{G}}^p$ ). By requiring the statistical closeness of the induced distribution to uniform in the aforementioned groups allows flexibility in proofs. Note that the assumptions do not depend on the choice of these two distributions. Further, the order  $s$  of  $H$  and group exponent  $\bar{\omega}$  of  $\widehat{\mathbb{G}}^p$  are unknown and the upper bound  $\bar{s}$  is used to instantiate the aforementioned distribution. Specifically, looking ahead we will set  $\mathcal{D}_H$  to be the uniform distribution over the set of integers  $[B]$  where  $B = 2^{\kappa_s} \cdot \bar{s}$ . Using Lemma 8, we get that the distribution is less than  $2^{-\kappa_s}$  away from uniform distribution in  $H$ . In our constructions we will set  $\kappa_s = 40$ . We will make this sampling more efficient for our later constructions. We refer the readers to Tucker [87, §3.1.3, §3.7] for more discussions about this instantiation. Finally, as stated we will also set  $\widehat{\mathcal{D}} = \mathcal{D}$  and  $\widehat{\mathcal{D}}_H = \mathcal{D}_H$ .

We also have the following lemma from Castagnos, Imbert, and Laguillaumie [31] which defines how to sample from a discrete Gaussian distribution.

**Lemma 8.** *Let  $\mathbb{G}$  be a cyclic group of order  $n$ , generated by  $g$ . Consider the random variable  $X$  sampled uniformly from  $\mathbb{G}$ ; as such it satisfies  $\Pr[X = h] = \frac{1}{n}$  for all  $h \in \mathbb{G}$ . Now consider the random variable  $Y$  with values in  $\mathbb{G}$  as follows: draw  $y$  from the discrete Gaussian distribution  $\mathcal{D}_{\mathbb{Z}, \sigma}$  with  $\sigma \geq n \sqrt{\frac{\ln(2(1+1/\epsilon))}{\pi}}$  and set  $Y := g^y$ . Then, it holds that:*

$$\Delta(X, Y) \leq 2\epsilon$$

**Remark 5.** By definition, the distribution  $\{g^x : x \leftarrow \mathcal{D}\}$  is statistically indistinguishable from  $\{g^y : y \leftarrow \{0, \dots, p \cdot s - 1\}\}$ . Therefore, it follows that  $\{x \bmod p \cdot s : x \leftarrow \mathcal{D}\}$  is statistically indistinguishable from  $\{x : x \leftarrow \{0, \dots, p \cdot s - 1\}\}$ . Similarly,  $\{x \bmod s : x \leftarrow \mathcal{D}_p\}$  is statistically indistinguishable from  $\{x : x \leftarrow \{0, \dots, s - 1\}\}$ . Furthermore, sampling a value  $x$  corresponding to  $\mathcal{D}$  is statistically indistinguishable from the uniform distribution in  $\{0, \dots, s - 1\}$  because  $s$  divides  $p \cdot s$ .

<sup>4</sup>Recall that  $p$  and  $\widehat{s}$  are co-prime.

<sup>5</sup>This follows from the property that the exponent of a finite Abelian group is the least common multiple of its elements.

**Definition 10** (Class Group Framework). *The framework is defined by two algorithms (CLGen, CLSolve) such that:*

- $\text{pp}_{\text{CL}} = (p, \kappa_c, \kappa_s, \bar{s}, f, h, \widehat{\mathbb{G}}, \mathbb{F}.\mathcal{D}, \mathcal{D}_H, \widehat{\mathcal{D}}, \widehat{\mathcal{D}}_H, \rho) \leftarrow_{\$} \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s}, p; \rho)$
- *The DL problem is easy in  $\mathbb{F}$ , i.e., there exists a deterministic polynomial algorithm CLSolve that solves the discrete logarithm problem in  $\mathbb{F}$ :*

$$\Pr \left[ x = x' \mid \begin{array}{l} \text{pp}_{\text{CL}} = \leftarrow_{\$} \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s}, p; \rho) \\ x \leftarrow_{\$} \mathbb{Z}/p\mathbb{Z}, X = f^x; \\ x' \leftarrow \text{CLSolve}(\text{pp}_{\text{CL}}, X) \end{array} \right] = 1$$

**Definition 11** (Hard Subgroup Membership Assumption ( $\text{HSM}_{\mathbb{M}}$ ) Assumption [33]). *Let  $\kappa$  be a the security parameter with prime  $p$  such that  $|p| \geq \kappa$ . Let (CLGen, CLSolve) be the group generator algorithms as defined in Definition 10, then the  $\text{HSM}_{\mathbb{M}}$  assumption requires that that  $\text{HSM}_{\mathbb{M}}$  problem be hard in  $\mathbb{G}$  even with access to the Solve algorithm. More formally, let  $\mathcal{D}$  (resp.  $\mathcal{D}_p$ ) be a distribution over the set of integers such that the distribution  $\{g^x : x \leftarrow_{\$} \mathcal{D}\}$  (resp.  $\{g_p^x : x \leftarrow_{\$} \mathcal{D}_p\}$ ) is at most distance  $2^{-\kappa}$  from the uniform distribution over  $\mathbb{G}$  (resp.  $H$ ). Then, we say that the  $\text{HSM}_{\mathbb{M}}$  problem is hard if for all PPT adversaries  $\mathcal{A}$ , there exists a negligible function  $\text{negl}(\kappa)$  such that:*

$$\Pr \left[ b = b' \mid \begin{array}{l} \text{pp}_{\text{CL}} \leftarrow_{\$} \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s}, p; \rho) \\ x \leftarrow_{\$} \mathcal{D}, x' \leftarrow_{\$} \mathcal{D}_p \\ b \leftarrow_{\$} \{0, 1\}; Z_0 = g^x; Z_1 = g_p^{x'} \\ b' \leftarrow_{\$} \mathcal{A}^{\text{Solve}(\text{pp}_{\text{CL}}, \cdot)}(\text{pp}_{\text{CL}}, Z_b) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\kappa)$$

When dealing with groups of known order, one can sample elements in a group  $\mathbb{G}$  easily by merely sampling exponents modulo the group order and then raising the generator of the group to that exponent. Unfortunately, note that here neither the order of  $\mathbb{G}$  (i.e.,  $ps$ ) nor that of  $H$  (i.e.,  $s$ ) is known. Therefore, we instead use the knowledge of the upper-bound  $\bar{s}$  of  $s$  to instantiate the distributions  $\mathcal{D}$  and  $\mathcal{D}_p$  respectively. This choice of choosing from the distributions  $\mathcal{D}$  and  $\mathcal{D}_p$  respectively allows for flexibility of various proofs.

## D Naive Construction of OPA

We present a naive construction that achieves our desired one-shot behavior. This construction is built from Shamir's Secret Sharing algorithm, as recalled in Construction 6. This shall serve as the baseline construction.

- Setup: Simply runs the setup algorithm for the secret sharing algorithm SS
- Enc: Client  $i$ , with inputs  $x_i$  secret shares the input, for each element in  $\mathbf{x}_i$ . In other words, for  $i = 1, \dots, L$ , do  $\{x_i^{(j)}\}_{j \in [m]} \leftarrow_{\$} \text{SS.Share}(x_i, t, r, m)$

It sets  $\mathbf{k}_i := \emptyset$ ,  $\mathbf{ct}_i := \emptyset$  and  $\text{aux}_i^{(j)} := \{x_i^{(j)}\}_{i=1}^L$

- C-Combine: Committee member  $j$ , simply adds up all the shares it has received, element-by-element.  
 $\text{AUX}^{(j)} \sum_{i=1}^n \text{aux}_i^{(j)}$
- S-Combine: Upon receiving from at least  $r$  committee members, it runs  $\text{SS.Reconstruct}(\{ \text{AUX}^{(j)} \}_{j \in \mathcal{S}})$ , element-by-element, to get the sum of vectors.
- Aggregate: The previous algorithm already gives the sum.

This algorithm has the following communication:

- Client  $i$  to Committee:  $L$  field elements per committee member for a total of  $mL$  field elements.
- Client  $i$  to Server: 0
- Committee  $j$  to Server:  $L$  field elements.
- Total Sent/Received per Client:  $mL$  field elements.

## Protocol Naive Construction of OPA

### One-Time System Parameters Generation

Run  $pp \leftarrow \text{SS.Setup}(1^\kappa, 1^t, 1^r, 1^m)$   
 Output Committee of size  $m$ .

### Data Encryption Phase by Client $i$ in iteration $\ell$

Let  $\mathbf{x}_{i,\ell}$  be the input of Client  $i$  in iteration  $\ell$ .  
**for**  $in = 1, \dots, L$  **do**  
      $ct_{i,\ell} = \perp, \left\{ \text{aux}_{i,\ell}^{(j)}[in] \right\}_{j \in [m]} \leftarrow \text{SS.Share}(x_{i,\ell}[in], t, m, \ell)$   
 Send  $ct_{i,\ell}$  to the Server  
 Send  $\text{aux}_{i,\ell}^{(j)}$  to committee member  $j$  for each  $j \in [m]$ .

### Set Identification Phase by Committee Member $j$ in iteration $\ell$

Let  $\left\{ \text{aux}_{i,\ell}^{(j)} \right\}$  be the inputs received by Committee member  $j$  from Clients  $i \in \mathcal{C}^{(j)}$   
 Send  $\mathcal{C}^{(j)}$  to server // This can be a bit string of length  $n$  where  $i$ -th bit is 1 if  $i \in \mathcal{C}^{(j)}$

### Set Intersection Phase by Server in iteration $\ell$

Let  $\left\{ \mathcal{C}^{(j)} \right\}_{j \in \mathcal{S}}$  be the client sets received from committee members  $j \in \mathcal{S}$ .  
**assert**  $|\mathcal{S}| \geq r$   
 Compute  $\mathcal{C} := \bigcap_{j \in \mathcal{S}} \mathcal{C}^{(j)}$  // This is bit-wise AND operation of the  $r$  bit strings.  
 Send  $\mathcal{C}$  to committee members in  $\mathcal{S}$

### Data Combination Phase by Committee Member $j$ in iteration $\ell$

Let  $\left\{ \text{aux}_{i,\ell}^{(j)} \right\}$  be the inputs received by Committee member  $j$  from Clients  $i \in \mathcal{C}$   
**for**  $in = 1, \dots, L$  **do**  
     Compute  $\text{AUX}^{(j)}[in] \leftarrow \sum_{i \in \mathcal{C}} \text{aux}_{i,\ell}^{(j)}[in]$   
 Send  $\text{AUX}^{(j)}$  to server

### Data Aggregation Phase by Server in iteration $\ell$

**for**  $in = 1, \dots, L$  **do**  
     Let  $\left\{ \text{AUX}_\ell^{(j)}[in] \right\}_{j \in \mathcal{S}}$  be the inputs received by the server with  $|\mathcal{S}| \geq t$ .  
     Run  $X_\ell[in] \leftarrow \text{SS.Reconstruct}\left(\left\{ \text{AUX}_\ell^{(j)}[in] \right\}_{j \in \mathcal{S}}\right)$

Figure 7: Our Naive Construction of OPA using just a secret sharing scheme.

- Total Sent/Received by Server:  $mL$  field elements.
- Total Sent/Received per Committee Member:  $L + nL$  field elements.

**Remark 6.** It is important to make a few observations. This solution resembles a multi-server solution where the clients simply secret-share the inputs with each server, and the servers collaborate to recover the sum. However, this observation also implies that the role of the committee members, while running a deterministic procedure, is proportional to the length of the input. This is undesirable as these committee members are simply other clients and such a heavy computation is undesirable.

It is also to be noted that even a simple solution such as using a one-time pad to mask the input, and then secret share the mask suffers from the same bottleneck as the mask has to be at least the length of the input.

**Computation Performance of  $\text{OPA}_{\text{LWR}}$  vs Naive Solution.** We now list the computation costs that were measured. Note that these experiments were done for  $L = 1$ . The motivation behind these experiments was to study how the two constructions fared across the computation time incurred by the Naive Solution and  $\text{OPA}_{\text{LWR}}$ , without any optimizations including packed secret sharing.

Our experiments indicate that, in addition to the significant savings in communication cost,  $\text{OPA}_{\text{LWR}}$  outperforms the Naive Solution when it comes to the performance of the committee. While the role of the committee is the same in both constructions, the cost of the committee for  $\text{OPA}_{\text{LWR}}$  does not scale linearly with the input, while the naive solution incurs a linear growth cost. Further, note that currently,  $\text{OPA}_{\text{LWR}}$  requires the committee to add shares of  $2^{11}$  elements. It is clear when




Figure 8: Comparison of Running Times between  $\text{OPA}_{\text{LWR}}$  and the Naive Solution. Values that scale linearly with  $n$  are marked with a slash on the entries.

aggregating vectors of length  $> 2^{11}$ ,  $\text{OPA}_{\text{LWR}}$  will outperform the naive solution. Meanwhile, it seems clear that  $\text{OPA}_{\text{LWR}}$ 's client cost which involves secret sharing the seed and then masking a vector of length  $L$  will be slower than simply secret sharing a vector of length  $L$ . Similarly, the cost of simply reconstructing the summed up vectors of length for the server in the naive solution will be better than the multi-step process of  $\text{OPA}_{\text{LWR}}$  which includes reconstructing the seed, then evaluating the said seed, before finally unmasking.

For completeness, we plot the server and client costs comparison in Figure 8. Note that we have scaled the costs, in the case of the naive solution, by 100 (i.e., if the actual cost was 1 unit, we plotted it as 100 units).

## E Constructions in CL Framework

**Construction 10** (PRF in CL Framework). Let  $(\text{CLGen}, \text{CLSolve})$  be the class group framework as defined in Definition 10. Then, let  $\text{pp}_{\text{CL}} \leftarrow_s \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s})$ . Further, let  $H : \mathcal{X} \rightarrow \mathbb{H}$ ,  $H' : \mathcal{X} \rightarrow \mathcal{K}$  be a hash function. Then, consider the following definition of  $\mathcal{K} = \mathcal{D}_H$ ,  $\mathcal{X} = \{0, 1\}^*$ ,  $\mathcal{Y} = \mathbb{H}$ ,  $F_{\text{CL}}(k, x) = \mathbb{H}(x)^{k \cdot H'(x, 1)}$ .

**Remark 7.** Note that the order of  $\mathbb{H}$  is unknown. Therefore, one has to rely on  $\mathcal{D}_H$  to hash into  $\mathbb{H}$ . Most recently, [79] showed how to hash into groups of unknown order to ensure that discrete logarithm is unknown. However, for our applications, this is not a concern. Indeed, one can simply compute the hash function as  $h^{H'(x)}$  where  $H'(x)$  hashed in  $\mathcal{D}_H$ .

**Theorem 9.** Construction 10 is a secure PRF where  $H$  is modeled as a random oracle under the  $\text{HSM}_{\text{M}}$  assumption.

*Proof.* We denote the challenger by  $\mathcal{B}$ . Let  $S_j$  be the event that the adversary wins in  $\text{Hybrid}_j$  for each  $j \in \{0, \dots, 2\}$ . Let  $q_e$  (resp.  $q_h$ ) denote the number of evaluation queries (resp. hash oracle queries) that the adversary makes. We use an analysis similar to the technique by Coron [39].

$\text{Hybrid}_0(\kappa)$ : Corresponds to the security game as defined for the security of PRF. It follows that the advantage of the adversary is

$$\text{Adv}_0 = 2 \cdot |\Pr[S_0] - 1/2| = \text{Adv}_{\mathcal{A}}^{\text{PRF}}$$

$\text{Hybrid}_1(\kappa)$ : This game is identical to  $\text{Hybrid}_1$  with the following difference. The challenger tosses biased coin  $\delta_t$  for each random oracle query  $H(t)$ . The biasing of the coin is as follows: takes a value 1 with probability  $\frac{1}{q_e+1}$  and 0 with probability  $\frac{q_e}{q_e+1}$ . Then, one can consider the following event  $E$ : that the adversary makes a query to the random oracle with  $x_i$  as an input where  $x_i$  was one of the evaluation inputs and for this choice we have that  $\delta_t$  was flipped to 0.

If  $E$  happens, the challenger halts and declares failure. Then, we have that:

$$\Pr[\neg E] = \left( \frac{q_e}{q_e + 1} \right)^{q_e} \geq \frac{1}{e(q_e + 1)}$$

where  $e$  is the Napier's constant. Finally, we get that:

$$\Pr[S_1] = \Pr[S_0] \cdot \Pr[\neg E] \geq \frac{\Pr[S_0]}{e(q_e + 1)}$$

**Hybrid<sub>2</sub>( $\kappa$ ):** This game is similar to **Hybrid<sub>1</sub>** with the following difference: we modify the random oracle outputs.

- If  $\delta_t = 0$ , the challenger samples  $w_t \leftarrow \mathcal{D}_H$  and sets  $H(t) = h^{w_t}$
- If  $\delta_t = 1$ , the challenger samples  $w_t \leftarrow \mathcal{D}_H, u_t \leftarrow \mathbb{Z}/M\mathbb{Z}$  and sets  $H(t) = h^{w_t} \cdot f^{u_t}$

Note that, under the HSM assumption, an adversary cannot distinguish between the two hybrids. Therefore, we get:

$$|\Pr[S_2] - \Pr[S_1]| \leq \epsilon_{\text{HSM}_M}$$

where  $\epsilon_{\text{HSM}_M}$  is the advantage that an adversary has in the  $\text{HSM}_M$  game. Note that **Hybrid<sub>2</sub>** corresponds to the case where the outputs are all random elements in  $\mathbb{G}$ . Therefore, the inputs are sufficiently masked and leak no information about the key. Therefore,  $\Pr[S_2] = 0$ . Then,

$$\text{Adv}_{\mathcal{A}}^{\text{PRF}} \leq (e \cdot (q_e + 1)) \cdot \epsilon_{\text{HSM}_M}$$

□

**Remark 8.** Note that the above scheme is simply an adaptation of the famous DDH-based construction of a key-homomorphic PRF that was shown to be secure by Naor *et al.* [70]. It is easy to verify that our construction is also key homomorphic as  $H(x)^{(k_1+k_2)} = H(x)^{k_1} \cdot H(x)^{k_2}$ .

## E.1 Distributed PRF in CL Framework

We build our construction of Distributed PRF from the Linear Secret Sharing Scheme LISS := (Share, GetCoeff, Reconstruct) with  $\text{pp}_{\text{SS}}$  denoting the public parameters of the LISS scheme. We specifically employ the Shamir Secret Sharing scheme over the Integers, as defined in [21].

**Construction 11** (Distributed PRF in CL Framework). A  $(r, m)$ -Distributed PRF is a tuple of PPT algorithms  $\text{DPRF} := (\text{Gen}, \text{Share}, \text{Eval}, \text{P-Eval}, \text{Combine})$  with the algorithms as defined in Figure 9. For simplicity, in the construction below we will set the corruption threshold  $t = r - 1$ . Though, the construction also holds for a lower  $t$ .

**Theorem 10.** *In the Random Oracle Model, if Construction 10 is a secure pseudorandom function if Integer Secret Sharing is statistically private, then Construction 11 is pseudorandom in the static corruptions setting.*

**Correctness.** For a polynomial  $f \in \mathbb{Z}[X]$ , every  $f(i)$  leaks information about the secret  $s \bmod i$  leading to a choice of polynomial  $f$  such that  $f(0) = \Delta \cdot s$ . For our use case, the secret is the PRF key  $k$ . Let us consider a set  $\mathcal{S} = \{i_1, \dots, i_r\}$  of indices and corresponding evaluations of the polynomial  $f$  at  $i_1, \dots, i_r$  giving us key shares:  $k^{(i_1)}, \dots, k^{(i_r)}$ . To begin with, one can compute the Lagrange coefficients corresponding to the set  $\mathcal{S}$  as:  $\forall i \in \mathcal{S}, \lambda_i(X) := \prod_{j \in \mathcal{S} \setminus \{i\}} \frac{x_j - X}{x_j - x_i}$ . This implies that the resulting polynomial is  $f(X) := \sum_{j=1}^r \lambda_{i_j}(X) \cdot k^{(i_j)}$ .

However,  $\lambda_i(X)$  requires one to perform a division  $x_j - x_i$  which is undefined as  $H$  hashes to  $\mathbb{G}$  whose order is unknown. To avoid this issue, a standard technique is to instead compute coefficient  $\Lambda_i(X) := \Delta \cdot \lambda_i(x)$ . Thereby, the resulting polynomial that is reconstructed if  $f'(X) = \Delta \cdot f(X) = \sum_{j=1}^r \Lambda_{i_j}(X) \cdot k^{(i_j)}$ . Consequently,

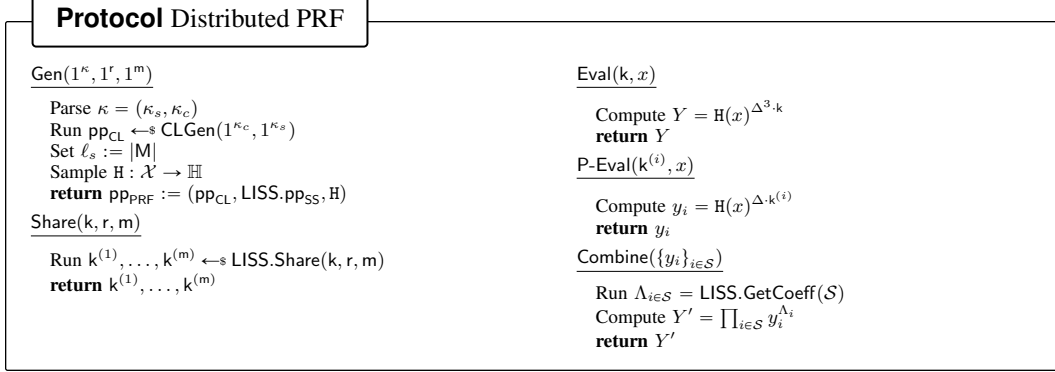


Figure 9: Construction of Distributed PRF based on the LISS := (Share, GetCoeff, Reconstruct) scheme of [21], with  $\text{pp}_{\text{SS}}$  denoting the public parameters of the LISS scheme. Recall that the offset  $\Delta := m!$  where  $m$  is the number of shares generated.

$$\begin{aligned}
 \mathbb{H}(x)^{\Delta^3 \cdot k} &= \mathbb{H}(x)^{\Delta \cdot f'(0)} = \mathbb{H}(x)^{\Delta \cdot \sum_{j=1}^r \Lambda_{i_j}(0) \cdot k^{(i_j)}} \\
 &= \prod_{j=1}^r \left( \mathbb{H}(x)^{\Delta \cdot k^{(i_j)}} \right)^{\Lambda_{i_j}(0)} \\
 &= \prod_{j=1}^r \left( \text{P-Eval}(k^{(i_j)}, x) \right)^{\Lambda_{i_j}(0)}
 \end{aligned}$$

Thus, our protocol is correct.

**Pseudorandomness.** Next, we consider the pseudorandomness property of our construction.

**Theorem 10.** *In the Random Oracle Model, if Construction 10 is a secure pseudorandom function if Integer Secret Sharing is statistically private, then Construction 11 is pseudorandom in the static corruptions setting.*

Boneh *et al.* [19] showed that from any Key Homomorphic PRF (which Construction 10), one can build a Distributed PRF. The proof of the following theorem follows the template of this scheme with certain important adaptations as our secret sharing scheme is over integers. The proof technique is to show that if there exists an adversary  $\mathcal{A}$  that can break the DPRF security, one can then use it to build an adversary  $\mathcal{B}$  to break the pseudorandomness of our original PRF, as defined in Construction 10. The idea behind the proof is for  $\mathcal{B}$ , upon receiving choice of  $r - 1$  corruptions as indices  $i_1, \dots, i_{r-1}$ , to then choose a random index  $i_r$  and implicitly set  $k^{(i_r)}$  to be the PRF key chosen by its challenger. Therefore, the  $\mathcal{B}$  now has knowledge of  $r$  indices, with which it can sample the Lagrange coefficients as before:

**for**  $j = 1, \dots, r$  **do**

$$\Lambda_{i_j}(X) := \prod_{\zeta \in \{1, \dots, r\} \setminus j} \frac{i_\zeta - X}{i_\zeta - i_j} \cdot (\Delta)$$

Now,  $\mathcal{B}$  with knowledge of the keys for indices  $i_1, \dots, i_{r-1}$  along with access to Oracle needs to simulate valid responses to P-Eval queries for an unknown index. Call this index  $i^*$ . Then, we have:

$$\begin{aligned}
 \text{P-Eval}(k^{(i^*)}, x) &:= \mathbb{H}(x)^{\Delta \cdot k^{(i^*)}} = \mathbb{H}(x)^{\sum_{j=1}^r \Lambda_{i_j}(i^*) \cdot k^{(i_j)}} \\
 &= \mathbb{H}(x)^{\sum_{i=1}^{r-1} \Lambda_{i_j}(i^*) \cdot k^{(i_j)}} \cdot \left( \mathbb{H}(x)^{k^{(i_r)}} \right)^{\Lambda_{i_r}(i^*)}
 \end{aligned}$$

The last term is simulated using  $\mathcal{B}$ 's own oracle access.



*Proof.* Let  $\mathcal{A}$  be a PPT attacker against the pseudorandomness property of DPRF, having advantage  $\epsilon$ .

$\mathcal{A}$  first chooses  $r - 1$  indices  $\mathbf{K} = \{i_1, \dots, i_{r-1}\}$  where each index is a subset of  $\{1, \dots, m\}$ .  $\mathcal{A}$  receives the shares of the keys  $k^{(i_1)} = f(i_1), \dots, k^{(i_{r-1})} = f(i_{r-1})$  (for unknown polynomial  $f$  of degree  $r$  such that  $f(0) = k \cdot \Delta$  with  $\Delta := \Delta$ ). Further,  $\mathcal{A}$  has access to  $\mathcal{O}_{\text{Eval}}(i, x)$  receiving P-Eval( $k_i, x$ ) ins response. Additionally,  $\mathcal{A}$  expects to have oracle access to the random oracle  $H$ .

Using this attacker  $\mathcal{A}$ , we now define a PPT attacker  $\mathcal{B}$  which will break the pseudorandomness property of Construction 10. Note that  $\mathcal{B}$  is given access to the oracle that either outputs the real evaluation of the PRF on key  $k^*$  or a random value. Additionally,  $\mathcal{B}$  expects to have oracle access to the random oracle  $H$ .

- **Setup:**  $\mathcal{B}$  does the following during Setup.
  - Receive set  $S = \{i_1, \dots, i_{r-1}\}$  from  $\mathcal{A}$ .
  - Next  $\mathcal{B}$  generates the key shares and public key as follows:
    - \* Sample  $k^{(i_1)}, \dots, k^{(i_{r-1})} \in \mathbb{Z}$ .
    - \*  $\mathcal{B}$  picks an index  $i_r$  at random and implicitly sets the PRF key chosen by its challenger as  $k^{(i_r)}$ .
    - \* Immediately, given the  $r$  indices, one can construct the secret sharing polynomial  $f \in \mathbb{Z}[X]$  as described earlier, but instead recreating the polynomial  $f'(X)$  using the coefficients  $\Lambda_{i_j}(X)$  for  $j = 1, \dots, r$  with  $k^{(i_r)}$  being unknown to  $\mathcal{B}$  and using its challenger to simulate a response.
    - \*  $\mathcal{B}$  gives  $k^{(i_1)}, \dots, k^{(i_{r-1})}$  to  $\mathcal{A}$ .
- **Queries to H:**  $\mathcal{B}$  merely responds to all queries from  $\mathcal{A}$  to  $H$  by using its oracle access to  $H$ .
- **Queries to Partial Evaluation:**  $\mathcal{B}$  receives as query input, some choice of key index specified by  $i^*$  and input  $x_j$  for  $i = 1, \dots, Q$ . In response  $\mathcal{B}$  does the following:
  - Forward  $x_j$  to its challenger. In response it implicitly receives P-Eval( $k^{(i_r)}, x_j$ ), but off by a factor of  $\Delta$  in the exponent. Call this  $h_{j,r}$ .
  - Compute:  $h_{j,i^*} = H(x_j)^{\sum_{i=1}^{r-1} \Lambda_{i_j}(i^*) \cdot k^{(i_j)}} \cdot (h_{j,r})^{\Lambda_{i_r}(i^*)}$  where  $\mathcal{B}$  uses its own access to hash oracle to get  $H(x_j)$ .
  - It returns  $h_{j,i^*}$  to  $\mathcal{A}$ .
- **Challenge Query:** On receiving the challenge input  $x^*$ ,  $\mathcal{B}$  does the following:
  - Ensure that it is a valid input, i.e., there is no partial evaluation queries on  $x^*$  at any unknown index point.
  - If not,  $\mathcal{B}$  forwards to its challenger  $x^*$ . In response it implicitly receives P-Eval( $k^{(i_r)}, x^*$ ), but off by a factor of  $\Delta$  in the exponent. Call this  $h^*$ .
  - It also uses its oracle access to  $H$  to receive  $h = H(x^*)$ .
  - It finally computes  $y = H(x_j)^{\Delta^2 \cdot \sum_{i=1}^{r-1} \Lambda_{i_j}(0) \cdot k^{(i_j)}} \cdot (h^*)^{\Delta^2 \cdot \Lambda_{i_r}(0)}$  and outputs  $y$  to  $\mathcal{A}$
- **Finish:** It forwards  $\mathcal{A}$ 's guess as its own guess.

**Analysis of the Reduction.** Note that for the case when  $b = 0$ ,  $\mathcal{A}$  expects to receive  $H(x^*)^{\Delta^3 \cdot k}$  where  $k$  is defined at the point 0. So, we get:

$$\begin{aligned} H(x^*)^{\Delta^3 \cdot k^{(0)}} &= H(x^*)^{\Delta^2 \cdot f'(0)} = H(x)^{\Delta^2 \sum_{j=1}^r \Lambda_{i_j}(0) \cdot k^{(i_j)}} \\ &= H(x)^{\Delta^2 \sum_{i=1}^{r-1} \Lambda_{i_j}(0) \cdot k^{(i_j)}} \cdot \left( H(x)^{k^{(i_r)}} \right)^{\Delta^2 \Lambda_{i_r}(0)} \end{aligned}$$

This shows that the returned value  $y$  is consistent when  $b = 0$ . Meanwhile, when  $b = 1$ ,  $h^*$  is a random element in the group and then  $y$  is a truly random value which means that  $\mathcal{B}$  has produced a valid random output for  $\mathcal{A}$ . Similarly, when  $b = 0$ , every response to partial evaluation is also done

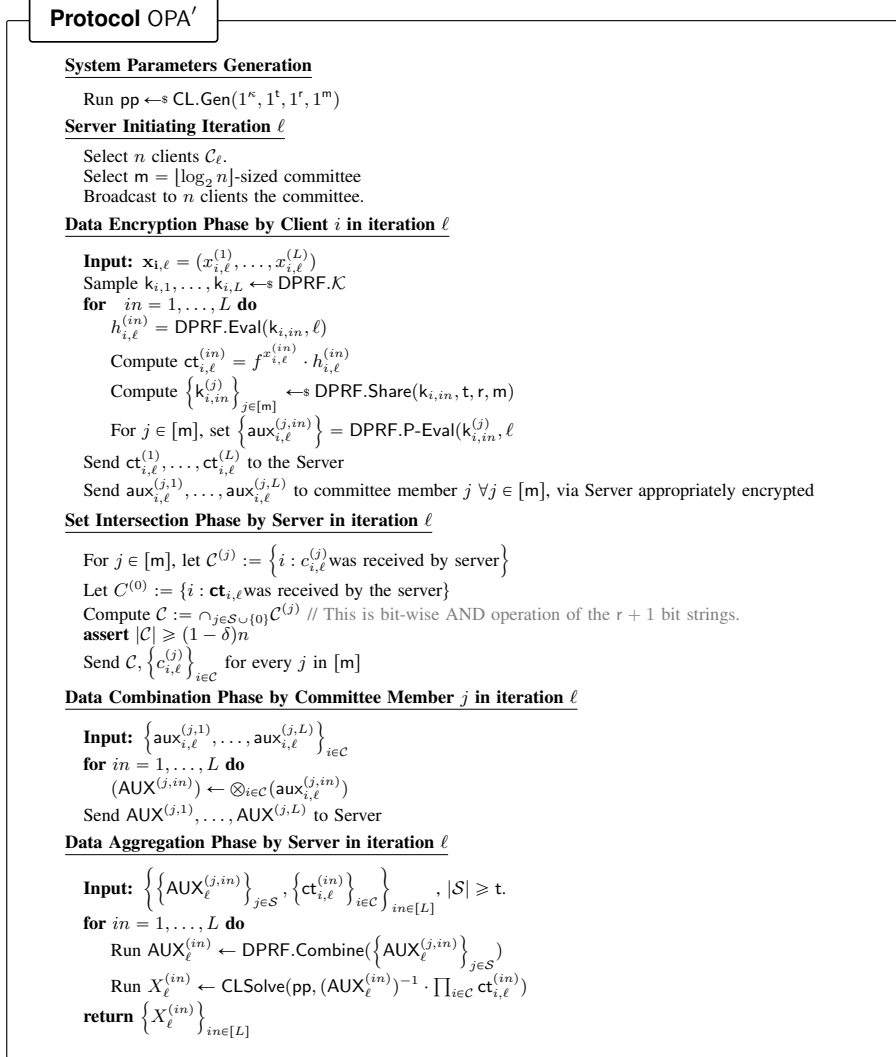


Figure 10: In this figure, we present the modified steps for OPA'. For simplicity, we present only the semi-honest construction. For malicious security, we augment similarly with a second mask.

consistently by correctness of the underlying secret sharing scheme. Meanwhile, when  $b = 1$ , we can rely on the statistical privacy preserving guarantee of the underlying secret sharing scheme to argue that the difference that the adversary can notice is statistically negligible. This concludes the proof where  $\mathcal{B}$  can only succeed with advantage  $\epsilon$ .  $\square$

## E.2 Construction of One-shot Private Aggregation without Leakage Simulation in CL Framework

**Construction 12.** We present the construction of One-shot Private Aggregation without Leakage Simulation, in the CL Framework in Figure 10.

## F Construction from LWR

### F.1 Distributed PRF from LWR

Let us revisit Construction 14. First, observe that the key space is from  $\mathbb{Z}_q^p$  which implies that the order of  $\mathcal{K}$  is known. Further, the computation occurs over a group whose structure and order is

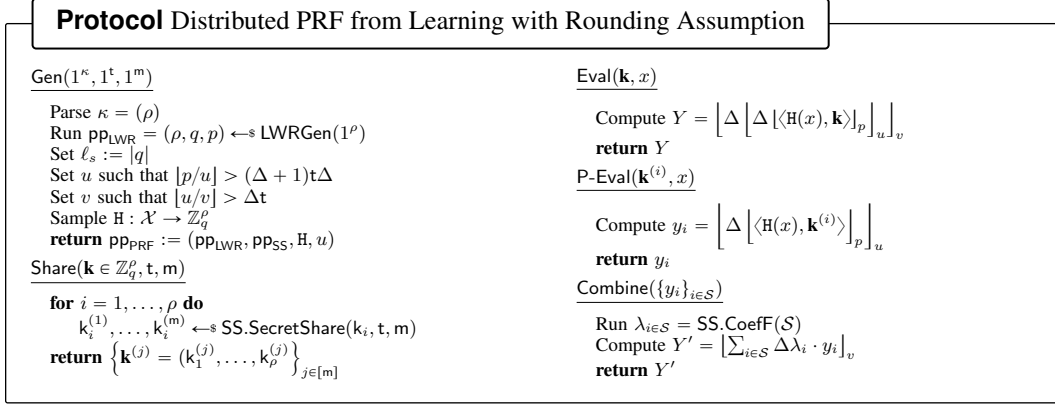


Figure 11: Construction of Distributed PRF based on the Secret Sharing scheme of Construction 6 where  $\text{SS} = (\text{SecretShare}, \text{Coeff})$  with  $\text{pp}_{\text{SS}}$  denoting the public parameters of the secret sharing scheme.

known. This is a departure from the construction based on the  $\text{HSM}_M$  assumption. Consequently, by assuming that both  $p$  and  $q$  are primes, one can avoid integer secret sharing but instead rely on traditional Shamir’s Secret Sharing over a field, which we defined earlier (see Construction 6).

We saw earlier that Construction 14 was only almost key homomorphic, i.e.:

$$F(\mathbf{k}_1 + \mathbf{k}_2, x) = F(\mathbf{k}_1, x) + F(\mathbf{k}_2, x) + e$$

where  $e \in \{0, 1\}$ . It also follows that:

$$T \cdot F(\mathbf{k}_1, x) = F(T \cdot \mathbf{k}_0, x) - e_T$$

where  $e_T \in \{0, \dots, T\}$ . This becomes a cause for concern as, in the threshold construction using the Shamir Secret Sharing over the field as shown in Construction 6, one often recombinces by multiplying with a Lagrange coefficient  $\lambda_{i_j}$ . Unfortunately, multiplying the result by  $\lambda_{i_j}$  implies that the error term  $e_{\lambda_{i_j}} \in \{0, \dots, i_j\}$ . The requirement is that this error term should not become “too large”. However, interpreting Lagrange coefficients as elements in  $\mathbb{Z}_p$  results in the error term failing to be low-norm leading to error propagation. To mitigate this, we use techniques quite similar to Construction 8 by essentially clearing the denominator by multiplying with  $\Delta := m!$ . This is a technique made popular by the work of Shoup [82] and later used in several other works including in the context of lattice-based cryptography by Agrawal *et al.* [4] and later to construct a distributed key homomorphic PRF from any almost key homomorphic PRF by Boneh *et al.* [19].<sup>6</sup> Then, the combine algorithm will simply multiply all partial evaluations with  $\Delta$  as well.

**Construction 13** (Distributed Almost Key Homomorphic PRF from LWR). A  $(t, m)$ -Distributed PRF is a tuple of PPT algorithms  $\text{DPRF} := (\text{Gen}, \text{Share}, \text{Eval}, \text{P-Eval}, \text{Combine})$  with the algorithms as defined in Figure 11.

**Issues with the Construction from Boneh *et al.* [19, §7.1.1].** As remarked earlier, their generic construction suffers from issues stemming from their security reduction. Specifically, their security reduction proceeds similarly to the proof of Theorem 10 and requires  $\mathcal{B}$  to answer honest evaluation queries for key indices for which it does not know the actual key share. Their explanation suggests that we again use the “clearing out the denominator” trick by multiplying with  $\Delta$ . However, the issue is that the resulting response will be of the form  $\Delta \cdot F(k_{i^*}, x)$  for  $i^*$ , unknown to  $\mathcal{B}$ . Consequently, one has to change the partial evaluation response to also include this offset to ensure the correctness of reduction. This would imply that the Combine algorithm will multiply with  $\Delta$  again, which would thus result in the actual Eval algorithm having an offset of  $\Delta^2$ . Furthermore, the partial evaluation algorithm should also have to round down to the elements in  $[0, u - 1]$  for the same reason that the Combine algorithm required this fix.

<sup>6</sup>However, their generic construction is incorrect, owing to issues in their security proof which is not entirely sketched out. We fix the issues in our construction.

**Correctness.**

$$\begin{aligned}
\left[ \Delta \left[ \langle \mathbf{H}(x), \mathbf{k} \rangle \right]_p \right]_u &= \left[ \Delta \left[ \sum_{z=1}^{\rho} H(x, z) \cdot \mathbf{k}_z \right]_p \right]_u \\
&= \left[ \Delta \left[ \sum_{z=1}^{\rho} H(x, z) \sum_{i_j \in \{i_1, \dots, i_t\}} \lambda_{i_j} s_z^{(i_j)} \right]_p \right]_u \\
&= \left[ \Delta \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \sum_{z=1}^{\rho} H(x, z) \cdot \lambda_{i_j} \cdot s_z^{(i_j)} \right]_p \right]_u \\
&= \left[ \Delta \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \langle \mathbf{H}(x), \lambda_{i_j} \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u \\
&= \left[ \Delta \left( e_t + \sum_{i_j \in \{i_1, \dots, i_t\}} \left[ \lambda_{i_j} \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right) \right]_u \\
&= \left[ \Delta \left( e_t + \sum_{i_j \in \{i_1, \dots, i_t\}} e_{\lambda_{i_j}} + \lambda_{i_j} \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right) \right]_u \\
&= \left[ \Delta \left( e_t + \sum_{i_j \in \{i_1, \dots, i_t\}} e_{\lambda_{i_j}} \right) + \sum_{i_j \in \{i_1, \dots, i_t\}} \Delta \cdot \lambda_{i_j} \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u \\
&= \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \Delta \cdot \lambda_{i_j} \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u
\end{aligned}$$

The last step follows provided the error term is small. Recall that  $e_t \in \{0, \dots, t\}$  and  $e_{\lambda_{i_j}} \in \{0, \dots, \lambda_{i_j}\}$ . Now observe that we multiply with  $\Delta$  and  $\lambda_{i_j}$  has a maximum value  $\Delta$ . Therefore,  $\Delta \cdot e_{\lambda_{i_j}} < \Delta^2$ . Therefore, the size of the error term is  $\leq t \cdot \Delta + t \cdot \Delta^2$ . Therefore, provided  $u$  is chosen such that  $\lfloor p/u \rfloor > (\Delta + 1) \cdot t \cdot \Delta$ , then the last step is correct. Now, we have:

$$\left[ \Delta \left[ \langle \mathbf{H}(x), \mathbf{k} \rangle \right]_p \right]_u = \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \Delta \cdot \lambda_{i_j} \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u$$

Therefore,

$$\begin{aligned}
\left[ \Delta \left[ \Delta \left[ \langle \mathbf{H}(x), \mathbf{k} \rangle \right]_p \right]_u \right]_v &= \left[ \Delta \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \Delta \cdot \lambda_{i_j} \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u \right]_v \\
&= \left[ \Delta \left( e_t + \sum_{i_j \in \{i_1, \dots, i_t\}} \lambda_{i_j} \left[ \Delta \cdot \left[ \langle \mathbf{H}(x), \mathbf{k}^{(i_j)} \rangle \right]_p \right]_u \right) \right]_v \\
&= \left[ \Delta \left( e_t + \sum_{i_j \in \{i_1, \dots, i_t\}} \lambda_{i_j} \text{P-Eval}(\mathbf{k}^{i_j}, x) \right) \right]_v \\
&= \left[ \sum_{i_j \in \{i_1, \dots, i_t\}} \lambda_{i_j} \cdot \Delta \cdot \text{P-Eval}(\mathbf{k}^{(i_j)}, x) \right]_v \\
&= \text{Combine} \left( \left\{ \text{P-Eval}(\mathbf{k}^{(i_j)}, x)_{i_j \in \{i_1, \dots, i_t\}} \right\} \right)
\end{aligned}$$

provided  $\lfloor u/v \rfloor > t\Delta$ .

**Pseudorandomness.** The proof of pseudorandomness follows the outline of the proof of Theorem 10 but with some important differences. First, we do not rely on integer secret sharing but rather plain secret sharing over the field. Therefore, the Lagrange coefficients correspond to  $\lambda_{i_j}$ . Or more formally, to respond to a partial evaluation query at point  $x_j$  with target key index  $i^*$ , the adversary  $\mathcal{B}$  does the following:

- Use its oracle to get partial evaluation on  $x_j$  at  $i_t$ , which we call as  $h_{j,t}$ .
- Then, use Lagrange coefficients but with suitably multiplying with  $\Delta$  to compute the correct distribution by rounding down to  $u$ . The choice of  $u$  guarantees that the response is correct.

For challenge query, it simply does two rounding down, first to  $u$  and then to  $v$ .

**Verification of Almost Key Homomorphism.** Let  $\mathbf{k}_1, \mathbf{k}_2$  be two keys that are shared. Now, let the key shares received by some party  $i_j$  be  $\mathbf{k}_1^{(i_j)}$  and  $\mathbf{k}_2^{(i_j)}$ . Then,

$$\begin{aligned} \text{P-Eval}(\mathbf{k}_1^{(i_j)}, x) + \text{P-Eval}(\mathbf{k}_2^{(i_j)}, x) &= \left\lfloor \Delta \left[ \langle H(x), \mathbf{k}_1^{(i_j)} \rangle \right]_p \right\rfloor_u + \left\lfloor \Delta \left[ \langle H(x), \mathbf{k}_2^{(i_j)} \rangle \right]_p \right\rfloor_u \\ &= \left\lfloor \Delta \left[ \langle H(x), \mathbf{k}_1^{(i_j)} \rangle \right]_p + \Delta \left[ \langle H(x), \mathbf{k}_2^{(i_j)} \rangle \right]_p \right\rfloor_u - e_1 \\ &= \left\lfloor \Delta \left[ \langle H(x), \mathbf{k}_1^{(i_j)} \rangle + \langle H(x), \mathbf{k}_2^{(i_j)} \rangle \right]_p \right\rfloor_u - 2e_1 \\ &= \text{P-Eval}(\mathbf{k}_1^{(i_j)} + \mathbf{k}_2^{(i_j)}, x) - 2e_1 \end{aligned}$$

It follows that for  $n$  such keys:

$$\sum_{i=1}^n \text{P-Eval}(\mathbf{k}_i^{(i_j)}, x) = \text{P-Eval}\left(\sum_{i=1}^n \mathbf{k}_i^{(i_j)}, x\right) - n \cdot e_1$$

This shows that the P-Eval is almost key-homomorphic. Consequently, one can verify that the whole Eval procedure is almost key homomorphic for the appropriate error function. To do this, recall that from correctness of our algorithm:

$$\text{Share}(\mathbf{k}, m, t) = \left\{ \mathbf{k}^{(i)} \right\}_{i \in [m]}, \text{Combine}\left(\left\{ \text{Eval}(\mathbf{k}^{(i_j)}, x) \right\}_{i \in \{i_1, \dots, i_t\}}\right) = \text{Eval}(\mathbf{k}, x)$$

In other words, for  $\mathbf{k}_1, \mathbf{k}_2 \in \mathcal{K}$ ,  $\text{Share}(\mathbf{k}_1, m, t) = \left\{ \mathbf{k}_1^{(i)} \right\}_{i \in [m]}$  and  $\text{Share}(\mathbf{k}_2, m, t) = \left\{ \mathbf{k}_2^{(i)} \right\}_{i \in [m]}$ , we will have for  $i \in [m]$ ,  $\text{Eval}(\mathbf{k}_1^{(i)}, x), \text{Eval}(\mathbf{k}_2^{(i)}, x) = \text{Eval}(\mathbf{k}_1^{(i)} + \mathbf{k}_2^{(i)}, x) - 2 \cdot e_1$ .

## F.2 One-shot Private Aggregation Construction based on LWR Assumption

We build OPA based on the LWR Assumption, building it based on the Key Homomorphic, Distributed PRF as presented in Construction 14. However, our construction is largely different from the template followed to build OPA from the HSM<sub>M</sub> assumption. This is primarily because of the growth in error when combining partial evaluations. Specifically, will get that  $\text{P-Eval}(\sum_{i=1}^n \mathbf{k}_i^{(j)}, x) = \sum_{i=1}^n \text{P-Eval}(\mathbf{k}_i^{(j)}, x) + e$  where  $e \in \{0, \dots, n-1\}$  where  $n$  is the number of clients participating for that label. This would require us to round down to a new value  $u'$  such that  $\lfloor u/u' \rfloor > n-1$ . Therefore, while we still employ the underlying functions of the distributed, key-homomorphic PRF based on LWR, we have to open up the generic reduction. For simplicity, we detail the construction for  $L = 1$ . Then, these are the differences:

- As done fore Construction 1, the input is encoded as  $x_i \cdot n + 1$ .
- Specifically, the client's share to the committee will only be the first level of the evaluation, i.e., rounded down to  $p$ .
- Then, committee member  $j$  will then add the shares up, multiply with the offset, and then round down to  $u$ . We will show that provided  $\lfloor p/u \rfloor > \Delta \cdot n$ , this is consistent with  $\text{P-Eval}(\sum_{i=1}^n \mathbf{k}_i^{(j)}, x)$ .

- The decoding algorithm is also similar to the one from Construction 1.

Recall that DPRF correctness requires that  $\lfloor p/u \rfloor > t \cdot \Delta + t \cdot \Delta^2$ , and so one just needs  $\lfloor p/u \rfloor > \max(t \cdot \Delta + t \cdot \Delta^2, n \cdot \Delta)$ . Then, one can rely on the correctness of DPRF as shown below to argue that when the server runs DPRF.Combine, the output is  $\text{Eval}(\sum_{i=1}^n \mathbf{k}_i, x)$ .

**Correctness.** We showed how the output of the server's invocation of DPRF.Combine is  $\text{Eval}(\sum_{i=1}^n \mathbf{k}_i, x)$ . Now, let us look at the remaining steps:

$$\begin{aligned}
X_\ell &= \sum_{i=1}^n \text{ct}_{i,\ell} - \text{AUX}_\ell = \sum_{i=1}^n (x_{i,\ell} * n + 1) + \text{Eval}(\mathbf{k}_i, \ell) - \text{Eval}(\sum_{i=1}^n \mathbf{k}_i, \ell) \bmod v \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n + \sum_{i=1}^n \text{Eval}(\mathbf{k}_i, \ell) - \text{Eval}(\sum_{i=1}^n \mathbf{k}_i, \ell) \bmod v \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n + \text{Eval}(\sum_{i=1}^n \mathbf{k}_i, \ell) - \text{Eval}(\sum_{i=1}^n \mathbf{k}_i, \ell) - e_{n-1} \bmod v \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n - e_{n-1} \bmod v \\
&= n \cdot \sum_{i=1}^n x_{i,\ell} + n - e_{n-1}
\end{aligned}$$

For the last step to hold, we need that

$$0 \leq n \cdot \sum_{i=1}^n x_{i,\ell} + n - e_{n-1} < v$$

$e_{n-1}$  the small value is 0 and the largest value is  $n - 1$  which requires that  $\sum_{i=1}^n x_{i,\ell} < (v - n)/n$ . Note that we already require  $\lfloor p/u \rfloor > n\Delta$ ,  $\lfloor u/v \rfloor > t\Delta \implies \lfloor p/v \rfloor > nt\Delta^2$ . In other words,  $\sum x_i < \frac{p}{n^2 t \Delta^2}$ . Finally,  $X'_\ell = n \cdot \sum_{i=1}^n x_{i,\ell} + n$  and that completes the remaining steps.

## G Constructions of Leakage Resilient Key-Homomorphic Pseudorandom Functions

### G.1 Construction from LWR Assumption

We also have the construction from Boneh *et al.* [19] of an *almost* Key Homomorphic PRF from LWR in the Random Oracle model which was later formally proved secure by Ernst and Koch [48] with  $\gamma = 1$ .

**Construction 14** (Key Homomorphic PRF from LWR). Let  $H : \mathcal{X} \rightarrow \mathbb{Z}_q^\rho$ . Then, define the efficiently computable function  $F_{\text{LWR}} : \mathbb{Z}_q^\rho \times \mathcal{X} \rightarrow \mathbb{Z}_p$  as  $\lfloor \langle H(x), \mathbf{k} \rangle \rfloor_p$ .  $F_{\text{LWR}}$  is an almost key homomorphic PRF with  $\gamma = 1$ .

**Construction 15** (Length Extended Key-Homomorphic from LWR). Let  $F_{\text{LWR}}$  be the function as defined in Construction 14. Then, consider  $F_{\text{LWR}}^L := (F_{\text{LWR}}(\mathbf{k}, (x, 1)), \dots, F_{\text{LWR}}(\mathbf{k}, (x, L)))$ .

It is easy to see that Construction 15 is a secure pseudorandom function. An adversary that can break the security of Construction 15 can be used to break the security of Construction 14.

**Theorem 11** (Leakage Resilience of Construction 15). *Let PRF be the PRF defined in Construction 15. Recall that  $\text{PRF}.\mathcal{K} = \mathbb{Z}_q^\rho$ . Then, it is leakage resilient in the following sense:*

$$\{\text{PRF}(\mathbf{k}, x), (\mathbf{k} + r) \bmod q : \mathbf{k}, r \leftarrow_s \mathcal{K}\} \approx_c \{Y, (\mathbf{k} + r) \bmod q : Y \leftarrow_s \mathcal{Y}, \mathbf{k}, r \leftarrow_s \mathcal{K}\}$$

*Proof.* The proof proceeds through a sequence of hybrids.

Hybrid<sub>0</sub>( $\kappa$ ): The left distribution is provided to the adversary. In other words, the adversary gets:

$$\{\text{PRF}(\mathbf{k}, x), (\mathbf{k} + r) \bmod q : \mathbf{k}, r \leftarrow_s \mathcal{K}\}$$

Hybrid<sub>1</sub>( $\kappa$ ): In this hybrid, we replace  $(k + r) \bmod q$  with a uniformly random value  $k' \leftarrow_s \mathcal{K}$ .

$$\{\text{PRF}(k, x), k' : k, k' \leftarrow_s \mathcal{K}\}$$

Note that  $(k + r) \bmod q$  and  $k'$  are identically distributed. Therefore, the Hybrid<sub>0</sub>, Hybrid<sub>1</sub> are identically distributed.

Hybrid<sub>2</sub>( $\kappa$ ): In this hybrid, we will replace the PRF computation with a random value from the range.

$$\{Y, k' : Y \leftarrow_s \mathcal{Y}, k' \leftarrow_s \mathcal{K}\}$$

Under the security of the PRF, we get that Hybrid<sub>1</sub>, Hybrid<sub>2</sub> are computationally indistinguishable.

Hybrid<sub>3</sub>( $\kappa$ ): We replace  $k'$  with  $(k + r) \bmod q$ .

$$\{Y, (k + r) \bmod q : Y \leftarrow_s \mathcal{Y}, k, r \leftarrow_s \mathcal{K}\}$$

As argued before Hybrid<sub>2</sub>, Hybrid<sub>3</sub> are identically distributed.

Note that Hybrid<sub>3</sub> is the right distribution from the theorem statement. This completes the proof.  $\square$

## G.2 Learning with Errors Assumption

**Construction 16** (Key Homomorphic PRF from LWE). Let  $H_{\mathbf{A}} : \mathcal{X} \rightarrow \mathbb{Z}_q^{L \times \lambda}$ . Then, define the efficiently computable function  $F_{\text{LWE}} : (\mathbb{Z}_q^\lambda \times \chi^L) \times \mathcal{X} \rightarrow \mathbb{Z}_q^L$  as  $F_{\text{LWE}}((\mathbf{k}, \mathbf{e}), x) := H_{\mathbf{A}}(x)\mathbf{k} + \mathbf{e}$ .  $F_{\text{LWE}}$  is an almost key homomorphic PRF.

**Remark 9.** We note that it has been shown that one can also sample  $\mathbf{x}$  from the same distribution as  $\mathbf{e}$ . This is known as the short-secret LWE assumption and was employed in the encryption scheme of Lyubashevsky *et al.* [66]. An immediate consequence of this assumption is that one can set  $\mathbf{A}$ 's dimension  $m$  to be much smaller than what is required for the LWE assumption.

Similar to the LWR construction, we need to show that the PRF based on LWE assumption is also leakage resilient. Unfortunately, a similar proof technique does not work because the construction also suffers from leakage on the error vector which are usually Gaussian secrets. Instead, we rely on the following assumption:

## H Deferred Proofs

**Theorem 4** (Leakage Resilience of Construction 4). *Let  $\text{PRG}_{\text{LWR}}$  be the PRG defined in Construction 4. Then, it is leakage resilient in the following sense:*

$$\{\text{PRG}_{\text{LWR}}(\text{sd}) \bmod p, \text{sd} + \text{sd}' \bmod q : \text{sd}, \text{sd}' \leftarrow_s \mathbb{Z}_q^{n_1}\} \approx_c \{\mathbf{y}, \text{sd} + \text{sd}' \bmod q : \mathbf{y} \leftarrow_s \mathbb{Z}_p^L, \text{sd}, \text{sd}' \leftarrow_s \mathbb{Z}_q^{n_1}\}$$

*Proof.* The proof proceeds through a sequence of hybrids.

Hybrid<sub>0</sub>( $\kappa$ ): The left distribution is provided to the adversary. In other words, the adversary gets:

$$\{\text{PRG}_{\text{LWR}}(\text{sd}) \bmod p, \text{sd} + \text{sd}' \bmod q : \text{sd}, \text{sd}' \leftarrow_s \mathbb{Z}_q^{n_1}\}$$

Hybrid<sub>1</sub>( $\kappa$ ): In this hybrid, we replace  $\text{sd} + \text{sd}' \bmod q$  with a uniformly random value  $\text{sd}'' \leftarrow_s \mathbb{Z}_q^{n_1}$ .

$$\{\text{PRG}_{\text{LWR}}(\text{sd}) \bmod p, \text{sd}'' : \text{sd}, \text{sd}'' \leftarrow_s \mathbb{Z}_q^{n_1}\}$$

Note that  $(\text{sd} + \text{sd}') \bmod q$  and  $\text{sd}''$  are identically distributed. Let us assume that there exists a leakage function oracle  $L$  that can be queried with an input  $\text{sd}$ , for which it either outputs  $\text{sd} + \text{sd}' \bmod q$  for a randomly sampled  $\text{sd}' \leftarrow_s \mathbb{Z}_q^{n_1}$  or outputs  $\mathbf{s}' \leftarrow_s \mathbb{Z}_q^{n_1}$ . If one could distinguish between hybrids Hybrid<sub>0</sub>, Hybrid<sub>1</sub>, then one could distinguish between the outputs of the leakage oracle, but the outputs are identically distributed. Therefore, the Hybrid<sub>0</sub>, Hybrid<sub>1</sub> are identically distributed. Therefore, the Hybrid<sub>0</sub>, Hybrid<sub>1</sub> are identically distributed.



Hybrid<sub>2</sub>( $\kappa$ ): In this hybrid, we will replace the PRG computation with a random value from the range.

$$\{\mathbf{y}, \text{sd}'' : \mathbf{y} \leftarrow \mathbb{Z}_p^L, \text{sd}'' \leftarrow \mathbb{Z}_q^{n_1}\}$$

Under the security of the PRG, we get that Hybrid<sub>1</sub>, Hybrid<sub>2</sub> are computationally indistinguishable.

Hybrid<sub>3</sub>( $\kappa$ ): We replace  $\text{sd}''$  with  $(\text{sd} + \text{sd}') \bmod q$ .

$$\{\mathbf{y}, (\text{sd} + \text{sd}') \bmod q : \mathbf{y} \leftarrow \mathbb{Z}_p^L, \text{sd}, \text{sd}' \leftarrow \mathbb{Z}_q^{n_1}\}$$

As argued before Hybrid<sub>2</sub>, Hybrid<sub>3</sub> are identically distributed.

Note that Hybrid<sub>3</sub> is the right distribution from the theorem statement. This completes the proof.  $\square$

**Theorem 1.** *Let  $\kappa_s$  and  $\kappa_c$  be the statistical and computational security parameters. Let  $L$  be the input dimension and  $n$  be the number of clients, that are  $\text{poly}(\kappa_c)$ . Let  $\delta$  be the dropout threshold and  $\eta$  be the corruption threshold such that  $\delta + \eta < 1$ . Then, there exists an efficient simulator Sim such that for all  $K \subset [n]$  such that  $|K| \leq \eta n$ , inputs  $X = \{\mathbf{x}_{i,\ell}\}_{i \in [n] \setminus K}$ , and for all adversaries  $\mathcal{A}$  against Construction 1, that controls the server and the set of corrupted clients  $K$ , which behave semi-honestly (resp. maliciously), the output of Sim is computationally indistinguishable from the joint view of the server and the corrupted clients. Sim is allowed to query  $F_{D,\delta}^L(X)$  (defined in Equation ??) once, per iteration.*

*Proof.* We will prove the theorem statement by defining a simulator Sim, through a sequence of hybrids such that the view of the adversary  $\mathcal{A}$  between any two subsequent hybrids are computationally indistinguishable. Let  $H = [n] \setminus K$ , which are the set of honest clients. Further, let  $\mathcal{C} = [n] \setminus D$  where  $D$  is the set of dropout clients.

It is important to note that the server is semi-honest. Therefore, it is expected to compute the set intersection of online clients  $\mathcal{C}$ , as expected. In other words, all committee members (and specifically the honest committee members) receive the same  $\mathcal{C}$ . This is an important contrast from active adversaries as a corrupt and active server could deviate from expected behavior and send different  $\mathcal{C}^{(j)}$ , for different committee members. This could help it glean some information about the honest clients.

We now sketch the proof below:

Hybrid<sub>0</sub>( $\kappa$ ): This is the real execution of the protocol where the adversary  $\mathcal{A}$  is interacting with the honest parties.

Hybrid<sub>1</sub>( $\kappa$ ): In this hybrid, we will rely on the security of the secret sharing scheme to do two things:

- On the one-hand, all corrupt committee members receive a random share from the honest client's seed. Note that there can be only a maximum of  $t$  corrupt committee members. By appropriately choosing  $m$ , conditioned on  $\eta$ , we can guarantee that this holds with overwhelming probability. Then, for an honest client  $i$ , these are the shares denoted by  $\left\{ \text{sd}_i^{(j)} \right\}_{j \in [m] \cap K}$  and are generated randomly.
- On the other hand, all the honest committee members receive a valid share of the honest client's seeds. However, each honest client  $i$  need to generate this from a polynomial  $p(X)$  that satisfies  $p(0) = \text{sd}_i$ , while also ensuring  $p(j) = \text{sd}_i^{(j)}$  for  $j \in [m] \cap K$ . Note that this is a polynomial time operation and is similar to the way packed secret sharing is done where multiple secrets are embedded at distinct points of the polynomial. See Construction 7 for how to build such a polynomial.

It is clear that by relying on the privacy of the secret sharing scheme, Hybrid<sub>0</sub>, Hybrid<sub>1</sub> are indistinguishable for the adversary.

Hybrid<sub>2</sub>( $\kappa$ ): In this hybrid, we change the definition of the last honest party's ciphertext. WLOG, let  $n$  be the last honest party in  $\mathcal{C}$ . Then, we will set  $\mathbf{ct}_n := \text{PRG.Expand}(\text{sd}_\ell) + \mathbf{x}_\tau - \sum_{i \in \mathcal{C} \cap H} \mathbf{ct}_i$ . Here,  $\mathbf{x}_\tau$  is the sum of the honest clients inputs. Note that we are still in the hybrid where Sim knows all the inputs.

It is clear that  $\text{Hybrid}_1, \text{Hybrid}_2$  are identically distributed, by the almost seed-homomorphism property of PRG, provided Sim chooses the inputs for the honest parties such that they sum up to the value in  $\mathbf{x}_\tau$ .

$\text{Hybrid}_3(\kappa)$ : Again, without loss of generality, let client 1 be the first honest client in  $\mathcal{C} \cap H$ . We will modify the way  $\mathbf{ct}_{1,\ell}$  is generated. We will set it as  $\mathbf{ct}_{1,\ell} := \mathbf{x}_{1,\ell} + u$  where  $u \leftarrow \text{PRG}(\mathcal{Y})$ .

$\text{Hybrid}_2, \text{Hybrid}_3$  are indistinguishable, provided Theorem 4 holds. In the reduction, we will implicitly set  $\text{sd}_1 + \text{sd}_n$  to be the leakage obtained from the Theorem 4's challenger. In this hybrid, Sim still continues to know all the inputs. If it was a real PRG output, then we can simulate  $\text{Hybrid}_2$ , while simulating  $\text{Hybrid}_3$  in the random case.

$\text{Hybrid}_4(\kappa)$ : In this hybrid, we will replace  $\mathbf{ct}_{1,\tau} := u'$  for  $u' \leftarrow \text{PRG}(\mathcal{Y})$ . It is clear that  $\text{Hybrid}_3, \text{Hybrid}_4$  are identically distributed.

At this point, observe that we have successfully replace the first honest client's ciphertext, with a uniformly random value that is independent of its input. Sim will continue to do this modification for every non-dropout honest client  $i \in \mathcal{C} \cap H$ . This leaves the clients with all-but-the-last honest clients' ciphertext to be independent of the input, while leaving the last honest client's ciphertext to be only a function of the sum of the inputs, which can be obtained by Sim's query to the functionality. Sim beings its interaction with the functionality. After all the honest clients have provided inputs to the trusted party  $\mathcal{T}$ , in Step ??, Sim does not instruct any corrupted client to abort but rather set their inputs to be 0. Then in Step ??, Sim does not abort the server. Therefore, in Step ??, Sim will learn the sum of the honest parties inputs. Denote it as  $\mathbf{x}_\ell$ , which is also the sum of the inputs of the honest, surviving clients. With this information, Sim uses the last hybrid to interact with the adversary  $\mathcal{A}$ , who's expecting the real world interaction. This will enable Sim to run  $\mathcal{A}$  internally. This is crucial to ensure that Sim can get the output of  $\mathcal{A}$ , in the real world, which might depend on its view (including the output) of the server. This view will, in turn, depend on the the honest clients' inputs. Since Sim sets the honest inputs, in this internal execution, to match the sum of inputs in the real world, we can guarantee that the output of  $\mathcal{A}$  in the internal simulation is indistinguishable from  $\mathcal{A}$ 's interaction in the real world by the aforementioned hybrid arguments.

**Proof of Security against Active Server.** Our constructions so far have relied on providing security against a semi-honest server. Note that, as shown in the proof of security for Theorems 1, we can use the functionality query to obtain the sum of all the honest non-drop out clients, as before.

In the semi-honest setting, it is easy to see that the set  $\mathcal{C}$ , with respect to which aggregation is performed, includes *all* the honest, non-dropout clients' inputs. Therefore, querying the functionality, Sim does indeed get the sum of all the honest clients' inputs that are also included in the summation in the real world. This is imperative to ensure that Sim, when internally invoking  $\mathcal{A}$ , can get the output of  $\mathcal{A}$  which should be indistinguishable from  $\mathcal{A}$ 's output in the real world. Specifically, this output of  $\mathcal{A}$  (in either the internal invocation or the actual execution) will depend on the view which consists of the output of the server. Therefore, if the output of the server in the real world does not include any of these honest clients' inputs, then the output produced by the internal invocation of  $\mathcal{A}$  can be different from that in the real world.

Let us look at the case when the server is corrupted. Such a server can mount an attack whereby the real-world execution of the protocol may exclude inputs of some of those honest parties but actually included in the output of the ideal functionality. The proof of malicious security is tricky in this setting. Specifically, a malicious server can drop clients after seeing the honest input. This is an issue in the simulation as the simulator has to generate the masked inputs for the honest clients without knowing which of them would be dropped later.

Prior works, beginning with that of Bonawitz *et al.* [17] have relied on using signatures to ensure that a malicious server does not compromise the privacy of an honest user. Fortunately, for OPA, we can rely on the one-shot nature of communication flow to secure messages and avoid using signatures.

As before, let  $K$  denote the corrupted clients. Then,  $H_{\text{Cli}} := [n] \setminus K$  is the set of honest clients,  $H_{\text{Com}} := [m] \setminus K$  is the set of honest committee members. Let  $K_{\text{Cli}} := [n] \cap K$  denote the corrupted clients and  $K_{\text{Com}} := [m] \cap K$  denotes the corrupted committee members.

Note that we do not rely on signatures. To achieve a protocol with signatures, there needs to be an additional round of communication between the committee members and the server. First, the

server forwards the message to the committee members. Then, the committee members responds with their set  $\mathcal{C}^{(j)}$ , which is also duly signed. Then, the server performs the intersection and contacts the committee member with this intersection along with signatures. A committee member then only aggregates if there are  $(m + t)/2$  valid signatures.

Our focus is to ensure that the committee members only speaks once. In other words, our construction currently has the server identify  $\mathcal{C}^{(j)}$ , for each  $j \in [m]$ , based on the information it has received from the client. Then, the server forwards the message to the committee member along with its computed intersection. This setting allows the server to selectively forward shares to committee member and also choose different sets for different committee members. We will show that if  $r > (m + t)/2$  where  $r$  is the reconstruction threshold in the committee and  $t$  is the corruption threshold, then the server doing so will receive meaningless information. Formally, we will show that there does not exist two sets of users  $\mathcal{C} \neq \mathcal{C}'$  such that the server can reconstruct the shares over these two sets.

Observe that the server controls  $t$  committee members. We require each honest committee member to participate once, per iteration. This is easily enforced as the share from the honest client encrypts, along with the share for the honest committee member, also the identity of the honest client and the iteration count. Therefore, a server cannot replay the same share, in another iteration. With this guarantee, a malicious server, in order to reconstruct the shares of two distinct sets  $\mathcal{C}, \mathcal{C}'$ , will require the cooperation of at least  $r - t$  honest users, while there are  $r - t$  honest users present. We will therefore need  $2(r - t) > m - t$ . Or,  $r > (m + t)/2$ . This ensures that the server can only effectively reconstruct with respect to a unique set  $\mathcal{C}$  and  $H^*$  is the set of honest users in this set. Note that the above inequality holds for  $r = 2 * m/3, t < m/3$ . Indeed, prior works such as Bonawitz *et al.* [17] and most recently LERNA [62] also tolerated only upto a  $m/3$  corruption threshold.

While we have shown that there is a unique set  $H^*$  of honest users,  $H^*$  is only revealed after all the honest clients have sent their inputs. Therefore, the simulator, during its internal execution of  $\mathcal{A}$ , needs to be able to generate the masked inputs for the honest users and it only knows the sum of *all* the honest clients that have not dropped out. This set may be distinct from  $H^*$ . Therefore, we need a way for the simulator to generate masked inputs, independent of the sum of the inputs, and then ensure that the correct sum is computed during reconstruction.

The simulator does the following:

- For every honest client that hasn't dropped out, i.e., for all  $i \in H_{\text{Cli}}$ , the simulator does the following:
  - Samples  $\text{dig}_{i,\ell} \leftarrow_{\$} \{0, 1\}^{\log q}$
  - Samples  $\text{sd}_{i,\ell} \leftarrow_{\$} \text{PRG}.\mathcal{K}$
  - It computes  $\text{mask}'_{i,\ell} := H(\text{dig}_{i,\ell})$
  - Computes  $\text{mask}_{i,\ell} = \text{PRG}.\text{Expand}(\text{sd}_{i,\ell})$
  - Sets  $\text{ct}_{i,\ell} := \text{mask}_{i,\ell} + \text{mask}'_{i,\ell}$
  - Like shown in proof of Theorem 1, the shares of  $\text{sd}_{i,\ell}^{(j)}, \text{dig}_{i,\ell}^{(j)}$  for corrupt committee members  $j \in K_{\text{Com}}$  are chosen at random. Meanwhile, the shares for the honest committee members are to be sampled in the second phase, with a specific purpose. However, the server still expects an encryption of shares from honest client to honest committee members. Therefore, it simply encrypts some random shares for the honest committee members too and sends it to the server.
  - It sends to  $\mathcal{A}$ ,  $\text{ct}_{i,\ell}$  and  $\text{sd}_{i,\ell}^{(j)}$  and  $\text{dig}_{i,\ell}^{(j)}$  for  $j \in [m]$ , which is encrypted appropriately.
- This concludes the client phase of the operation. Then, comes the interaction with the committee. Note that the simulator is also required to simulate the honest committee member  $j$ .
- The simulator, which has received  $\mathcal{C}^{(j)}$  for each honest committee member  $j$  does the check to make sure that there exists at least  $r - t$  such committee members with the same  $\mathcal{C}^{(j)}$ . We will call this client set as  $\mathcal{C}$ , while calling the set of these committee members to be  $C_{\text{good}}$ . Meanwhile, it records those committee members with a different  $\mathcal{C}^{(j)}$ . We will call this as some set  $C_{\text{bad}}$ . Looking ahead, for those honest committee members in  $C_{\text{bad}}$ , the

shares of the honest clients that are to be added up is going to be random values. Note that  $|C_{\text{bad}}| \leq m - r$ .

- The simulator now operates in two phases for honest committee member  $j \in H_{\text{Com}}$ . First is the share generation phase for honest clients  $i$ . It does the following:
  - If  $j \in C_{\text{bad}}$ , then for honest client  $i \in \mathcal{C}^{(j)} \cap H_{\text{Cli}}$ , set  $\text{sd}_{i,\ell}^{(j)}, \text{dig}_{i,\ell}^{(j)}$  to be random values.
  - Now, the simulator computes the shares for all honest clients  $i$  to  $j \in C_{\text{good}}$ . These are valid shares of  $\text{sd}_{i,\ell}, \text{dig}_{i,\ell}$  subject to the constraint that random values were fixed for those  $j \in C_{\text{bad}}$  where  $i \in \mathcal{C}^{(j)}$ , and for those  $j \in K_{\text{Com}}$ .
  - The honest committee member  $j$  receives from  $\mathcal{A}$ ,  $\text{sd}_{i,\ell}^{(j)}$  and  $\text{dig}_{i,\ell}^{(j)}$  for  $i \in \mathcal{C}^{(j)}$ . Note that the maximum number of prefixed values is  $m - r + t$ , and by our constraint  $r > m - r + t$  which guarantees that these prefixed values cannot uniquely determine a polynomial of degree  $r$ .
- The second phase, is the combination phase. It responds, as expected, subject to the set  $\mathcal{C}^{(j)}$  that it receives.
- Sim now queries the functionality. First, it provides  $[n] \setminus \mathcal{C}$  as the set of dropped out clients. Then, it sends for those corrupted clients  $K \cap \mathcal{C}$ , input as 0 to the functionality. In response, it gets  $\sum_{\mathcal{C} \cap H} \mathbf{x}_{i,\ell}$ . Call this  $\mathbf{x}_H$ .
- Sim now picks  $i^* \in H^*$ . It programs the random oracle by setting  $H(\text{dig}_{i^*,\ell}) = \mathbf{x}_H - \mathbf{mask}'_{i^*,\ell}$ .
- Sim continues to respond, on behalf of the honest committee member, as expected.
- Finally,  $\mathcal{A}$  (which controls the server) will make queries to random oracle and it answers as expected. Sim outputs whatever  $\mathcal{A}$  outputs at the end.

We will now need to show that the above simulation is indistinguishable from the real world execution that  $\mathcal{A}$  expects when it is internally run. The hybrids proceeds as follows:

Hybrid<sub>0</sub>( $\kappa$ ): This is the real world execution.

Hybrid<sub>1</sub>( $\kappa$ ): In this execution, we replace the shares sent by the honest client  $i$  to honest committee member  $j$ , which are encrypted under  $\text{pk}_j$  with a random value. Under the semantic security of this encryption scheme, we can guarantee that this is indistinguishable from the previous hybrid. Meanwhile, these honest committee members (which the simulator controls) will receive the shares directly from the simulator. The view of  $\mathcal{A}$ , in this hybrid, is indistinguishable from the real world execution, under the semantic security of the encryption scheme.

Hybrid<sub>2</sub>( $\kappa$ ): We will rely on the security of the secret sharing scheme to sample the shares for the honest clients, similar to Hybrid<sub>1</sub> of semi-honest security. For those  $j \in K_{\text{Com}}$ , the shares are randomly chosen. Furthermore, for those  $j \in C_{\text{bad}}$  also the shares are randomly chosen. Finally, for those  $j \in C_{\text{good}}$  it gets a valid share subject to those previously chosen random values. This is similar to Hybrid<sub>1</sub> in the proof of semi-honest security.

Hybrid<sub>3</sub>( $\kappa$ ): In this hybrid, for all those honest clients  $i$  that are not in  $\mathcal{C}$ , we will set  $\mathbf{ct}_{i,\ell} = \mathbf{mask}_{i,\ell} + \mathbf{mask}'_{i,\ell}$ , effectively setting the input to be 0. Observe that the view of  $\mathcal{A}$  remains unchanged as these honest clients inputs were never incorporated in the final sum anyway. Furthermore, if any of these  $i \in \mathcal{C}^{(j)}$  for  $j \in C_{\text{bad}}$ , the shares from these honest clients  $i$  to these  $j$  are completely random *and* independent of  $\mathbf{mask}_{i,\ell}$  and  $\mathbf{mask}'_{i,\ell}$ .

Hybrid<sub>4</sub>( $\kappa$ ): In this hybrid, we pick an honest surviving client  $i^* \in H^*$ . It sets the inputs for all  $i \neq i^* \in H^*$  to be 0. Then sets  $\mathbf{x}_{i^*,\ell}$  to be the sum of all the  $i \in H^*$ . Call this sum as  $\mathbf{x}_H$ . Observe that the values are still correlated and pseudorandom.

Hybrid<sub>5</sub>( $\kappa$ ): In this hybrid, we will program  $H(\text{dig}_{i^*,\ell}) = \mathbf{x}_H - \mathbf{mask}'_{i^*,\ell}$ , while setting  $\mathbf{ct}_{i^*,\ell} = \mathbf{mask}_{i^*,\ell} + \mathbf{mask}'_{i^*,\ell}$ . Note that because  $\text{dig}_{i^*,\ell}$  is chosen uniformly at random from  $q$  values where  $q$  is a large prime. The probability of collision is negligible. There is only negligible difference in the view of  $\mathcal{A}$ .

Hybrid<sub>6</sub>( $\kappa$ ): In this hybrid, we will set  $\mathbf{ct}_{i,\ell}$  for  $i \neq i^*$  to be some random term in the ciphertext space. Then, we will set  $\mathbf{ct}_{i^*,\ell} = \mathbb{H}(\text{dig}_{i^*,\ell}) - \sum_{i \neq i^*} \mathbf{ct}_{i,\ell}$ .

Note that under the leakage resilience property of the seed-homomorphic PRG, we can conclude that the two hybrids are computationally indistinguishable.

Hybrid<sub>7</sub>( $\kappa$ ): In this hybrid, we will replace  $\mathbf{ct}_{i,\ell} = \mathbf{mask}_{i,\ell} + \mathbf{mask}'_{i,\ell}$  for  $i \neq i^*$ .

Observe that this last hybrid is exactly what the simulator produces. This concludes the proof.  $\square$

**Theorem 2.** *Let  $\kappa_s$  and  $\kappa_c$  be the statistical and computational security parameters. Let  $L$  be the input dimension and  $n$  be the number of clients, that are  $\text{poly}(\kappa_c)$ . Let  $\delta$  be the dropout threshold and  $\eta$  be the corruption threshold such that  $\delta + \eta < 1$ . Then, there exists an efficient simulator Sim such that for all  $K \subset [n]$  such that  $|K| \leq \eta n$ , inputs  $X = \{\mathbf{x}_{i,\ell}\}_{i \in n \setminus K}$ , and for all adversaries  $A$  against Construction 2, that controls the server and the set of corrupted clients  $K$ , which behave semi-honestly (resp. maliciously), the output of Sim is computationally indistinguishable from the joint view of the server and the corrupted clients. Sim is allowed to query  $F_{D,\delta}(X)^\ell$  (defined in Equation ??) once, per iteration.*

*Proof.* The proof proceeds similar to that of Theorem 1, through a sequence of hybrids. However, there are a few differences. Construction 5 has the error vector  $\mathbf{e} \leftarrow \chi$ . However, we will replace  $\mathbf{e} = \mathbf{e}' + \mathbf{f}'$  where  $\mathbf{e}', \mathbf{f}' \leftarrow \chi'$ , the distribution present in Hint-LWE Assumption (see Definition 6). The hybrid descriptions are similar, so we only specify the differences:

- In Hybrid<sub>2</sub> we will set:

$$\mathbf{ct}_{n,\ell} = \mathbf{A} \cdot \text{sd}_\ell - \sum_{i \in (C \cap H) \setminus \{n\}} \mathbf{ct}_i + \mathbf{e}_\ell + \Delta \mathbf{x}_\ell$$

- We will argue that Hybrid<sub>2</sub>, Hybrid<sub>3</sub> are indistinguishable under Hint-LWE Assumption. We will sketch the reduction now.
  - Recall that, from the Hint-LWE Challenge, we get  $(\mathbf{A}, \mathbf{u}^*, \mathbf{s}^* := \mathbf{s} + \mathbf{r}, \mathbf{e}^* := \mathbf{e}' + \mathbf{f}')$ .
  - As done for the LWR construction, we will set the  $\mathbf{s}_1 + \mathbf{s}_n = \mathbf{s}^*$ , the leakage on key.
  - For generating  $\mathbf{ct}_{1,\ell}$  we will use  $\mathbf{u}^*$ , while also sampling a separate  $\mathbf{f}_1 \leftarrow \chi'$ . This gives us:  $\mathbf{ct}_{1,\ell} = \mathbf{u}^* + \mathbf{f}'_{n,\ell} + \Delta \cdot \mathbf{x}_{1,\ell}$
  - We will set  $\mathbf{ct}_{n,\ell} := \mathbf{A} \cdot \text{sd}_\ell - \sum_{i \in (C \cap H) \setminus \{n\}} \mathbf{ct}_i + \mathbf{e}^* + \sum_{i \in (C \cap H) \setminus \{1\}} (\mathbf{e}'_{i,\ell} + \mathbf{f}'_{i,\ell}) + \Delta \cdot \mathbf{x}_\ell$
  - When  $\mathbf{u}^*$  is the real sample, then  $\mathbf{ct}_{1,\ell}$  satisfies Hybrid<sub>2</sub>'s definition. Meanwhile, the  $\mathbf{ct}_{n,\ell}$  is also correctly simulated. Similarly, the case when it a random sample.

The proof of security against malicious server also follows that of the previous theorem.  $\square$

## I Extensions of One-shot Private Aggregation

### I.1 Security Against Committee Members

Hitherto, we have only considered the indistinguishability of information from the perspective of the server. However, one can consider the requirement to hold for even corrupt committee members. Specifically, should their entire committee collude (or at least  $t$  of them), then the client's input remains hidden. It is easy to observe that Figure 10 does not satisfy the stronger security definition. Specifically, if we had a single committee member, then the auxiliary information (which is available to the committee member) simply masks the input and therefore can be unmasked. Thus, to accommodate security against collusion of all committee members we modify OPA syntax and construction to include a key from the server to keep client privacy. Informally, we do the following:

- First, the server or the aggregator has a secret key as well. This is denoted by  $k_0$ .
- Second, for each label  $\ell$ , the server first publishes a “public key”, as a function of the following algorithm  $\text{aux}_{0,\ell} \leftarrow \text{PublicKeyGen}(k_0, \ell)$

- Third, the encrypt procedure takes into account this auxiliary information, i.e.,  $\left( \text{ct}_{i,\ell}, \left\{ \text{aux}_{i,\ell}^{(j)} \right\}_{j \in [m]} \right) \leftarrow_s \text{Enc}(\text{pp}, k_i, x_i, \text{aux}_{0,\ell}, t, m, \ell)$ . In other words, the server publishes the public key and then the client can begin encrypting to a label.
- Fourth, the decrypt procedure takes  $k_0$  as input too.

## NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The paper sets out to solve a critical problem in prior work on secure aggregation. In this work, we demonstrate how to reduce the synchronization by employing one additional party. We demonstrate experiments to show competitive performance over prior work. In addition, we also train machine learning models to justify that our protocol can be used for its intended purpose.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?



Answer: [Yes]

Justification: We have a conclusion paragraph that draws attention to some of the limitations while identifying how they can be remedied in future work.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The paper introduces all the necessary theoretical framework and assumptions for security of the construction. There's detailed proof deferred to the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The protocols are well detailed, including the parameter settings for our classifier.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

**5. Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: Unfortunately, there was no support for supplementary material upload. However, we are happy to furnish the anonymized code for interested reviewers.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

#### 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Clear tabulation provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

#### 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: Does not apply. We report mean of 30 iterations.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Clear details provided for reproducibility.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Code of Ethics compliant.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We focus on privacy-preserving federated learning which guarantees privacy of client-held data.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Does not apply for us.

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We specify the details for all the five datasets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

## 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets are being released.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

## 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our project does not involve any human subjects.

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Per the above question, we do not have any human subjects and therefore IRB approvals are not needed.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.