USENIX THE ADVANCED COMPUTING SYSTEMS ASSOCIATION

ACORN: Input Validation for Secure Aggregation

James Bell and Adrià Gascón, *Google LLC;* Tancrède Lepoint, *Amazon;*
Baiyu Li, Sarah Meiklejohn, and Mariana Raykova, *Google LLC;* Cathie Yun

https://www.usenix.org/conference/usenixsecurity23/presentation/bell

This paper is included in the Proceedings of the
32nd USENIX Security Symposium.

August 9–11, 2023 • Anaheim, CA, USA

978-1-939133-37-3

Open access to the Proceedings of the
32nd USENIX Security Symposium
is sponsored by USENIX.

ACORN: Input Validation for Secure Aggregation

James Bell[◇], Adrià Gascón[◇], Tancrède Lepoint[†,*] Baiyu Li[◇], Sarah Meiklejohn[◇], Mariana Raykova[◇], Cathie Yun[*]
[◇]Google [†]Amazon

Abstract

Secure aggregation enables a server to learn the sum of client-held vectors in a privacy-preserving way, and has been applied to distributed statistical analysis and machine learning. In this paper, we both introduce a more efficient secure aggregation protocol and extend secure aggregation by enabling *input validation*, in which the server can check that clients' inputs satisfy constraints such as L_0, L_2, and L_∞ bounds. This prevents malicious clients from gaining disproportionate influence on the aggregate statistics or machine learning model.

Our new secure aggregation protocol improves the computational efficiency of the state-of-the-art protocol of Bell et al. (CCS 2020) both asymptotically and concretely: we show via experimental evaluation that it results in 2-8X speedups in client computation in practical scenarios. Likewise, our extended protocol with input validation improves on prior work by more than 30X in terms of client communication (with comparable computation costs). Compared to the base protocols without input validation, the extended protocols incur only 0.1X additional communication, and can process binary indicator vectors of length 1M, or 16-bit dense vectors of length 250K, in under 80s of computation per client.

1 Introduction

Single-server secure aggregation, which enables a server to learn the sum of client-held vectors in a privacy-preserving way, can be used for the secure computation of distributed histograms or for averaging model updates in federated learning systems. As some concrete examples, it supports cryptographic protocols for recommendation systems [36] and time series analysis [44], and is also used in large-scale real-world deployments for predictive typing and selection [39, 27, 25].

In the single-server setting, a powerful server talks to a large number of (resource-constrained) clients with limited connectivity. Along with limited bandwidth, the latter constitutes a central challenge in production systems [27, 8]. The server might be corrupted and even collude with a subset of the clients. This threat model strikes a good balance between trust and efficiency for large-scale distributed computation, and is used by several existing aggregation protocols [6, 7, 42, 44, 41, 36, 45] and more general results [4, 13].

To achieve acceptable levels of accuracy and privacy, the *minimum* number of clients contributing to an aggregation ranges between 100 and 10,000 [27], depending on the application, with larger numbers resulting in better privacy or better trade-offs between privacy and accuracy. On the other hand, input vector sizes correspond to model sizes or histogram sketches, so lengths are easily in the range of hundred of thousands or millions. Therefore, a secure aggregation protocol suitable for practical applications must be *scalable* in terms of being able to tolerate large inputs and a large number of clients, and *dropout-robust* in terms of tolerating a relatively high fraction of clients that abort during the protocol execution. While client computation is a natural concern addressed in several previous works, bandwidth consumption (both in download and upload) is often a determining factor in practice [27]. Achieving both practical computational and communication efficiency is the main focus of this work.

Privacy-preserving input validation. Another aspect crucial for deploying secure aggregation in practice is *correctness in the face of corrupted clients*. This corresponds to enhancing protocols with defenses against *malicious* clients who seek to bias the aggregate data. We discuss such attacks in the context of federated learning in Section 1.2, along with the role that bounding the L_2 or L_∞ norm of client inputs (*norm-bounding*) plays as a defense. In statistics applications such as frequency counting, malicious clients should be prevented from having a disproportionate influence on the output, e.g. contributing a value other than 1 to a histogram bucket or contributing to a large number of buckets. This corresponds to a k-hotness check (i.e. a check that at most k entries are 1 and the rest are 0, which is an L_0 bound) on the input vectors of clients.

To implement these defenses, the server can perform *input validation* on the data sent by clients, relying on zero-

[*]Work done while employed at Google.

knowledge proofs to preserve privacy. Crucially, this must be done without requiring significant client computation. Input validation comes in different forms, ranging from detection of malicious client behavior (but still having it cause the protocol to abort) to both identifying misbehaving clients and removing their contributions from the final aggregate statistics.

1.1 Our Contributions

We introduce and evaluate three protocols in this paper: RLWE-SecAgg, ACORN-detect, and ACORN-robust. Their security and efficiency properties, as well as comparison with existing approaches, are presented in Table 1.

Our first contribution is RLWE-SecAgg, a new secure aggregation protocol based on lattice cryptography that improves the state of the art protocol due to Bell et al. [6] in terms of both concrete and asymptotic efficiency. More precisely, it retains the low communication of this protocol but achieves optimal computational costs.

Our second contribution includes protocol variants ACORN-detect and ACORN-robust with input validation based on zero-knowledge proofs that are practical in terms of both computation and communication. For example, in under 80 seconds, a client running ACORN-detect on a standard laptop can (1) show k-hotness of a binary input of length 1M, or (2) show that a dense vector of length 250K has its L_∞ norm bounded by 2^{16}. In terms of communication, the overhead of ACORN-detect over (non-validated) secure aggregation is roughly 0.1X. This is in contrast with previous works [33, 14] with double-digit factor overheads (see Table 1). To enable this, we provide a new zero-knowledge construction with logarithmic proof size for proving an L_∞ bound on a private vector (committed to in a constant-size commmitment).

In our evaluation of these protocols, we consider two scenarios: *analytics* and *learning*. The former corresponds to the secure computation of a size-ℓ histogram with inputs from $n = 10^4$ devices, where the protocol ensures that each clients contributes no more than once to a bounded number of buckets. The latter corresponds to a federated learning application, where the goal is to average length-ℓ model updates from $n = 500$ devices, while showing a bound on the norm of each client's input. We also provide benchmarks and overheads of the end-to-end performance using four real-world tasks and datasets, showing a bandwidth overhead of at most 1.05X and a manageable computational overhead.

RLWE-SecAgg: Secure aggregation from (R)LWE. As a starting point, we formulate a generalization of the Bell et al. secure aggregation protocol [6], which we refer to as PRG-SecAgg. Our generalized protocol recovers the original PRG-SecAgg construction if we instantiate it using a PRG-based encoding of the input, but we also present a new instantiation—RLWE-SecAgg—that uses a lattice-based encoding. This construction reduces client and server computation costs, both asymptotically and in terms of concrete efficiency. In more detail, for n clients and length-ℓ input vectors, PRG-SecAgg requires clients to do $O(\ell \log n)$ work and the server to do $O(n\ell \log n)$ work. In RLWE-SecAgg these costs improve to $O(\ell + \log n)$ and $O(n(\ell + \log n))$, respectively. This means that for $\log n \leq \ell$, which is the case in all known applications, our new protocol's client computation and communication costs are $O(\ell)$, which is optimal in that it matches the insecure baseline where clients just send their data.

ACORN: Practical private input validation. We propose secure aggregation protocols that support two types of input validation: ACORN-detect supports *detection* of malicious client behavior, while ACORN-robust provides *robustness* to misbehaving clients, as it has the ability to identify them and adaptively exclude their inputs from the final sum.

Our protocols extend our generalized SecAgg protocol and thus can be instantiated with both PRG-SecAgg and RLWE-SecAgg. One complicating factor is that, in SecAgg, clients encode their inputs using pairwise correlated keys. This design decision is justified by its communication efficiency [7], as the correlated randomness can be computed in an input-independent way using constant-sized seeds. Previous works like EIFFeL [14] and RoFL [33] use alternative underlying aggregation schemes (with quadratic and linear dependence in the number of clients, see Table 1) that result in much higher communication than ACORN.

A consequence of using correlated keys is that enforcing correctness becomes complex: clients must *individually* prove a norm bound on the input being encoded but *collectively* prove that the keys used in the encoding step of the protocol are correctly correlated. This latter property would be guaranteed if each client proved individually that it formed its key honestly, but this would be expensive. Instead, in ACORN-detect we use a distributed proof that does not require clients to interact. In ACORN-robust, we instead require neighboring clients to form *identical* commitments to their pairwise shared masks. As long as one of the pair is honest, the server can thus identify a mismatch and exclude the cheating client. This sort of client-aided verification is efficient, but does not work if two malicious clients are neighbors. We thus require clients to commit to shares of their correlated randomness before knowing who their neighbors will be, and also need a logarithm number of rounds to recursively perform this exclusion.

Succinct ZK proofs of bounded norm. Besides using an appropriate underlying SecAgg protocol, an important technique to achieve efficient communication is *ciphertext packing*: encoding several plaintext elements in a single ciphertext. While this keeps ciphertext expansion low even when working in a large group, it complicates the client's proof of correct encoding, as it needs to show an L_∞ bound for correctness of the (linear) packing function. For this we rely on Bulletproofs [10, 11], a discrete log-based zero-knowledge proof

			Clear	Bell et al. [6]	RLWE-SecAgg (Sec. 3)	RoFL [33]	EIFFeL [14]	ACORN-detect (Sec. 4)	ACORN-robust (Sec. 4)
Client communication			ℓ	$\ell+\log(n)$	$\ell+\log(n)$	$\ell+n$	ℓn^2	$\ell+\log(n)$	$\ell+\log^2(n)$
n	ℓ	$\gamma=\delta$							
10^2	10^4	0.05	31KB	45KB	45KB	>2MB*†	94MB	47KB	365KB
10^3	2^{18}	0.05	705KB	1030KB	1111KB	>51MB†	240GB*	1032KB	1350KB
10^3	2^{18}	0.33	705KB	1062KB	1144KB	>51MB†	N/A	1064KB	14MB
Client computation			ℓ	$\ell\log(n)$	$\ell+\log(n)$	ℓn	ℓn^2	$\ell\log(n)$	$\ell\log(n)+\log^2(n)$
Secure against									
server & up to γn clients			○	●	●	◐	●	●	◐
up to n clients			✓	✓	✓	✓	×	✓	✓
Input validation			✓	×	×	✓	✓	✓	✓
Robustness									
δn dropouts			✓	✓	✓	✓	✓	✓	✓
δn invalid inputs			✓	×	×	×	✓	×	✓

Table 1: The concrete communication, security properties, and communication and computation asymptotics of various secure aggregation algorithms. The concrete rows are all based on proving an L_∞ bound of $2^{32}/n$, with security against a semi-honest server and malicious clients. The bottom rows show which of the approaches protect client inputs, validate client inputs, function in the presence of (limited) dropouts, and function in the presence of (limited) malicious clients. *Values extrapolated via asymptotics from others in the column. † RoFL doesn't provide details for how to agree on a sharing of zero, so these communication statements don't include that (which is likely only a small fraction of the total cost). By ◐ and ● we denote that the protocol provides semi-honest and malicious security, respectively.

system with logarithmic proof size.

Using Bulletproofs to prove a bound on each entry of the input vector (as in RoFL [33]) results in linear communication; for a vector of length 2^{20}, for example, the commitments alone would require over 33 MB. The same is true of representing the desired bound as an arithmetic circuit or rank-1 constraint system (R1CS) and then proving it using a state-of-the-art proof system: this compilation process requires one commitment per input (meaning, in our case, one commitment for every entry in the vector), which again results in linear communication. Moreover, proving things entry-wise is not compatible with ciphertext packing.

Instead, we show how the recent techniques by Gentry et al. [19] for approximate proofs of L_∞ bounds via random projections, combined with known tricks for range proofs [22] and other optimizations, allow us to reduce our correct encoding proof to a single linear constraint that can be proved using Bulletproofs. Our approach thus allows us to prove *exact* L_∞ bounds, and furthermore to commit to an arbitrary-length vector using only 256 bits. Moreover, verification of multiple such proofs can be *batched*, which is crucial for our protocol to scale to large cohort sizes.

1.2 Secure Aggregation for FL

Federated learning (FL) is a distributed setting [27] where many clients collaboratively learn a model under the coordination of a central server. In each round of FL, (1) the server broadcasts the current model to all clients involved in that round. Then, in a (2) local training step, clients update the model using their local dataset of examples. Finally, the clients engage in (3) an aggregation step where all locally trained models are pre-processed and aggregated for the server to obtain an updated model. These steps (1-3) are iterated for a number of training rounds.

The role of secure aggregation, and thus of our work, in FL is in step (3) above. A secure summation protocol significantly reduces the leakage to the server with respect to a system where model updates are made available in the clear. Next we discuss poisoning attacks in FL, and how input validation – and in particular norm-bounding – helps to mitigate them.

Poisoning attacks in FL. Attacks can be divided into *untargeted* attacks, in which the goal is to generally degrade the quality of the model, and *targeted* attacks. In targeted attacks (also known as *backdoor* attacks), the goal of the attacker is to induce a given behavior in a particular subtask; e.g., for classification, have examples in a given class (cars) be misclassified (as, for example, birds) while retaining accuracy for the rest of the examples.

Several recent studies have empirically evaluated the effectiveness of imposing bounds on the L_2 norms of clients' model updates as a defense against these types of attacks. Intuitively, the norm bound limits the influence of malicious clients when trying to derail learning. Sun et al. [47] focus on model replacement attacks, where a malicious client provides a scaled malicious model update to effectively replace the current model by a backdoored one. They run a comprehensive study on the EMNIST dataset [12]—a real-life, user-partitioned, and non-IID dataset—and conclude that norm-bounding is an effective defense against backdoor attacks here.

Chowdhury et al. [14] empirically evaluate 7 attacks in the literature (5 untargeted, 2 targeted) on image classification tasks. The results show how norm-bounding helps as a defense, with gaps of more than 20% in accuracy being recovered (for targeted attacks). For the investigated backdoor attacks, the gap in accuracy between the main and backdoor tasks drops from roughly 10% to over 80% when norm-bounding is applied.

Shejwalkar et al. [43] identify parameters where defenses become (in)effective with respect to both existing attacks and new ones they propose. The emphasis of their work is on untargeted attacks operating within real-world deployments. Regarding norm-bounding specifically, they conclude that it "can effectively protect cross-device FL in practice" and more concretely that the "Norm-bounding Aggregation Rule (AGR) is enough to protect production FL against untargeted poisoning, questioning the need for the more sophisticated (and costlier) AGRs".

Finally, Lycklama et al. [33] characterize the classes of targeted attacks that norm-bounding can defend against and provide extensive empirical evaluation and an open-source experimentation framework. The authors conclude that while norm-bounding significantly decreases the available surface for adversarial attacks, it is *not a silver bullet*. In particular, continuous attacks on tail targets [48] remain effective even under norm bounds. Subsequent works have developed defenses against attacks on the tails based on sparsification combined with norm-bounding [49].

Server attacks in FL. The common model for secure aggregation in FL assumes a server that honestly runs step (1) of each FL round, i.e. that broadcasts the model resulting from the previous round to all clients. As shown by Pasquini et al. [37], a malicious server can instead send carefully crafted models to specific clients in order to extract their input. We prove our protocols secure assuming this attack does not take place, but the defenses proposed by Pasquini et al., and in particular model hashing, can be directly applied to our protocols. More concretely, clients can share with each other a hash of the model that they received from the server and verify that it matches their own before proceeding with the protocol.

1.3 Other Related Work

There are several well-known works on verifiable secure aggregation in the two-server or multi-server models [9, 15, 2], but we focus our discussion on the single-server trust model.

Stevens et al. perform differentially private secure aggregation (without input validation) using an LWE-based protocol [46]. This work is similar to our first contribution, RLWE-SecAgg. However, they overlook a subtlety in the security of their scheme, claiming that "[a secure aggregation of keys] reveals nothing about their individual [key] values." This is untrue, because the output itself conveys information even if computed securely. Our security proof addresses this issue.

Lycklama et al. [33] and Chowdhury et al. [14] introduce secure aggregation protocols with input validation called RoFL and EIFFeL, respectively. RoFL requires each client to send commitments to each vector entry to the server. For vectors of length 262,000, they report a 48x increase in required communication (to 51MB) when proving an L_∞ bound, compared to sending the vector in the clear. EIFFeL shares the computation amongst the clients, using them to replace the servers in Prio [15]. This allows them to deal with a constant fraction of malicious clients and dropouts. However, their communication scales quadratically in the number of clients and linearly in the vector length. Thus even for a vector of length 10^4 and 100 clients they report 94MB of communication. This is about three orders of magnitude greater than the cost in the clear. EIFFeL and RoFL suffer from the difficulties of balancing input validation with communication costs, which is a major focus of our work. In Table 1 we offer a detailed comparison in terms of both asymptotic and concrete efficiency.

Ghodsi et. al. [20] propose zPROBE, a secure aggregation protocol that checks that each entry of a client's masked input is constructed honestly from an input of bounded size. They do this by putting the circuit for a pseudorandom generator in a generic proof framework, but as this is very costly they check only a random subset of entries in their experiments. This is enough to prevent submissions where an appreciable fraction of the entries lie outside the desired bound, but this is not sufficient in FL as a model can be corrupted by changing only one or a few entries by a large amount.

Karakoç et. al. [28] also provide secure aggregation with range validation using an oblivious programmable pseudorandom function. They describe this work as a proof of concept and provide experiments only for vectors of length 16 due to the currently prohibitive computational costs.

2 Preliminaries

We denote by $x \leftarrow \chi$ sampling according to a distribution χ. If X is a finite set, we denote by $x \leftarrow X$ uniform sampling from X. By $\approx_{\sigma,\lambda}$ we denote indistinguishability with computational parameter λ and statistical parameter σ; i.e., $D_n \approx_{\sigma,\lambda} F_n$ if there exists another distribution E_n such that D_n is statistically close to E_n ($|\Pr_{x\leftarrow D_n}[A(x)=1]-\Pr_{x\leftarrow E_n}[A(x)=1]| < \sigma(n)$ for an unbounded adversary A) and E_n is computationally indistinguishable from F_n ($|\Pr_{x\leftarrow E_n}[A(x)=1]-\Pr_{x\leftarrow F_n}[A(x)=1]| < \lambda(n)$ for a PPT adversary A). We use the standard simulation-based formalism [21, 32] in our security proofs. We assume key agreement, authenticated encryption, and signature primitives, which we denote as KA, $\mathsf{E}_{\mathsf{auth}}$, and Sig.

2.1 Setting and Threat Model

We consider n clients $1,\ldots,n$, each holding a private vector $\mathbf{x_i} \in \mathbb{Z}_t^\ell$, and a server with communication channels established with all clients. The goal is for the server to obtain the sum of all client vectors ($\sum_i^n \mathbf{x_i}$), with robustness to a certain fraction of client dropouts. To make this concrete, the functionality is parameterized by a maximum fraction of dropouts $\delta \in [0,1]$, defined as follows:

$$f(\mathbf{x_1},\ldots,\mathbf{x_n}) = \begin{cases} \sum_{i\in[n]\setminus D} \mathbf{x_i} & \text{if } |D| \le \delta n \\ \perp & \text{otherwise} \end{cases} \quad (1)$$

where $D \subseteq [n]$ is the set of clients that dropped out during the protocol execution and the sum happens in $\mathbb{Z}_t^\ell$. We aim to withstand an adversary consisting of a coalition of γn clients, for $\gamma \in [0,1]$, and possibly also colluding with the server. As in previous works [7, 6], we assume that corrupt clients are fully malicious. For RLWE-SecAgg and ACORN-detect, we also assume the server is fully malicious,[1] but for ACORN-robust we prove security only in the case of a semi-honest server.

2.2 Lattices and Polynomial Rings

A lattice is a discrete subgroup $\Lambda \subset \mathbb{R}^N$, and it can be represented as the set of all integer combinations of a basis $\mathbf{B}$ such that $\Lambda = \mathbf{B}\mathbb{Z}^N$. We use the cyclotomic ring $R = \mathbb{Z}[X]/(X^N+1)$ for a power-of-two N, and write $R_q = \mathbb{Z}[X]/(q, X^N+1)$ for the residual ring of R modulo q. The *coefficient embedding* of a polynomial $a = \sum_{i=0}^{N-1} a_i X^i \in R$ is the vector $(a_0, a_1, \ldots, a_{N-1})$, and we define the L_∞ norm of a as $\|a\|_\infty = \|(a_0,a_1,\ldots,a_{N-1})\|_\infty = \max_i |a_i|$. As a convention, we use bold $\mathbf{a}$ to denote the coefficient embedding of a polynomial $a \in R$. We also define the negacyclic matrix representation of $a \in R$ as

$$\varphi(a) = \begin{pmatrix} a_0 & -a_{N-1} & \cdots & -a_1 \\ a_1 & a_0 & \cdots & -a_2 \\ \vdots & \vdots & \ddots & \vdots \\ a_{N-1} & a_{N-2} & \cdots & a_0 \end{pmatrix} \in \mathbb{Z}^{N\times N}.$$

We can naturally extend the map φ to vectors $\mathbf{a}$ over R such that $\varphi(\mathbf{a})$ is a matrix produced by vertically concatenating $\varphi(a_i)$ for all $a_i \in \mathbf{a}$. Without loss of generality, since the product of two polynomials $a,b \in R$ has the coefficient embedding $\mathbf{A}\cdot\mathbf{b}$ for $\mathbf{A} = \varphi(a)$, we represent $a \cdot b$ as a matrix-vector product $\mathbf{A}\cdot\mathbf{b}$. When $q = 1 \pmod{2N}$, this computation can be done more efficiently via Number Theoretic Transformation (NTT) than a naïve matrix-vector multiplication.

2.3 Ring LWE and Encryption

The *ring learning-with-errors* (RLWE) assumption [35], parameterized by a ring R of degree N over $\mathbb{Z}$, an integer modulus $q > 0$ defining a quotient ring $R_q = R/qR$, distributions χ_s, χ_e over R, and an integer m, states that for a secret $s \leftarrow \chi_s$, given $m = \mathsf{poly}(N)$ many independent samples from the distribution $A^{\mathsf{RLWE}}_{N,q,\chi_e}(s) = \{(a, as+e) \in R_q^2 \mid a \leftarrow R_q, e \leftarrow \chi_e\}$, it is computationally hard to distinguish them from m uniformly random samples over R_q. As we sometimes work with coefficient embedding of polynomials, we can rewrite an RLWE sample as $(\mathbf{A}, \mathbf{As}+\mathbf{e})$ for $\mathbf{s} \leftarrow \chi_s \subseteq \mathbb{Z}_q^N$, a matrix $\mathbf{A} = \varphi(a) \in \mathbb{Z}_q^{N\times N}$ for $a \leftarrow R_q$, and an error vector $\mathbf{e} \leftarrow \chi_e \subseteq \mathbb{Z}^N$. In our protocol we treat $\mathbf{A}$ as a public parameter. To encrypt a plaintext message $\mathbf{x} \in \mathbb{Z}_T^N$, we sample $\mathbf{s} \leftarrow \chi_s$ and compute

$$\mathsf{Enc}(\mathbf{s},\mathbf{x}) = (\mathbf{A}, \mathbf{As} + T\cdot\mathbf{e} + \mathbf{x} \bmod q), \quad (2)$$

and decrypt using $\mathsf{Dec}(\mathbf{s},(\mathbf{A},\mathbf{y})) = (\mathbf{y} - \mathbf{As}) \bmod T$. For longer messages $\mathbf{x} \in \mathbb{Z}_T^\ell$ such that $\ell > N$, we can naturally extend this encryption scheme by using multiple $\mathbf{A}_1,\ldots,\mathbf{A}_{\ell/N}$ where $\mathbf{A}_i = \varphi(a_i)$ for $a_i \leftarrow R_q$; equivalently, we sample $\mathbf{s} \leftarrow \chi_s$ and compute an extended ciphertext $\mathsf{Enc}(\mathbf{s}, \mathbf{x} \in \mathbb{Z}_T^\ell) = (\mathbf{A}, \mathbf{As} + T\cdot\mathbf{e} + \mathbf{x} \bmod q)$, where $\mathbf{A} = \varphi(\mathbf{a})$ is the vertical concatenation of $\mathbf{A}_1,\ldots,\mathbf{A}_{\ell/N}$, and $\mathbf{e} \leftarrow \chi_e^{\ell/N}$. When the plaintext modulus T is coprime to q, the distribution on $(\mathbf{A}, \mathbf{As}+T\cdot\mathbf{e})$ with $\mathbf{A} = \varphi(\mathbf{a})$ for $\mathbf{a} \leftarrow R_q^{\ell/N}$ and $\mathbf{e} \leftarrow \chi_e^{\ell/N}$ is indistinguishable from uniform under the RLWE assumption. For simplicity, we omit the public parameter $\mathbf{A}$ from ciphertexts.

Two important properties that we use in our protocol are key homomorphism and message homomorphism, i.e. (informally) $\mathsf{Enc}(\mathbf{s}_1,\mathbf{x}_1) + \mathsf{Enc}(\mathbf{s}_2,\mathbf{x}_2) = \mathsf{Enc}(\mathbf{s}_1+\mathbf{s}_2, \mathbf{x}_1+\mathbf{x}_2)$.

2.4 Commitment and Zero-Knowledge Proofs

Let G be a cyclic group of order q. The vector Pedersen commitment of $\mathbf{v} \in \mathbb{Z}_q^n$ using generators $g_0, g_1, \ldots, g_n \leftarrow \mathsf{G}$ and randomness $r \in \mathbb{Z}_q$ is $C = g_0^r \prod_{i=1}^n g_i^{v_i} \in \mathsf{G}$, and we denote the commitment algorithm using the notation $\mathsf{com} = \mathsf{Commit}(\mathbf{v}; r)$. It is perfectly hiding and computationally binding under the discrete logarithm assumption. We build our zero-knowledge proofs using Bulletproofs [11], which we describe in more detail in Section 5. Bulletproofs satisfies zero knowledge, meaning a simulator without knowledge of a witness can produce proofs that are indistinguishable from honest ones, and knowledge soundness, meaning it is possible to extract a valid witness from any proof that verifies. We

[1] We technically assume that the server behaves semi-honestly in a key distribution phase but otherwise maliciously, which is implied by assuming a fully malicious server with a PKI.

describe our proof systems as interactive, but make them non-interactive via the Fiat-Shamir heuristic [17], which means we operate in the random oracle model.

3 Generalized Secure Aggregation

In this section we present a generalized version of the secure aggregation protocol of Bell et al. [6] (SecAgg), where we abstract the method used by each party to hide its input as an encoding scheme $(\mathsf{Encode}, \mathsf{Decode})$. This encoding scheme should be additively homomorphic in both keys and values, meaning $\sum_i \mathsf{Encode}(\mathbf{sk}_i, \mathbf{x}_i) = \mathsf{Encode}(\sum_i \mathbf{sk}_i, \sum_i \mathbf{x}_i)$, and thus $\mathsf{Decode}(\sum_i \mathbf{sk}_i, \sum_i \mathsf{Encode}(\mathbf{sk}_i, \mathbf{x}_i)) = \sum_i \mathbf{x}_i$.

A simplified version of SecAgg is in Figure 1, and a full formal specification is in Algorithm 6 (in the appendix). This also contains the additional interactions needed to support input validation, which we ignore for now but describe in the next section. We then provide two examples of how this encoding can be instantiated: the first allows us to recover the original Bell et al. PRG-SecAgg construction, while the second provides a more efficient construction, RLWE-SecAgg (as we confirm experimentally in Section 6).

Commitments, distributed graph generation, and seed sharing. At its heart, SecAgg consists of two interactions between a set of clients and a server. In the first, `ShareSeeds`, each client i takes as input some randomness and learns four pieces of information: (1) a *pairwise seed* $\mathbf{seed}_{i,j}$ that it shares with each neighbor j in some defined communication graph, (2) a *self* seed $\mathbf{seed}_i$, and sets of shares (3) $\mathsf{shares}_{i,\mathcal{D}}$, corresponding to shares of these different seeds that this client should provide to the server for neighbors that *drop out* and (4) $\mathsf{shares}_{i,\mathcal{S}}$ that the client should provide to the server for neighbors that do not. In a slight abuse of notation, we write this as $(\varepsilon, \{\{\mathbf{seed}_{i,j}\}_{j \in \mathsf{N}(i)}, \mathbf{seed}_i, \mathsf{shares}_{i,\mathcal{D}}, \mathsf{shares}_{i,\mathcal{S}}\}_i) \leftarrow$ $\mathtt{ShareSeeds}(\varepsilon, \{\mathsf{rand}_i\}_i)$, where the first input (and output) denote the input (and output) of the server, which in this protocol should learn nothing, and the remaining sets denote the individual inputs (and outputs) of the clients. Some of the main challenges of this first protocol lie in ensuring that honest clients do not have too many malicious neighbors in the communication graph, and that $\mathbf{seed}_{i,j} = \mathbf{seed}_{j,i}$ for all pairs of honest neighbors i and j. This latter property is crucial in ensuring that the derived *masks* cancel when masked inputs are aggregated by the server.

Masking. Using the information learned in this first interactive protocol, client i can then mask its input $\mathbf{x}_i$ using an encoding key computed as

$$\mathbf{sk}_i = \mathbf{s_i} + \sum_{j \in A_i, j<i} \mathbf{s_{ij}} - \sum_{j \in A_i, i<j} \mathbf{s_{ij}}, \tag{3}$$

where $\mathbf{s_{ij}} = \mathsf{F}(\mathbf{seed}_{ij})$, $\mathbf{s_i} = \mathsf{F}(\mathbf{seed}_i)$ for a length-expanding function F, and A_i are the neighbors that i believes to be survivors at this step in the protocol. It then encodes its input as $\mathbf{y}_i = \mathsf{Encode}(\mathbf{sk}_i, \mathbf{G}\mathbf{x}_i)$, where $\mathbf{G}$ is a matrix that allows us to *pack* multiple entries of $\mathbf{x}_i$ into a single plaintext slot; we discuss this in more detail below.

Dropout agreement and unmasking. If all clients are honest and do not drop out, then all their pairwise masks cancel, meaning $\sum_i \mathbf{sk}_i = \sum_i \mathbf{s_i}$. In this case, each client could just provide their individual self mask $\mathbf{s_i}$ to the server at the end of the protocol, who could then take advantage of the dual-homomorphic property of the encoding scheme to compute $\sum_i \mathbf{x}_i = \mathbf{G}^{-1}(\mathsf{Decode}(\sum \mathbf{s_i}, \sum \mathbf{y}_i))$. To account for clients who drop out, however, the server must have a way to recover their pairwise masks in order to cancel them out itself from the keys of *surviving* clients (e.g., if a surviving client i used $\mathbf{s_{ij}}$ for a dropped out client j in forming $\mathbf{sk}_i$, there is no corresponding $\mathbf{sk}_j$ containing $-\mathbf{s_{ij}}$ to cancel out the masks "naturally").

The second interactive protocol that SecAgg provides is thus $\sum \mathbf{s}_i, \{\mathcal{D}_i, \mathcal{S}_i\} \leftarrow \mathtt{RecoverAggKey}(\varepsilon,$ $\{\mathsf{shares}_{i,\mathcal{D}}, \mathsf{shares}_{i,\mathcal{S}}\})$, which allows the server to recover the aggregate key and thus compute the aggregated input as described above. Intuitively, in this protocol each client i sends a share of the self seed for each surviving neighbor (in $\mathcal{S}_i$) and a share of the pairwise seed for each dropped out neighbor (in $\mathcal{D}_i$), which allows the server to recompute the mask and learn the aggregate encoding key. In more detail, the server computes the aggregate key $\mathbf{sk}$ as

$$\mathbf{sk} = \sum_{i \in \mathcal{S}} (\mathbf{s_i} + \sum_{j \in \mathcal{D}_i, j<i} \mathbf{s_{ij}} - \sum_{j \in \mathcal{D}_i, i<j} \mathbf{s_{ij}}),$$

where $\mathbf{s_i}$ is reconstructed from the shares provided by the neighbors of a surviving client i and $\mathbf{s_{ij}}$ is reconstructed from the shares provided by i for dropped out neighbors j. Crucially, this process requires clients and the server to agree on the set of dropouts and survivors, as otherwise even honest clients could inadvertently reveal information that would allow the server to unmask an individual honest contribution.

3.1 PRG-SecAgg

We can recover the original PRG-SecAgg protocol [6] by instantiating F as a seed-stretching PRG and the encoding scheme as follows, for $\mathbf{sk}_i, \mathbf{x}_i, \mathbf{y} \in \mathbb{Z}_q^{\ell}$:

$$\begin{aligned} \mathsf{Encode}(\mathbf{sk}_i, \mathbf{x}_i) &:= \mathbf{sk}_i + \mathbf{x}_i \bmod q \\ \mathsf{Decode}(\mathbf{sk}_i, \mathbf{y}) &:= \mathbf{y} - \mathbf{sk}_i \bmod q \end{aligned} \tag{4}$$

3.2 RLWE-SecAgg

Our second SecAgg instantiation, RLWE-SecAgg, leverages an encoding based on RLWE. In this case, the key expansion algorithm samples a key from the appropriate RLWE secret distribution χ_s and then generates the masks as RLWE samples using the expanded key. This combined process of key sampling and mask generation is much more computationally efficient than the key expansion in PRG-SecAgg.

Public parameters: Vector length ℓ, input domain $\mathbb{X}^\ell$, secret distribution χ_s, and seed expansion function $\mathsf{F}\colon \{0,1\}^\lambda \mapsto \mathsf{supp}(\chi_s)^\ell$
Client i's input: $\mathbf{x}_i \in \mathbb{X}^\ell$
Server output: $z \in \mathbb{X}^\ell$

1. Using the server to send messages, clients engage in the `ShareSeeds` protocol, with each surviving client i learning $\{\mathbf{seed}_{i,j}\}_{j \in \mathsf{N}(i)}$, $\mathbf{seed}_i$, $\mathsf{shares}_{i,\mathcal{D}}$, and $\mathsf{shares}_{i,\mathcal{S}}$. The server aborts if there are fewer than $(1-\delta)n$ surviving clients.
2. Each surviving client i performs the following:
   - Computes its *packed encrypted input* $\mathbf{y}_i = \mathsf{Encode}(\mathbf{sk}_i, \mathbf{G}\mathbf{x}_i)$ with key defined as $\mathbf{sk}_i = \mathbf{s}_i + \sum_{j \in A_i, j<i} \mathbf{s}_{ij} - \sum_{j \in A_i, i<j} \mathbf{s}_{ij}$ for $\mathbf{s}_{ij} = \mathsf{F}(\mathbf{seed}_{i,j})$, $\mathbf{s}_i = \mathsf{F}(\mathbf{seed}_i)$ (as in Equation 3).
   - Forms commitments $\mathsf{com}_{\mathbf{sk},i}$ and $\mathsf{com}_{\mathbf{x},i}$ to its key and input respectively.
   - Computes proofs $\pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}$, $\pi^{0 \le \mathbf{x}_i < t}$, and $\pi^{\mathsf{valid}(\mathbf{x}_i)}$ of encoding, smallness, and validity.
   - Sends to the server $\mathbf{y}_i$, $\mathsf{com}_{\mathbf{sk},i}$, $\mathsf{com}_{\mathbf{x},i}$, $\pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}$, $\pi^{0 \le \mathbf{x}_i < t}$, $\pi^{\mathsf{valid}(\mathbf{x}_i)}$.
3. The server aborts if it receives fewer than $(1-\delta)n$ messages or if any of the proofs fail to verify. Otherwise, the server and the clients engage in the `RecoverAggKey` protocol, with the server taking as input the global sets $\mathcal{D}$ and $\mathcal{S}$ of dropouts and survivors and each client i taking as input its sets $\mathsf{shares}_{i,\mathcal{D}}$ and $\mathsf{shares}_{i,\mathcal{S}}$ and providing the appropriate shares to the server according to the status of their neighbors. At the end of the protocol the server learns the aggregate key $\mathbf{sk}$.
4. Each client, acting as a distributed prover, engages with the server (acting as the verifier) in the distributed key correctness protocol. The server aborts if the collective proof fails to verify.
5. The server outputs $\sum_{i \in \mathcal{S}} \mathbf{x}_i$ as $\mathbf{G}^{-1}(\mathsf{Decode}(\mathbf{sk}, \sum_{i \in \mathcal{S}} \mathbf{y}_i))$.

Figure 1: General SecAgg protocol with input verification.

Unlike PRG-SecAgg, this encoding requires a set of public parameters: a polynomial ring $R = \mathbb{Z}[X]/(X^N+1)$ and its residual ring $R_q = \mathbb{Z}[X]/(q, X^N+1)$ for a modulus q, a plaintext modulus T that is coprime to q, a plaintext dimension ℓ, a secret key distribution χ_s and an error distribution χ_e over R, and a matrix $\mathbf{A}$ generated as discussed in Section 2.3. These parameters can be distributed to the clients by the server or through a public channel. They are used in the encoding and decoding algorithms, defined as follows:

$$\begin{aligned} \mathsf{Encode}(\mathbf{sk}_i, \mathbf{x}_i) = \mathbf{y}_i &:= \mathbf{A} \cdot \mathbf{sk}_i + T(\mathbf{e}+\mathbf{f}) + \mathbf{x}_i \bmod q, \\ &\quad \text{where } \mathbf{e}, \mathbf{f} \leftarrow \chi_e^{\ell/N}. \qquad (5) \\ \mathsf{Decode}(\mathbf{sk}, \mathbf{y}) &:= (\mathbf{y} - \mathbf{A} \cdot \mathbf{sk} \bmod q) \bmod T. \end{aligned}$$

We present formal proofs of the correctness and security of this encoding in the full version of this paper [5], but provide some intuition for them here.

Correctness. To ensure that the obtained result is the sum of the $\mathbf{x}_i \in \mathbb{Z}_t$ over the integers we need that (i) the sum of errors and messages does not overflow the ciphertext modulus q, and (ii) the sum of the messages does not overflow the plaintext modulus T. These result in the constraints $2nTb_e < q$ and $nt < T$, where b_e is an L_∞ bound on the error $e \leftarrow \chi_e$.

Security. It is tempting to claim that all we need for security is to choose RLWE parameters in a way that ensures the individual encodings $\mathbf{y}_i$ of client contributions are pseudorandom in isolation. However, the server gets more information than just n independent RLWE ciphertexts, as it can also recover $\sum_i \mathbf{e}_i$ from $\tilde{\mathbf{y}} = \sum \mathbf{y}_i$. A common approach to eliminate leakage is to add a large noise to "drown" the error [18, Chapter 21], in a way analogous to how circuit privacy is achieved in some (R)LWE-based homomorphic encryption schemes. The resulting modulus q would be very large, however, which hurts both computation and communication.

Instead, we argue in the full version that the encodings of all clients' inputs are indistinguishable from random values that sum up to an encoding of the sum of all inputs. This property can be established from the hardness of an RLWE variant, *Hint-RLWE*, in which samples consist of standard RLWE pairs $(a, as+e) \in R_q^2$ and a "hint" $e+f$, where f is sampled from the same Gaussian distribution as e. The additional noise term f allows us to gradually break the correlation among the shared secrets used in the ciphertexts of neighboring clients, via a carefully constructed hybrid argument. Lee et al. [31] showed that the Hint-RLWE problem with error size σ is as hard as the standard RLWE with error size $(1/\sqrt{2})\sigma$. The error terms in our RLWE encodings are thus only slightly larger than standard RLWE encryption, avoiding the need for noise flooding.

Ciphertext expansion. PRG-SecAgg has very limited ciphertext expansion, which is optimal in the sense that the modulus q can be chosen to be exactly tn, to ensure that the result of adding all n values fits in the modulus. This results in only a $1 + \frac{\log_2 n}{\log_2 t}$ factor overhead with respect to an insecure solution where clients just send their values. A naïve encoding in RLWE-SecAgg that puts each entry of $\mathbf{x}_i$ in a polynomial coefficient would result in a $1 + \frac{\log_2 q}{\log_2 t}$ factor overhead. This can be quite wasteful, as q needs to be large ($\geq 2^{46}$) for security. However, we can use a larger plaintext modulus T to *pack* multiple entries of $\mathbf{x}_i$ in a plaintext slot.

In more detail, let $\mathbf{G}$ be the gadget matrix $\mathbf{G} = (1, t, t^2, \ldots, t^{p-1}) \otimes \mathbf{I}_{l/p}$ for $p = \lfloor \log T / \log(nt) \rfloor$. Then by computing $\mu_i = \mathbf{G}\mathbf{x}_i \in [T]^{l/p}$, we effectively pack every p entries of the input $\mathbf{x}_i$ into a single plaintext slot of μ_i while ensuring that the result of the *packed* sum fits in T. To decode from a packed slot, one can apply a digit extraction algorithm for base t, denoted by $\mathbf{G}^{-1}$, which can be naturally extended to a packed vector. Importantly, this packing operations is linear, and thus it can be incorporated into the input validity constraints we consider in the next section.

4 Adding Input Validation

In this section we present ACORN, an extension to the generalized SecAgg protocol that allows for client input validation. Specifically, we provide a way for the server to check that the (hidden) inputs of clients satisfy some pre-defined notion of validity and that their messages in the protocol have been computed according to its specification. We first present ACORN-detect, where the server can detect that misbehavior has occurred but cannot attribute it to an individual client or recover from it, and then present ACORN-robust, in which the server can both identify misbehaving clients and remove their contributions from the final sum.

To achieve this, as described below we require non-interactive zero-knowledge proofs of vector smallness and valid encoding, and an interactive proof for the correctness of an aggregated key. We instantiate these primitives in Section 5 with efficient discrete log-based protocols.

4.1 Detecting Client Misbehavior

We present our summary protocol of ACORN-detect in Figure 1 and our detailed protocol in Figure 6, where the additional steps required for input validation are in red. Across the entire protocol, we require a zero-knowledge proof of the following relation R:

$$\Big\{(x,w) \mid x = ((\mathbf{y}_i, \mathsf{com}_{\mathbf{x},i}, \mathsf{com}_{\mathbf{sk},i})_{i\in\mathcal{S}}, \mathbf{sk}, t, \ell, \mathbf{G}), w = (\mathbf{x}_i, \mathbf{sk}_i, r_i, s_i)_{i\in\mathcal{S}},$$
$$\forall i \in \mathcal{S} : \big(\mathsf{com}_{\mathbf{x},i} = \mathsf{Commit}(\mathbf{x}_i; r_i), \mathsf{com}_{\mathbf{sk},i} = \mathsf{Commit}(\mathbf{sk}_i; s_i),$$
$$\mathbf{y}_i = \mathsf{Encode}(\mathbf{sk}_i, \mathbf{G}\mathbf{x}_i), \mathbf{x}_i \in \mathbb{Z}_t^{\ell}, \mathsf{valid}(\mathbf{x}_i)\big), \sum_{i\in\mathcal{S}} \mathbf{sk}_i = \mathbf{sk}\Big\}$$

We first observe that the witness for this relation is distributed among the clients, with each client i holding $\mathbf{x}_i$ and $\mathbf{sk}_i$ (and the relevant randomness) but being unaware of the other inputs. All the conditions of the relation except the last one, however, are on the individual components and thus each client can prove them independently. This means forming

1. A proof $\pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}$ that $\mathbf{y}_i = \mathsf{Encode}(\mathbf{sk}_i, \mathbf{G}\mathbf{x}_i)$, where $\mathbf{sk}_i$ and $\mathbf{x}_i$ are the values contained in the relevant commitments.
2. A proof $\pi^{\mathsf{valid}(\mathbf{x}_i)}$ that $\mathsf{valid}(\mathbf{x}_i)$ holds.
3. A proof $\pi^{0\leq \mathbf{x}_i < t}$ that $\mathbf{x_i} \in \mathbb{Z}_t^{\ell}$. This condition is needed to prove that no wraparound happens modulo the plaintext space, and thus that the packed sum can be decoded using $\mathbf{G}^{-1}$.

These proofs and the two commitments $\mathsf{com}_{\mathbf{sk},i}$ and $\mathsf{com}_{\mathbf{x},i}$ are sent to the server at the same time as the masked input $\mathbf{y}_i$.

With individual proofs for these individual constraints, the only remaining requirement of R is that $\sum_i \mathbf{sk}_i = \mathbf{sk}$. Clients could prove this individually by proving that they formed $\mathbf{sk}_i$ as specified by the protocol (Equation 3), but as the formation of $\mathbf{sk}_i$ requires key agreements and applications of a length-expanding function F this would be highly inefficient.

Instead, we have the clients prove *collectively* that $\sum_i \mathbf{sk}_i = \mathbf{sk}$, which is the minimal requirement needed for the server to decode and recover the aggregated inputs. This is done by having each client i provide a (partial) proof $\pi_i^{\sum \mathbf{sk}_i}$, which they can do without interacting with other clients. These proofs collectively demonstrate the correctness of the aggregated key $\mathbf{sk}$. Unlike the individual proofs, this proof cannot be made non-interactive, so we instead consider it as an interaction between each client and the server. In our summarized presentation in Figure 1 we present this as a separate step (Step 4), but in our detailed presentation in Figure 6 we show how this protocol can be woven into the broader SecAgg protocol without requiring any additional rounds of interaction.

Security. We formally prove the security of ACORN-detect, following a simulation-based argument [21, 32], in the full version of the paper [5]. Briefly, security for an honest server follows from the knowledge soundness of the proofs, which gives the simulator the ability (acting as the server) to extract the underlying inputs. Acting as the knowledge extractor means the simulator here needs to rewind, and thus that ACORN-detect is not concurrently secure in the honest server setting. Security in the malicious server setting is largely orthogonal to the question of input validation, and thus our proof follows closely the one of Bell et al. [6], with the simulator additionally relying on zero knowledge to ensure that its interactions with the adversary are indistinguishable from what it expects.

Efficiency. The asymptotic costs for the protocol, using the instantiations in Section 5, are in Table 1. Step 1 requires $\log(n)$ work per client, and is concretely very cheap. Client costs are thus dominated by the following tasks in Step 2: (1) running Encode and expanding seeds to compute $\mathbf{sk_i}$, (2) committing to $\mathbf{x_i}$ and $\mathbf{sk_i}$, and (3) generating proofs. The server's work is dominated by the analogous tasks of (1) verifying proofs in Step 3 and (2) computing $\mathbf{sk}$ in Step 4.

In PRG-SecAgg, the encoding step corresponds to PRG expansions (implemented with AES), and in RLWE-SecAgg it corresponds to the noisy linear transformation in Equation 5. As we show in Section 5, we use a discrete log-based proof system, and thus commitment generation means computing two Pedersen vector commitments (requiring two length-ℓ group multi-exponentiations), and proof generation requires $O(\ell \log(\ell))$ computation to produce a logarithmic size proof.

4.2 Robustness in the Face of Misbehavior

In the full version of the paper [5], we present ACORN-robust, which allows the server to not only identify misbehaving clients, but also exclude their input from the result on-the-fly.

This property, sometimes refereed to as *guaranteed output delivery* [26], ensures in this context that as long as the number of cheating clients stays below a given threshold α and no more than $(\delta-\alpha)n$ other clients drop out, an honest server is guaranteed a valid output. We present and prove ACORN-robust secure assuming a semi-honest server; an extension to a malicious server seems possible, but we leave this as future work.

5 Zero-Knowledge Constructions

In this section we provide constructions of the zero-knowledge proofs required in ACORN. Specifically, we instantiate proofs of *aggregated keys*, *correct encoding*, *input smallness*, and *input validity*. For ACORN-detect, we present the distributed proof of aggregated key correctness below, which is a variant of a Schnorr proof with an additional step to ensure that the interactive protocol achieves (full) zero knowledge. For ACORN-robust, this proof is integrated into the protocol itself. For the latter three predicates, we leverage the Bulletproofs system, which we present before presenting our proofs for these individual predicates.

5.1 Distributed Proof of Aggregated Key Correctness

In ACORN-detect, each client i commits to its encoding key sk_i, which incorporates its self mask and pairwise masks, using some randomness r_i; in other words, it creates an Pedersen commitment $C_i = \mathbf{g}^{\mathsf{sk}_i} h^{r_i}$. At the end of the protocol, the server learns the sum of self masks $\mathbf{s}$ in the clear. Our distributed key correctness protocol, presented in Figure 2, thus aims to convince the server that $\prod_i C_i$ is a commitment to $\mathbf{s}$, where the witness for this proof (consisting of the opening of $\prod_i C_i$) is additively shared among all the clients.

Our protocol uses the Schnorr proof of knowledge, with each client proving knowledge of the randomness r_i used to form its commitment. The server, acting as the verifier, can combine these proofs using the fact that the challenge response in Schnorr is a linear function of the committed value.

To tolerate a malicious server, we need this proof to be (fully) zero knowledge, which the sigma protocol can be when instantiated as a non-interactive proof using the Fiat-Shamir transformation. However, this does not work in our setting since clients act as distributed provers and do not have the same view that could then be hashed to form the challenge. Trying to send the required information to all clients to obtain such a common view is not viable since it incurs a prohibitive communication overhead.

Instead, we modify the execution so that the server commits to its challenge ahead of time using parameters provided by each client. To retain the property that this is a proof of knowledge, we require that this commitment is *equivocable*, as the knowledge extractor needs to be able to send two different challenges consistent with the same commitment. For example, Pedersen commitments provide equivocation when the discrete logarithm between the generators g and h is known.

Public parameters: group $\mathbb{G},\mathbb{H}$ with generators $\mathbf{g}\in\mathbb{G}^{\ell}$ and $h\in\mathbb{H}$

Client i's input: $\mathbf{sk}_i$, r_i, $C_i = \mathsf{Commit}(\mathbf{sk}_i; r_i)$

Server's input: $\mathbf{s}\in|\mathbb{G}|^{\ell}$, $\{C_i\}_i$

1. Client i samples α_i randomly and sends $h_i = h^{\alpha_i}$ to the server.
2. The server samples a random challenge e and forms a commitment $\mathsf{com}_{i,\mathsf{chl}} = g^e h_i^{s_i}$ for each client i, using some randomness s_i. The server sends $\mathsf{com}_{i,\mathsf{chl}}$ to client i.
3. Each client i samples a random value $k_i \in |\mathbb{H}|$ and sends $K_i = h^{k_i}$ to the server.
4. The server opens the committed challenge by sending e and s_i to client i.
5. Client i checks that $\mathsf{com}_{i,\mathsf{chl}} = g^e h_i^{s_i}$, where $\mathsf{com}_{i,\mathsf{chl}}$ is the value it received in Step 2. If not, it aborts. Otherwise, it computes $t_i = r_i\cdot e + k_i$ and sends it and α_i to the server.
6. The server checks that $h_i = h^{\alpha_i}$. If so, it computes $C = \prod_{i\in[n]} C_i/\mathbf{g}^{\mathbf{s}}$ and $t=\sum_{i\in[n]} t_i$. It checks that $h^t = C^e\prod_{i\in[n]} K_i$ and outputs 1 if this check passes and 0 otherwise.

Figure 2: Distributed key correctness protocol for proving that $\sum_i \mathbf{sk}_i = \mathbf{s}$.

Correctness. We can verify that

$$h^t = \prod_{i\in[n]} h^{t_i} = (\prod_{i\in[n]} h^{r_i})^e (\prod_{i\in[n]} h^{k_i}) = C^e \prod_{i\in[n]} K_i.$$

Knowledge soundness. The soundness of the protocol follows similarly to the soundness of the single prover Schnorr protocol. The extractor can rewind the execution of steps 3 and 4 with the i-th client and provide two different openings e_1 and e_2 for the committed challenge using the equivocability of the commitment scheme, and obtain two different values $t_{i,1}$ and $t_{i,2}$. If the proof verifies in both cases, then $h^{t_{i,1}} = h^{t_{i,2}}$, and the extractor can compute $r_i = (t_1-t_2)(e_1-e_2)^{-1}$ since $e_1-e_2\neq 0$.

Zero knowledge. The simulator for client i rewinds step 2 and 3 after it has obtained the opening of the commitment for challenge e. It generates t_i at random and sets $K_i = h^{t_i}(h^{r_i})^{-e}$.

5.2 Inner Product Proofs

We build our zero-knowledge proofs using Bulletproofs [11]. In the context of this work, similarly to Gentry et al. [19], we regard Bulletproofs as a zero-knowledge proof of knowledge of vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$ that satisfy an inner product constraint $\langle \mathbf{x}, \mathbf{y} \rangle = a$ in an order-q group G, where a is a public scalar.

At a high level, to prove $\langle \mathbf{x}, \mathbf{y} \rangle = a$ the prover recursively computes a new equation $\langle \mathbf{x}', \mathbf{y}' \rangle = a'$ for vectors of half the length, and computes commitments given a challenge sent by the verifier. This requires both the prover and verifier to compute new generators at each recursive step, with the prover also computing new commitments to $\mathbf{x}'$ and $\mathbf{y}'$. In the non-interactive variant using the Fiat-Shamir heuristic the generator computation can be unfolded, resulting in a single multi-exponentiation of length $2n$. We state the concrete costs for Bulletproofs in terms of multi-exponentiation operations, for which efficient sublinear algorithms are known [38]. The full protocol details are given by Gentry et al. [19, Section E.2].

Lemma 1 ([11, 19])**.** *Let $C \in \mathsf{G}$ be a group element, and let $\mathbf{h} \in \mathsf{G}^2, \mathbf{g} \in \mathsf{G}^{2n}$ be sets of generators in G known to both the prover and verifier. Bulletproofs allows the prover to prove knowledge of vectors $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_q^n$ and randomness $r \in \mathbb{Z}_q$ such that $C = h_0^r h_1^{\langle \mathbf{x}, \mathbf{y} \rangle} \prod_{i=1}^n \mathbf{g}_i^{\mathbf{x}_i} \mathbf{g}_{i+n}^{\mathbf{y}_i}$. It satisfies perfect completeness, statistical zero knowledge, and computational knowledge soundness under the discrete logarithm assumption.*

When compiled into a NIZK proof, the prover performs $6n + 8(\log_2(n) + 1)$ group exponentiations (computed as several multi-exponentiations), and the verifier performs a single multi-exponentiation of length $2(n + \log_2(n) + 3)$. Moreover, batched verification of m proofs requires a single multi-exponentiation of length $2n + 2 + m \log_2(2n + 4)$. The proof size is $2 \log n + 4$ group elements.

5.3 Proofs of Smallness

We consider two variants of the problem of proving in zero knowledge that $\mathbf{x}_i \in [0, t-1]$ for all i: the first is efficient for $t = 2$, which is useful for applications that rely on binary k-hot encodings, while the second works for arbitrary values of t, which is useful in the learning setting and in our proof of correct encoding.

5.3.1 Proof for $t = 2$

We describe the protocol for $t = 2$ for simplicity, which closely follows the approach to range proofs in Bulletproofs [11]. Let $C \in \mathsf{G}$ be a group element, and let $h \in \mathsf{G}$ and $\mathbf{g}, \mathbf{h} \in \mathsf{G}^\ell$ be generators in G (public parameters).

1. The prover finds $\mathbf{y} \in \mathsf{G}^\ell$ satisfying (i) $\mathbf{x} \circ \mathbf{y} = 0$ and (ii) $\mathbf{x} = \mathbf{1}^\ell + \mathbf{y}$. These properties hold if and only if $\mathbf{x}$ is binary. The prover commits to $\mathbf{x}|\mathbf{y}$ as $C = h^r \prod_{i=1}^\ell \mathbf{g}_i^{\mathbf{x}_i} \mathbf{h}_i^{\mathbf{y}_i}$ and sends this to the verifier.

2. The verifier sends random challenge scalars $\tau, \rho \in \mathbb{Z}_q$ to the prover. Define $\mathbf{r} = (\tau^{i-1})_{i \in [\ell]}$.

3. By Schwartz-Zippel, $\langle \mathbf{x}, \mathbf{y} \circ \mathbf{r} \rangle + \rho(\langle \mathbf{x}, \mathbf{r} \rangle + \langle -\mathbf{1}^\ell, \mathbf{y} \circ \mathbf{r} \rangle) = \langle \mathbf{1}^\ell, \mathbf{r} \rangle$ holds if and only if (i) and (ii) hold, except with probability $(\ell+1)/q$. This can be rewritten as a single constraint $\langle \mathbf{x}', \mathbf{y}' \rangle = (1 - \rho^2)\langle \mathbf{1}^\ell, \mathbf{r} \rangle$, for $\mathbf{x}' := \mathbf{x} - \mathbf{1}^\ell \rho$ and $\mathbf{y}' := \mathbf{y} \circ \mathbf{r} + \rho \mathbf{r}$. The prover and verifier can obtain a commitment C' to $\mathbf{x}'|\mathbf{y}'$ by computing $C' = C \cdot \prod_{i=1}^\ell \mathbf{g}_i^{-\rho} \mathbf{h}_i^{\rho}$. This is a commitment to $\mathbf{x}'$ and $\mathbf{y}'$ using generators $(h, \mathbf{g}, \mathbf{h}')$ where $\mathbf{h}'_i = \mathbf{h}_i^{\tau^{1-i}}$. The prover then uses Bulletproofs as described in Lemma 1 to prove that $\langle \mathbf{x}, \mathbf{y}' \rangle = 0$ with respect to C'.

This proof is thus a reduction to Bulletproofs (Lemma 1), requiring three additional length-ℓ multi-exponentiations in Step 3 by both the prover and verifier. However, this overhead can be reduced to only two length-ℓ multi-exponentiations for the verifier, as both the generator switch and commitment update can be combined with the analogous operations in the outer loop of Bulletproofs. Moreover, the proof can be made non-interactive via Fiat-Shamir, by deriving τ, ρ from the protocol transcript and proof statement.

5.3.2 Proof for arbitrary t

To show that $\mathbf{x}_i \in [a, b]$ it suffices to show that $(\mathbf{x}_i - a)(b - \mathbf{x}_i)$ is non-negative. We thus show that $c_i := \mathbf{x}_i(t - 1 - \mathbf{x}_i) \geq 0$. A common way to do this is to exhibit a decomposition of c_i into four squares. However, a useful optimization consists of showing that $c'_i := 4c_i + 1 \geq 0$ [22]. These two conditions are equivalent over the integers, but because $c'_i \equiv 1 \bmod 4$ it can be written as a sum of three squares, where the three squares can be efficiently determined [40].

For convenience, we write c'_i as $1 + (t-1)^2 - (2\mathbf{x}_i - t + 1)^2$. The protocol thus proceeds by having the client prove that $c'_i \geq 0$ for all i, by showing that it knows $\mathbf{u}, \mathbf{v}, \mathbf{w}$ such that

$$\mathbf{x}' \circ \mathbf{x}' + \mathbf{u} \circ \mathbf{u} + \mathbf{v} \circ \mathbf{v} + \mathbf{w} \circ \mathbf{w} = \mathbf{a} \tag{6}$$

holds over the integers, where $\mathbf{x}' := 2 \cdot \mathbf{x} - (t-1) \cdot \mathbf{1}$ and $\mathbf{a} := (-(1 + (t-1)^2)) \cdot \mathbf{1}$ is public. The prover must also show that these computations do not wrap around the modulus q, which means showing that

$$\|\mathbf{x}'|\mathbf{u}|\mathbf{v}|\mathbf{w}\|_\infty < \sqrt{q}/4 \tag{7}$$

At this point it might seem that we're going in circles, as we reduced a range proof to two constraints, one of which is itself a range proof. However, this second bound is very loose, because $\sqrt{q}/4 \gg t$. In our implementation of Bulletproofs we have $q > 2^{250}$, while we are interested in values of t for natural datatypes, e.g. $t = 2^{16}$. Therefore we can take advantage of approximate range proofs introduced by Gentry et al. [19] and also used in [34], whose properties (assuming Fiat-Shamir) are summarized in the following lemma.

Lemma 2 ([19, Lemma 3.5])**.** *Fix a security parameter λ. Let $\mathbf{z} \in \mathbb{Z}^\ell$ be a vector such that $\|\mathbf{z}\|_\infty \leq t$, and let $\gamma > 1$ be* *such that $\gamma > 2500\sqrt{\ell}$. There is a ZK proof system to show $\|\mathbf{z}\|_\infty \leq \gamma \cdot t$ by proving a single constraint $\langle \mathbf{z}|\mathbf{y}, \mathbf{b}\rangle = s$ given vector commitments to $\mathbf{z}$ and $\mathbf{y}$, where $\mathbf{b} \in \mathbb{Z}_q^{\ell+\lambda}$ is a public vector, $\mathbf{y} \in [\pm\gamma t/2(1+1/\lambda)]^\lambda$, and $s \in \mathbb{Z}_q$.*

The requirement in Lemma 2 that $\sqrt{q}/(4t) > 2500\sqrt{4\ell}$ holds for Equation 7 as long as $\log_2(q) > 2\log_2(2500) + 2\log_2(t) + \log_2(\ell) + 6$, which combined with the fact that $q > 2^{250}$ means that even for $t = 2^{64}$ this approach can handle vectors of length $\ell > 2^{90}$ (well beyond realistic input sizes).

We can now describe our range proof protocol for large t. Let $h \in \mathsf{G}$ and $\mathbf{g} \in \mathsf{G}^{8\ell+\lambda}$ be public generators in G. We denote by $\mathsf{Commit}(h, \mathbf{g}[i : i+\ell-1]; \mathbf{v})$ the vector Pedersen commitment to $\mathbf{v}$ of length ℓ, using generators h and $\mathbf{g}_i, \ldots, \mathbf{g}_{i+\ell-1}$.

1. The prover finds auxiliary vectors $\mathbf{u}, \mathbf{v}, \mathbf{w}$ such that Equation 6 holds and "double-commits" to $\mathbf{z} := \mathbf{x}|\mathbf{u}|\mathbf{v}|\mathbf{w}$ as $C_1 := \mathsf{Commit}(h, \mathbf{g}[\ell : 4\ell]; \mathbf{z})$ and $C_2 := \mathsf{Commit}(h, \mathbf{g}[4\ell+1 : 8\ell]; \mathbf{z})$. Let $\langle \mathbf{z}|\mathbf{y}, \mathbf{b}\rangle = s$ be the constraint that proves Equation 7 as in Lemma 2. Then the prover computes $C_y = \mathsf{Commit}(h, \mathbf{g}[8\ell+1 : 8\ell+\lambda]; \mathbf{y})$ and sends C_1, C_2, and C_y to the verifier.

2. The verifier sends random challenge scalars $\sigma, \tau, \rho \in \mathbb{Z}_q$ to the prover. Let $\mathbf{r_x}|\mathbf{r_u}|\mathbf{r_v}|\mathbf{r_w} = (\sigma^{i-1})_{i\in[4\ell]}$ and $\mathbf{r} = (\tau^{i-1})_{i\in[\ell]}$ be the corresponding challenge vectors.

3. The following constraint is equivalent to Equation 6, except with probability bounded by $(\ell+4)/q$:

   $$\langle \mathbf{x}', \mathbf{x}' \circ \mathbf{r}\rangle + \langle \mathbf{u}, \mathbf{u} \circ \mathbf{r}\rangle + \langle \mathbf{v}, \mathbf{v} \circ \mathbf{r}\rangle + \langle \mathbf{w}, \mathbf{w} \circ \mathbf{r}\rangle = \langle \mathbf{a}, \mathbf{r}\rangle.$$

   Instead of proving this directly, we use a trick from Gentry et al. [19] to express it as

   $$\langle \mathbf{x}' + \mathbf{r_x}, (\mathbf{x}' - \mathbf{r_x}) \circ \mathbf{r}\rangle + \langle \mathbf{u} + \mathbf{r_u}, (\mathbf{u} - \mathbf{r_u}) \circ \mathbf{r}\rangle + \langle \mathbf{v} + \mathbf{r_v}, (\mathbf{v} - \mathbf{r_v}) \circ \mathbf{r}\rangle + \langle \mathbf{w} + \mathbf{r_w}, (\mathbf{w} - \mathbf{r_w}) \circ \mathbf{r}\rangle = s,$$

   where $s = \langle \mathbf{a}, \mathbf{r}\rangle - (\|\mathbf{r_x}\|^2 + \|\mathbf{r_u}\|^2 + \|\mathbf{r_v}\|^2 + \|\mathbf{r_w}\|^2)$ is a public value. This constraint can be rewritten as a single inner product $\langle \bar{\mathbf{x}}|\bar{\mathbf{u}}|\bar{\mathbf{v}}|\bar{\mathbf{w}}, \bar{\mathbf{x}}'|\bar{\mathbf{u}}'|\bar{\mathbf{v}}'|\bar{\mathbf{w}}'\rangle = a$ for some $a \in \mathbb{Z}_q$, where $\bar{\mathbf{x}} := \mathbf{x} + \mathbf{r_x}$ and $\bar{\mathbf{x}}' := (\mathbf{x} - \mathbf{r_x}) \circ \mathbf{r}$, and analogously for the rest. By proving these constraints against C_1 and C_2, the prover is effectively showing that C_1 and C_2 are commitments to the same vector $\mathbf{z}$. Next, to incorporate the constraint $\langle \mathbf{z}|\mathbf{y}, \mathbf{b}\rangle = s$ that proves Equation 7, the prover can replace $\mathbf{z}$ with $\bar{\mathbf{z}} := \bar{\mathbf{x}}|\bar{\mathbf{u}}|\bar{\mathbf{v}}|\bar{\mathbf{w}}$ and prove an equivalent constraint $\langle \bar{\mathbf{z}}|\mathbf{y}, \mathbf{b}\rangle = a'$, where $a' := s - \langle \mathbf{b}, \mathbf{r_x}|\mathbf{r_u}|\mathbf{r_v}|\mathbf{r_w}|\mathbf{0}^{|\mathbf{y}|}\rangle$.

   At this point, the prover has the constraints $\langle \bar{\mathbf{z}}, \bar{\mathbf{x}}'|\bar{\mathbf{u}}'|\bar{\mathbf{v}}'|\bar{\mathbf{w}}'\rangle = a$ and $\langle \bar{\mathbf{z}}|\mathbf{y}, \mathbf{b}\rangle = a'$, which can be merged into a single constraint using ρ, resulting in a single inner product:

   $$\langle \bar{\mathbf{z}}|\mathbf{y}, \bar{\mathbf{x}}'|\bar{\mathbf{u}}'|\bar{\mathbf{v}}'|\bar{\mathbf{w}}'|\rho\mathbf{b}\rangle = a + \rho a'. \tag{8}$$

   The prover and the verifier both obtain a commitment C' to $\bar{\mathbf{z}}|\mathbf{y}|\bar{\mathbf{x}}'|\bar{\mathbf{u}}'|\bar{\mathbf{v}}'|\bar{\mathbf{w}}'|\rho\mathbf{b}$ from C_1, C_2 and C_y, which can be done using only linear operations. The prover can then use Bulletproofs to prove the constraint in Equation 8.

In addition to using Bulletproofs, the above protocol requires three length-8ℓ multi-exponentiations by the prover and verifier to compute C' in step 3. As in the protocol for $t = 2$, the prover and verifier also need to switch the generators to match the commitment C'. The prover can again combine the process of both switching generators and updating the commitment with the analogous operations in the outer loop of Bulletproofs, thus computing them almost for free. The only remaining overhead is two multi-exponentiations of length 8ℓ for the verifier, but this overhead can be batched as it depends on public (not proof-specific) generators.

For $\mathbf{x} \in \mathbb{Z}_t^\ell$, the overall proof system $\pi^{0\leq\mathbf{x}<t}$ thus has the costs stated in Lemma 1, using $n = 8\ell + \lambda$, with two additional multi-exponentiations of length 8ℓ for the verifier and commitment cost (to vectors of total length $8\ell + \lambda$ with entries in $[t]$) for the prover.

5.4 Proofs of Validity of Encoding

In all variants of the protocol, $\pi^{\mathsf{Enc}(\mathbf{sk},\mathbf{x})}$ reduces to proving an inner product constraint involving *public* packing matrix $\mathbf{G} \in \mathbb{Z}_q^{\bar{\ell}\times\ell}$ and masked input $\mathbf{y_i} \in \mathbb{Z}_q^{\bar{\ell}}$, and a *private* committed input vector $\mathbf{x_i}$, for each client i. In the PRG-based ACORN-detect, for example, we rely on the constraint $\mathbf{y_i} = \mathbf{sk}_i + \mathbf{G}\mathbf{x}_i$ for a committed key $\mathbf{sk}_i \in \mathbb{Z}_q^{\bar{\ell}}$, while for the RLWE-based variant we rely on the constraint $\mathbf{y_i} = \mathbf{A}\mathbf{sk}_i + T\mathbf{e}_i + \mathbf{G}_k\mathbf{x}_i + q\mathbf{d}_i \wedge \|\mathbf{e}_i\|_\infty < b_e$ for committed key and error term $\mathbf{sk}_i, \mathbf{e_i} \in \mathbb{Z}_q^{\bar{\ell}}$ and for some private $\mathbf{b}_i \in \mathbb{Z}_q^{\bar{\ell}}$.

In general, these required constraints can be written as a single constraint $\langle \mathbf{x}|\mathbf{v_1}|\cdots|\mathbf{v_k}|, \mathbf{b}\rangle = a$ using Schwartz-Zippel as in the smallness proofs in Section 5.3, and then using the smallness proofs directly to prove that $\|\mathbf{e}_i\|_\infty < b_e$. For PRG-SecAgg, these reductions to a single constraint do not have any overhead. For RLWE-SecAgg, the combined proof requires additional multi-exponentiations due to the secret multiplier $\mathbf{d}_i$ and the smallness proof of $\mathbf{e}_i$.

5.5 Other Validity Predicates

We have already presented proofs of one useful validity predicate: $\mathsf{valid}(\mathbf{x}) := \mathbf{x} \in [0,t)^\ell$, which extends to $\mathsf{valid}(\mathbf{x}) := \|\mathbf{x}\|_\infty = t$. We now discuss useful variants related to bounding L_0 and L_2 norms. We first observe that proving k-hotness, i.e. $\mathsf{valid}(\mathbf{x}) := \mathbf{x} \in \{0,1\}^\ell \wedge \|\mathbf{x}\|_0 = k$, can be achieved by just merging the constraint $\langle \mathbf{x}, \mathbf{1}^\ell\rangle = k$ with the proof of Section 5.3.1, which does not add any overhead. Let us also consider how to prove $\mathsf{valid}(\mathbf{x}) := \mathbf{x} \in [0,t) \wedge \|\mathbf{x}\|_2 \leq b$ for some public bound b, where t could be replaced by some natural bit-width like 2^{16} or 2^{32}. This can be done in two steps: first we establish that $\|\mathbf{x}\|_2 \leq \eta b$ using an approximate L_2 proof such that $\eta b < q/2$ for some gap parameter $\eta > 1$, and then we apply Lagrange's four-square theorem and prove that

$$\langle \mathbf{x}|\upsilon_0|\upsilon_1|\upsilon_2|\upsilon_3, \mathbf{x}|\upsilon_0|\upsilon_1|\upsilon_2|\upsilon_3\rangle = b \tag{9}$$

where $\upsilon_0, \ldots, \upsilon_3$ are integers guaranteed to exist if $\|x\|_2 \leq b$. Recently, Lyubashevsky et al. [34, Lemma 2.9] showed that the approximate L_2 bound proof can be adopted from the approximate L_∞ bound proof of Lemma 2. These two proofs can be combined with the proofs of Section 5.3 to show that Equation 9 holds over the integers, and the overhead is the additional commitments to $\upsilon_0, \ldots, \upsilon_3$ (twice using different sets of generators) and the increased inner product constraint length (by 4). The details of this extension were given by Gentry et al. [19, Section 3.5].

6 Implementation and Evaluation

In this section we present experimental results for our new protocols, focusing on RLWE-SecAgg (in Section 6.1) and ACORN-detect (in Section 6.2). We do not benchmark ACORN-robust, but as described in Section 4.2 these costs can be derived from those of ACORN-detect (taking into account the additional vector commitments clients must form).

We focus on two scenarios when setting experiment parameters: (1) federated learning (FL) applications with $n = 500$ clients and input vectors $\mathbf{x}$ containing 16-bit integers (i.e. $t = 2^{16}$); and 2) federated aggregation (FA) applications with $n = 10000$ clients and input vectors $\mathbf{x}$ containing binary values (i.e. $t = 2$). In both settings, we consider input vectors of length ℓ ranging from 2^{10} to 2^{20}, which covers a wide range of real-world scenarios. Our experiments were performed on a laptop with an Intel i7-1185G7 CPU running at 3GHz and with 16GB memory, in single thread mode, and we take advantage of SIMD instructions such as AVX512. For the input validation steps we also performed experiments on a Pixel 6 Pro smartphone.

6.1 RLWE-SecAgg

RLWE parameters. We first set the error distribution χ_e for sampling $\mathbf{e}$ and $\mathbf{f}$ in our encryption scheme to be a discrete Gaussian D_{σ_1} with standard deviation $\sigma_1 = 4.5$. As shown in the full version of the paper [5], the security level of our RLWE encryption with this error distribution can be derived from the hardness of solving RLWE with a discrete Gaussian error distribution of standard deviation $\sigma = 3.2$. In the implementation we use a tail-cut discrete Gaussian with support $[-60, 60]$ to sample $\mathbf{e} + \mathbf{f}$, which is statistically close to $D_{2\sigma_1}$ with distance at most 2^{-30}. When input validation is not required, we choose a power-of-2 ring degree $N \in \{2^{11}, 2^{12}, 2^{13}\}$, and we pick a prime modulus $q = 1 \pmod{2N}$ with 155 bits of security according to the lattice estimator [3]; such q takes advantage of Number Theoretic Transform (NTT) for fast polynomial multiplication. With

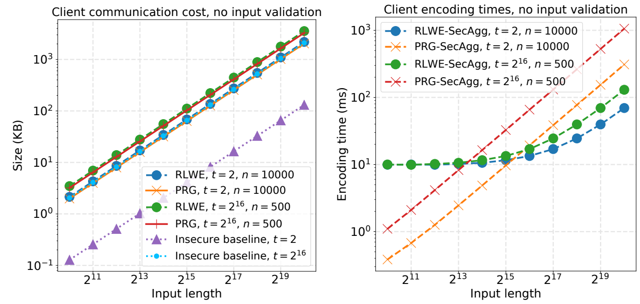

Figure 3: Averaged over 100 runs and plotted on a logarithmic scale, the message size (on the left) and encoding times (on the right) for our two scenarios for RLWE-SecAgg, PRG-SecAgg, and the insecure baseline where clients send the (uncompressed) input vector in the clear.

these parameters our RLWE encoding scheme achieves at least 128 bits of security with $n \leq 10^4$ clients. Furthermore we pick ring parameters such that it can optimally accommodate input messages via packing. With input validation, we choose $N \in \{2^{11}, 2^{12}\}$ and a prime q of at most 96 bits and achieve the same level of security as above. As shown in the full version of this paper [5], these parameters allow us to prove valid RLWE encodings using Bulletproofs based on curve25519.

Ciphertext expansion. Since the RLWE modulus q is usually much larger than the input bound t, we pack multiple input entries into a single plaintext slot. In addition, when the packed input vector has length $\bar{\ell} < N$, each client i sends just the first $\bar{\ell}$ coefficients $\mathbf{y}'_i = \mathbf{y}_i[1 \ldots \bar{\ell}]$ instead of the full $\mathbf{y}_i$. The server can still recover the aggregated input from $\sum_i \mathbf{y}'_i$ and the first ℓ rows of the public randomness $\mathbf{A}$. In contrast, in PRG-SecAgg we have more choices for the modulus q. For PRG-SecAgg without input validation, we can set the modulus $q = nt$ to achieve the optimal ciphertext expansion ratio, which is the total ciphertext bit-size over the input bit-size; when input validation is required, we set the modulus $q = P$.

Experimental results. We benchmarked the RLWE encoding step for individual clients, which involves expanding seeds to secret keys and encoding the packed input with the properly aggregated secret keys. For comparison, we also benchmarked the PRG encoding step, where seeds are expanded by repeatedly calling AES to the desired length, and masking is done via modular addition. The results are in Figure 3.

For example, in the FL use case where $t = 2^{16}$ and $n = 500$, when the input has length $\ell = 2^{16}$, RLWE encoding takes only 17ms while PRG encoding takes 65ms; for an input of length $\ell = 2^{20}$, RLWE encoding takes 130ms while PRG encoding takes 1.06s. Figure 3 shows our encoding benchmark results of both RLWE-SecAgg and PRG-SecAgg. Overall, RLWE encoding achieves roughly up to 5x speedup in the FA setting for $\ell \geq 2^{15}$, and up to more than 8x speedup in the FL setting

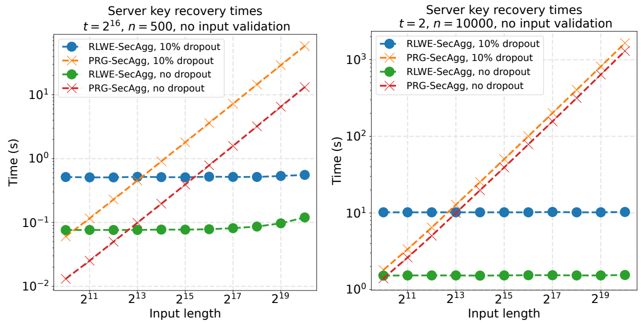

Figure 4: Server key recovery experiment results for RLWE-SecAgg and PRG-SecAgg. The runtimes are averaged over 100 runs and plotted on a logarithmic scale. The diagram on the left shows runtimes for the FL scenario ($t = 2^{16}$ and $n = 500$), and the diagram on the right for the FA scenario ($t = 2$ and $n = 10^4$). We consider cases with no dropouts and with a $\delta = 1/10$ fraction of dropouts.

for $\ell \geq 2^{13}$; for shorter input $\mathbf{x}$ the time spent on RLWE secret sampling is more significant than the PRG mask expansion.

We also benchmarked the server key recovery step for RLWE-SecAgg and PRG-SecAgg. The results are shown in Figure 4. For RLWE-SecAgg, the key recovery step includes expanding seeds to RLWE secrets of length N and decoding the RLWE masked sum, which involves an NTT operation per RLWE ciphertext. We see from the results that the RLWE key recovery times are dominated by seed expansion, which is independent of the input length ℓ, and the time spent on decoding the masked sum was not significant except when n is small and ℓ is very large. Compared to PRG-SecAgg, the RLWE key recovery step is much more efficient for long inputs: for example, with $n = 10000$ it takes only 10.3s to recover all secrets and decode the masked sum in RLWE-SecAgg when $\ell = 2^{20}$ for all $n = 10000$ clients with 10% dropout rate, while in PRG-SecAgg the same step requires 1650s.

6.2 ACORN-detect

While RLWE-SecAgg is more efficient without input validation, it becomes less efficient on the client when input validation is required due to additional proofs of smallness of the error terms. For example, when input $\mathbf{x}$ is a binary vector of length 2^{16}, generating the input validation proof for RLWE-SecAgg takes 8.4s whereas it takes 4.6s for PRG-SecAgg. For input $\mathbf{x}$ of length 2^{16} and l_∞ norm $t = 2^{16}$, it takes 26.07s to generate RLWE-SecAgg input validation proofs and 18.7s for PRG-SecAgg. We thus focus our results only on the PRG variant of ACORN-detect, and refer the readers to the full version for the RLWE variant.

We benchmarked the main components of ACORN-detect with graph parameters $\gamma = \delta = 1/10$. For the client, these consist of the encoding step and generating the necessary

	Client computation time			Client bandwidth	
Type	Round (s)	Total (m)	Factor	Total (MB)	Factor
MNIST (19k params, 160 rounds, training time per round 1.17s)					
SA	1.19	3.22	1x	5.41	1x
L_2	6.65	17.79	5.52x	5.68	1.05x
L_∞	6.59	17.63	5.47x	5.67	1.05x
CIFAR-10 S (62k params, 100 rounds, training time per round 0.83s)					
SA	0.86	1.48	1x	11.02	1x
L_2	18.81	31.39	21.28x	11.20	1.02x
L_∞	18.28	30.50	20.68x	11.19	1.02x
CIFAR-10 L (273k params, 160 rounds, training time per round 20.75s)					
SA	20.91	55.82	1x	77.66	1x
L_2	100.33	267.60	4.79x	77.97	1.01x
L_∞	96.55	257.53	4.61x	77.96	1.01x
Shakespeare (818k params, 20 rounds, training time per round 284.16s)					
SA	284.64	94.89	1x	29.08	1x
L_2	519.92	173.31	1.83x	29.13	1.01x
L_∞	510.42	170.15	1.79x	29.12	1.01x

Table 2: End-to-end model training with ACORN-detect. "SA" denotes experiments using PRG-SecAgg, which includes local training and secure aggregation. L_2 and L_∞ denote experiments using ACORN-detect with L_2 and L_∞ validity proofs. Computation time and bandwidth are reported per client, and overhead factors are with respect to secure aggregation (SA).

commitments to $\mathbf{sk}$ and $\mathbf{x}$ and proofs $\pi^{0 \leq \mathbf{x} < t}$ and $\pi^{\mathsf{Enc}(\mathbf{sk},\mathbf{x})}$. As discussed in Section 5.5, these costs also cover $\pi^{\mathsf{valid}(\mathbf{x})}$ for various validity predicates (e.g., one-hotness and both L_∞ and L_1 bounds). For the server, this consists of proof verification and key recovery steps. Figure 5 shows the client and server runtimes as well as the client communication costs for both settings we consider, where all benchmarks were run on a laptop. When running the client computations on the Pixel 6 Pro smartphone, we observed an average slowdown of 3X.

Encoding. We set the mask modulus q to the group size of curve25519 to match our Bulletproofs implementation. Comparing to PRG-SecAgg without input validation, this modulus q is less optimal in terms of packing capacity, and as a result, encoding times are increased by 40% to 70%. Regardless, encoding still takes less than one second for all but one input lengths (the exception being vectors of length 2^{20} in the federated learning use case).

Commitment generation. When $t = 2$, commitment generation is fast and grows slowly even for long inputs: for inputs of length $\ell = 2^{20}$ the commitments can be generated in 404ms. When input entries are large ($t = 2^{16}$), commitment generation is slower, but can still finish in 1.13s for $\ell = 2^{17}$. On the Pixel 6 Pro, commitments can be generated in 734ms for $t = 2$ and $\ell = 2^{20}$, and in 2.1s for $t = 2^{16}$ and $\ell = 2^{17}$.

Proof generation. We implemented the more efficient Bulletproofs variant due to Gentry et al. [19, Section E.2]. Our implementation [16] further optimizes proof generation by

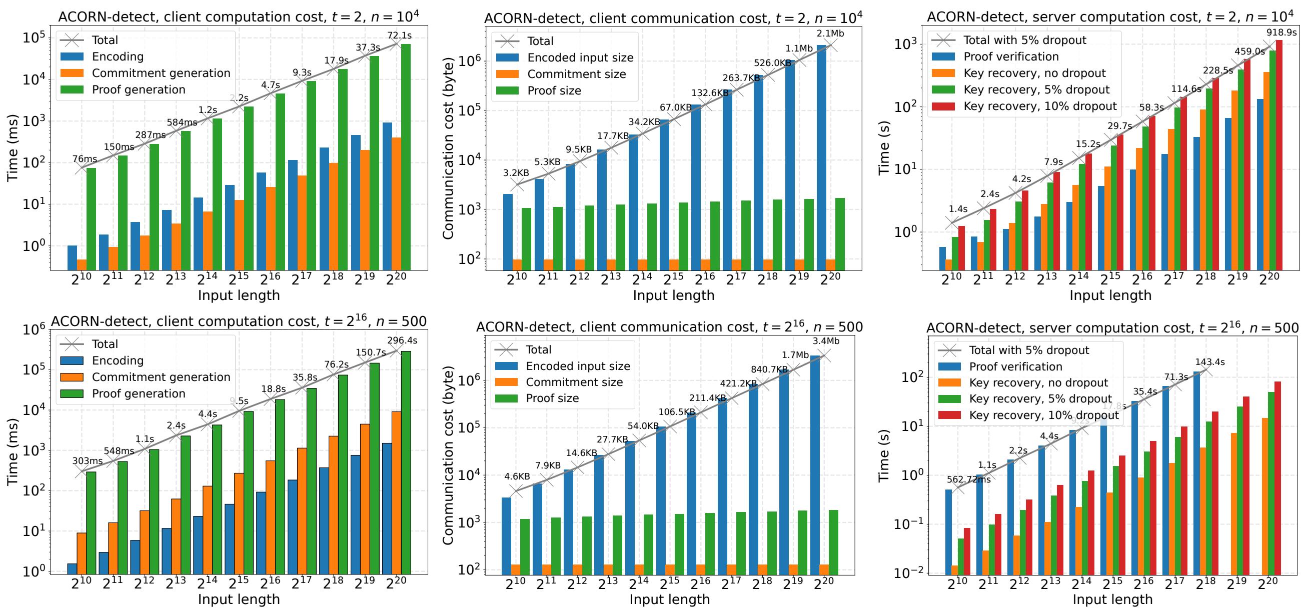

Figure 5: On a logarithmic scale, benchmarks for ACORN-detect for the FA use case where $t = 2$ and $n = 10000$ (on the top row) and the FL use case where $t = 2^{16}$ and $n = 500$ (on the bottom). In both cases we measured (1) the client runtime (on the left); (2) the client communication cost (in the middle); and (3) the server runtime (on the right), considering proof verification and key reconstruction for three different levels of dropouts. Our proof verification experiments for $t = 2^{16}$ and input length $\ell \geq 2^{19}$ ran out of memory, which is why these bars are missing in the bottom right diagram. The encoding experiments were repeated 100 times for each parameter set, and the other experiments were repeated for at least 5 seconds or 10 iterations; the average running times are estimated using the bootstrap resampling method with 95% confidence level.

not requiring the client to pad the inner product constraints to a power-of-2 length, which saves almost half of the proof generation time when the input is exactly or slightly longer than a power of 2. When $t = 2$, all proofs can be generated in 572ms for inputs of length $\ell = 2^{13}$ and in 70s for length $\ell = 2^{20}$. When $t = 2^{16}$, the combined linear constraint is roughly four times longer than in the $t = 2$ case, so proof generation is slower: it runs in 2.27s for inputs of length 2^{13} and in 285s for length 2^{20}. For comparison, proof generation on the Pixel 6 Pro for $\ell = 2^{13}$ takes 2.1s for $t = 2$ and 8.2s for $t = 2^{16}$.

Proof verification. The verification step also takes advantage of the lightweight linear proof optimization, and we benchmarked the batched verification of proofs from all n clients using the techniques mentioned in Section 5.2. As we can see, batched proof verification in the binary case is very efficient due to the smaller size of proofs and the SIMD acceleration: verifying all proofs from 10,000 clients takes 1.7s for inputs of length $\ell = 2^{13}$ and 133.2s for $\ell = 2^{20}$. When $t = 2^{16}$, the proof is longer and hence verification requires more time: it takes 9.1s for $\ell = 2^{13}$ and 131.1s for $\ell = 2^{18}$. Note that the server can divide client proofs in many small batches and fully parallelize the proof verification process.

End-to-end performance in FL. Following Lycklama et al. [33], we demonstrate the effectiveness of ACORN-detect for four practical use cases in federated learning. Concretely, we consider the following three tasks on image-based datasets: training (1) a convolutional neural network (CNN) on the Federated-MNIST dataset [12], (2) the LeNet-5 [30] and (3) the ResNet-20 CNN [23] on the CIFAR-10 dataset [29]. We also consider (4) the task of training an LSTM [24] model on the text-based Shakespeare dataset [12]. We use TensorFlow [1] to train neural networks, where the hyperparameters are set as in [33]. To account for network latency we introduce delays of 0.5ms as in [33]. We report the per-client running time that includes the local training and aggregation steps (as described in Section 1.2). The results are in Table 2, and a more detailed experimental setup can be found in Appendix A.

For each use case, we compare ACORN-detect with both L_∞- and L_2-norm based validity proofs against PRG-SecAgg without input validation. A single round on each client consists of receiving the updated model, performing local training using its samples, and participating in the secure aggregation protocol using its model updates as input. As we can see, the local training time dominates the running time of our SA baseline. Furthermore, in all cases the bandwidth overhead of ACORN-detect is very modest (at most 1.05x), and in most cases the computational overhead is also fairly low (at most 5.52x). The exception is for CIFAR-10 S, where the higher overhead (21.28x) is due to a relatively higher learning rate ($\eta = 0.01$) and the small sample count (1024) and epochs per round (2) required to reach the desired accuracy. This makes the local training highly efficient, and thus generating validity proofs becomes a more dominant cost.

7 Conclusion and Open Problems

We presented a new secure aggregation protocol, RLWE-SecAgg, along with extensions, ACORN, that allow the server to perform validity checks on the inputs provided by clients. Our benchmarks demonstrate that the overheads of these checks are practical. Other zero-knowledge protocols offer lower prover runtimes, however, and may do so without making other costs impractical for our setting. For example, lattice-based proofs may offer a better balance between computational and communication overheads, and would also offer the advantage when combined with RLWE-SecAgg of providing plausible post-quantum security.

Acknowledgements

We would like to thank Michael Specter for benchmarking our code on a Pixel device, and our anonymous reviewers and shepherd for their helpful comments and suggestions.

References

[1] M. Abadi et al. "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems". In: *CoRR* abs/1603.04467 (2016). arXiv: 1603.04467. URL: http://arxiv.org/abs/1603.04467.

[2] S. Addanki, K. Garbe, E. Jaffe, R. Ostrovsky, and A. Polychroniadou. "Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares". In: *SCN*. Vol. 13409. Lecture Notes in Computer Science. Springer, 2022, pp. 516–539.

[3] M. R. Albrecht, R. Player, and S. Scott. "On the concrete hardness of Learning with Errors". In: *J. Math. Cryptol.* 9.3 (2015), pp. 169–203.

[4] B. Alon, M. Naor, E. Omri, and U. Stemmer. "MPC for Tech Giants (GMPC): Enabling Gulliver and the Lilliputians to Cooperate Amicably". In: *IACR Cryptol. ePrint Arch.* (2022), p. 902.

[5] J. Bell et al. "ACORN: Input Validation for Secure Aggregation". In: *IACR Cryptol. ePrint Arch.* (2022). URL: https://eprint.iacr.org/2022/1461.

[6] J. H. Bell, K. A. Bonawitz, A. Gascón, T. Lepoint, and M. Raykova. "Secure Single-Server Aggregation with (Poly)Logarithmic Overhead". In: *CCS*. ACM, 2020, pp. 1253–1269.

[7] K. A. Bonawitz et al. "Practical Secure Aggregation for Privacy-Preserving Machine Learning". In: *CCS*. ACM, 2017, pp. 1175–1191.

[8] K. A. Bonawitz et al. "Towards Federated Learning at Scale: System Design". In: *MLSys*. mlsys.org, 2019.

[9] D. Boneh, E. Boyle, H. Corrigan-Gibbs, N. Gilboa, and Y. Ishai. "Lightweight Techniques for Private Heavy Hitters". In: *IEEE Symposium on Security and Privacy*. IEEE, 2021, pp. 762–776.

[10] J. Bootle, A. Cerulli, P. Chaidos, J. Groth, and C. Petit. "Efficient Zero-Knowledge Arguments for Arithmetic Circuits in the Discrete Log Setting". In: *Proceedings of Eurocrypt 2016*. 2016.

[11] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. "Bulletproofs: Short Proofs for Confidential Transactions and More". In: *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2018, pp. 315–334.

[12] S. Caldas et al. *Leaf: A benchmark for federated settings*. arXiv preprint arXiv:1812.01097. 2018.

[13] A. R. Choudhuri, A. Goel, M. Green, A. Jain, and G. Kaptchuk. "Fluid MPC: Secure Multiparty Computation with Dynamic Participants". In: *CRYPTO (2)*. Vol. 12826. Lecture Notes in Computer Science. Springer, 2021, pp. 94–123.

[14] A. R. Chowdhury, C. Guo, S. Jha, and L. van der Maaten. "EIFFeL: Ensuring Integrity for Federated Learning". In: *CCS*. ACM, 2022.

[15] H. Corrigan-Gibbs and D. Boneh. "Prio: Private, Robust, and Scalable Computation of Aggregate Statistics". In: *NSDI*. USENIX Association, 2017, pp. 259–282.

[16] *Efficient Bulletproofs in Rust*. Anonymized; available upon acceptance for double blinded review. 2022.

[17] A. Fiat and A. Shamir. "How to Prove Yourself: Practical Solutions to Identification and Signature Problems". In: *CRYPTO*. Vol. 263. Lecture Notes in Computer Science. Springer, 1986, pp. 186–194.

[18] C. Gentry. "A fully homomorphic encryption scheme". PhD thesis. Stanford University, USA, 2009. URL: https://searchworks.stanford.edu/view/8493082.

[19] C. Gentry, S. Halevi, and V. Lyubashevsky. "Practical Non-interactive Publicly Verifiable Secret Sharing with Thousands of Parties". In: *EUROCRYPT (1)*. Vol. 13275. Lecture Notes in Computer Science. Springer, 2022, pp. 458–487.

[20] Z. Ghodsi, M. Javaheripi, N. Sheybani, X. Zhang, K. Huang, and F. Koushanfar. *zPROBE: Zero Peek Robustness Checks for Federated Learning*. 2022. DOI: 10.48550/ARXIV.2206.12100. URL: https://arxiv.org/abs/2206.12100.

[21] O. Goldreich. *The Foundations of Cryptography - Volume 2: Basic Applications*. Cambridge University Press, 2004.

[22] J. Groth. "Non-interactive Zero-Knowledge Arguments for Voting". In: *International Conference on Applied Cryptography and Network Security (ACNS)*. Vol. 3531. Lecture Notes in Computer Science. 2005, pp. 467–482.

[23] K. He, X. Zhang, S. Ren, and J. Sun. *Deep Residual Learning for Image Recognition*. arXiv preprint arXiv:1512.03385. 2015.

[24] S. Hochreiter and J. Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (1997), pp. 1735–1780.

[25] *How Messages improves suggestions with federated technology*. https://support.google.com/messages/answer/9327902. Accessed: 2022-10-06.

[26] Y. Ishai, E. Kushilevitz, Y. Lindell, and E. Petrank. "On Combining Privacy with Guaranteed Output Delivery in Secure Multiparty Computation". In: *CRYPTO*. Vol. 4117. Lecture Notes in Computer Science. Springer, 2006, pp. 483–500.

[27] P. Kairouz et al. "Advances and Open Problems in Federated Learning". In: *Found. Trends Mach. Learn.* 14.1-2 (2021), pp. 1–210.

[28] F. Karakoç, M. Önen, and Z. Bilgin. "Secure Aggregation Against Malicious Users". In: *SACMAT '21: The 26th ACM Symposium on Access Control Models and Technologies, Virtual Event, Spain, June 16-18, 2021*. Ed. by J. Lobo, R. D. Pietro, O. Chowdhury, and H. Hu. ACM, 2021, pp. 115–124. DOI: 10.1145/3450569.3463572. URL: https://doi.org/10.1145/3450569.3463572.

[29] A. Krizhevsky. *Learning Multiple Layers of Features from Tiny Images*. 2009.

[30] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. "Gradient-based learning applied to document recognition". In: *Proceedings of the IEEE*. 1998.

[31] J. Lee, D. Kim, D. Kim, Y. Song, J. Shin, and J. H. Cheon. "Instant Privacy-Preserving Biometric Authentication for Hamming Distance". In: *IACR Cryptol. ePrint Arch.* (2018), p. 1214.

[32] Y. Lindell. "How to Simulate It - A Tutorial on the Simulation Proof Technique". In: *Tutorials on the Foundations of Cryptography*. Springer International Publishing, 2017, pp. 277–346.

[33] H. Lycklama, L. Burkhalter, A. Viand, N. Küchler, and A. Hithnawi. "RoFL: Attestable Robustness for Secure Federated Learning". In: *CoRR* abs/2107.03311 (2021).

[34] V. Lyubashevsky, N. K. Nguyen, and M. Plançon. "Lattice-Based Zero-Knowledge Proofs and Applications: Shorter, Simpler, and More General". In: Lecture Notes in Computer Science (2022).

[35] V. Lyubashevsky, C. Peikert, and O. Regev. "On Ideal Lattices and Learning with Errors over Rings". In: *J. ACM* 60.6 (2013), 43:1–43:35.

[36] L. Melis, G. Danezis, and E. D. Cristofaro. "Efficient Private Statistics with Succinct Sketches". In: *NDSS*. The Internet Society, 2016.

[37] D. Pasquini, D. Francati, and G. Ateniese. "Eluding Secure Aggregation in Federated Learning via Model Inconsistency". In: *CCS*. ACM, 2022, pp. 2429–2443.

[38] N. Pippenger. "On the Evaluation of Powers and Monomials". In: *SIAM J. Comput.* 9.2 (1980), pp. 230–250.

[39] *Predicting Text Selections with Federated Learning*. https://ai.googleblog.com/2021/11/predicting-text-selections-with.html. Accessed: 2022-10-06.

[40] M. O. Rabin and J. O. Shallit. "Randomized algorithms in number theory". In: *Communications on Pure and Applied Mathematics* 39.S1 (1986), S239–S256. DOI: https://doi.org/10.1002/cpa.3160390713. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpa.3160390713. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.3160390713.

[41] L. Reyzin, A. D. Smith, and S. Yakoubov. "Turning HATE Into LOVE: Homomorphic Ad Hoc Threshold Encryption for Scalable MPC". In: *IACR Cryptol. ePrint Arch.* 2018 (2018), p. 997.

[42] E. Roth, H. Zhang, A. Haeberlen, and B. C. Pierce. "Orchard: Differentially Private Analytics at Scale". In: *OSDI*. USENIX Association, 2020, pp. 1065–1081.

[43] V. Shejwalkar, A. Houmansadr, P. Kairouz, and D. Ramage. "Back to the Drawing Board: A Critical Evaluation of Poisoning Attacks on Federated Learning". In: (2022). To appear.

[44] E. Shi, T. H. Chan, E. G. Rieffel, R. Chow, and D. Song. "Privacy-Preserving Aggregation of Time-Series Data". In: *NDSS*. The Internet Society, 2011.

[45] J. So, B. Güler, and A. S. Avestimehr. "Turbo-Aggregate: Breaking the Quadratic Aggregation Barrier in Secure Federated Learning". In: *IEEE J. Sel. Areas Inf. Theory* 2.1 (2021), pp. 479–489.

[46] T. Stevens, C. Skalka, C. Vincent, J. Ring, S. Clark, and J. Near. "Efficient Differentially Private Secure Aggregation for Federated Learning via Hardness of Learning with Errors". In: *CoRR* abs/2112.06872 (2021). arXiv: 2112.06872. URL: https://arxiv.org/abs/2112.06872.

[47] Z. Sun, P. Kairouz, A. T. Suresh, and H. B. McMahan. "Can You Really Backdoor Federated Learning?" In: *CoRR* abs/1911.07963 (2019).

[48] H. Wang et al. "Attack of the Tails: Yes, You Really Can Backdoor Federated Learning". In: *NeurIPS*. 2020.

[49] C. Xie, M. Chen, P.-Y. Chen, and B. Li. "CRFL: Certifiably Robust Federated Learning against Backdoor Attacks". In: *Proceedings of the Thirty-eighth International Conference on Machine Learning (ICML)*. 2021.

A End-to-End Experiment Setup

Here we give more details about our end-to-end experiments on the four federated learning tasks in Section 6.2. Our experiments ran on a laptop with a quad-core Intel i7-1185G7 CPU and 16GB memory. Neural network training was performed using Python in the TensorFlow [1] framework without GPU acceleration. We used the same training parameters as Lycklama et al. [33] and simulated the same network latency (0.5ms). In particular, we used SDG as the training algorithm and used the same local learning rates η: $\eta = 0.05$ for MNIST, $\eta = 0.01$ for CIFAR-10 S, $\eta = 0.05$ for CIFAR-10 L, and $\eta = 0.3$ for Shakespeare. We set the number of epochs per round for MNIST to 5, for CIFAR-10 S and CIFAR-10 L to 2, and for Shakespeare to 1. During each epoch, a client runs the training algorithm on 1248 samples for MNIST, 1024 samples for CIFAR-10 S and L, and 69440 samples for Shakespeare; such training tasks are done in batches of size 32 for MNIST and 64 for others. Model parameter updates are converted to 8-bit fixed point values by applying 8-bit probabilistic quantization with 7 fractional bits. In each of our experiments, we ran the federated learning task with a full dataset using a corresponding secure aggregation protocol, and the per-client performance was measured using the average of five end-to-end runs.

Public parameters: Vector length ℓ, input domain $\mathbb{X}^\ell$, secret distribution χ_s, and seed expansion function $\mathsf{F}: \{0,1\}^\lambda \mapsto \mathsf{supp}(\chi_s)^\ell$
Client i's input: $\mathbf{x}_i \in \mathbb{X}^\ell$
Server output: $z \in \mathbb{X}$

Commitments

1. Client i generates keypairs $(\mathsf{sk}_{i,1}, \mathsf{pk}_{i,1}), (\mathsf{sk}_{i,2}, \mathsf{pk}_{i,2}) \leftarrow \mathsf{Sig.KeyGen}(1^\lambda)$ and sends $(\mathsf{pk}_{i,1}, \mathsf{pk}_{i,2})$ to the server. It performs the first step in the distributed key correctness protocol, which results in it sending a message h_i to the server.
2. The server commits to the public key vectors $\mathsf{pk}_1 = (\mathsf{pk}_{i,1})_i$ and $\mathsf{pk}_2 = (\mathsf{pk}_{i,2})_i$ using a Merkle tree. It sends the root hashes $h_{\mathsf{root},1}$ and $h_{\mathsf{root},2}$ to each client. It also performs the second step of the distributed key correctness protocol, which means sampling its random challenge e and forming and sending $\mathsf{com}_{i,\mathsf{chl}}$ to client i.

Distributed graph generation

3. Client i selects k neighbors by sampling randomly and without replacement k times from the set of n clients, and sends the resulting set $\mathsf{N}_\rightarrow(i)$ of outgoing neighbors to the server. Denote by $\mathsf{N}(i)$ all neighbors of client i (consisting of their outgoing edges and implicitly defined incoming edges).
4. The server sends $\mathsf{N}_\leftarrow(i), (j, \mathsf{pk}_{j,1}, \pi_{j,1}, \mathsf{pk}_{j,2}, \pi_{j,2})_{j \in \mathsf{N}(i)}$ to client i, where $\pi_{j,1}$ and $\pi_{j,2}$ are Merkle inclusion proofs with respect to roots $h_{\mathsf{root},1}$ and $h_{\mathsf{root},2}$.
5. Client i aborts if the server has sent more than $3k+k$ keys, if there is an index $j \in \mathsf{N}_\rightarrow(i)$ that is not reflected in the keys sent by the server, or if the Merkle inclusion proofs fail to verify.

Seed sharing

6. Each client i that has not dropped out performs the following:
   - Generates a random seed $\mathbf{seed}_i$.
   - Computes two sets of shares $H_i^{\mathbf{seed}} = \{h_{i,1}^{\mathbf{seed}}, \ldots, h_{i,k}^{\mathbf{seed}}\} = \mathtt{ShamirSS}(t, k, \mathbf{seed}_i)$ and $H_i^s = \{h_{i,1}^s, \ldots, h_{i,k}^s\} = \mathtt{ShamirSS}(t, k, sk_{i,1})$.
   - Sends to the server messages $m_j = (j, c_{i,j})$ for each $j \in \mathsf{N}_\rightarrow(i)$, where $c_{i,j} \leftarrow \mathsf{E}_{\mathsf{auth}}.\mathsf{Enc}(k_{i,j}, i \| j \| h_{i,j}^{\mathbf{seed}} \| h_{i,j}^s)$ for $k_{i,j} = \mathsf{KA.Agree}(\mathsf{sk}_{i,2}, \mathsf{pk}_{j,2})$.
7. If the server receives messages from fewer than $(1-\delta)n$ clients, it aborts. Otherwise, it sends all messages $(j, c_{i,j})$ to client j. Denote by $A_j \subseteq \mathsf{N}(j)$ the set of neighbors for whom client j received such a message.

Masking

8. Each client i that has not dropped out performs the following:
   - Computes a shared random seed $\mathbf{seed}_{i,j}$ as $\mathbf{seed}_{i,j} = \mathsf{KA.Agree}(sk_{i,1}, pk_{j,1})$.
   - Computes its *packed encrypted input* $\mathbf{y}_i = \mathsf{Encode}(\mathbf{sk}_i, \mathbf{Gx}_i)$ with key defined as $\mathbf{sk}_i = \mathbf{s}_i + \sum_{j \in A_i, j<i} \mathbf{s}_{ij} - \sum_{j \in A_1, i<j} \mathbf{s}_{ij}$ for $\mathbf{s}_{ij} = \mathsf{F}(\mathbf{seed}_{i,j})$, $\mathbf{s}_i = \mathsf{F}(\mathbf{seed}_i)$ (as in Equation 3).
   - Forms $\sigma_{i,j}^{\mathsf{incl}} \leftarrow \mathsf{Sig.Sign}(\mathsf{sk}_{i,2}, m_{i,j} = \text{“included”} \| i \| j)$ for all $j \in A_i$.
   - Forms commitments $\mathsf{com}_{\mathbf{sk},i} \leftarrow \mathsf{Commit}(\mathbf{sk}_i; r_i)$ and $\mathsf{com}_{\mathbf{x},i} \leftarrow \mathsf{Commit}(\mathbf{x}_i)$ to its key and input respectively.
   - Computes proofs $\pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}$, $\pi^{0 \le \mathbf{x}_i < t}$, and $\pi^{\mathsf{valid}(\mathbf{x}_i)}$ of encoding, smallness, and validity.
   - Performs the third step of the distributed key correctness protocol to form K_i.
   - Sends to the server $\mathbf{y}_i$, $(m_{i,j}, \sigma_{i,j}^{\mathsf{incl}})_j$, $\mathsf{com}_{\mathbf{sk},i}, \mathsf{com}_{\mathbf{x},i}, K_i, \pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}, \pi^{0 \le \mathbf{x}_i < t}, \pi^{\mathsf{valid}(\mathbf{x}_i)}$.

Dropout agreement and unmasking

9. The server collects packed encoded inputs for a determined time period. If it receives fewer than $(1-\delta)n$, it aborts. Otherwise, it defines a global set of dropouts $\mathcal{D}$ and a set of survivors $\mathcal{S}$. It then sends the messages and signatures $(m_{j,i}, \sigma_{j,i}^{\mathsf{incl}})$ to every client $i \in \mathcal{S}$, along with the sets $\mathcal{D}_i = \mathsf{N}(i) \cap \mathcal{D}$ (its incoming neighbors that are dropouts) and $\mathcal{S}_i = \mathsf{N}(i) \cap \mathcal{S}$ (its incoming neighbors that are not). It also sends the opening of the commitment for the distributed key correctness protocol (following the fourth step), containing its challenge e.
10. Each client i that has not dropped out performs the following:
    - Checks that $\mathcal{D}_i \cap \mathcal{S}_i = \emptyset$, that $\mathcal{S}_i, \mathcal{D}_i \subseteq \mathsf{N}(i) \cap A_i$, and that all signatures $\sigma_{j,i}^{\mathsf{incl}}$ are valid on message $m_{j,i}$ for all $j \in \mathcal{S}_i$, aborting if any of these checks fail.
    - Computes $\sigma_{i,j}^{\mathsf{ack}} \leftarrow \mathsf{Sign}(\mathsf{sk}_{i,2}, \text{“ack”} \| i \| j)$ for all $j \in \mathcal{S}_i$.
    - Performs the fifth step of the distributed key correctness protocol, which means forming values t_i and α_i.
    - Sends $(m_{i,j}, \sigma_{i,j}^{\mathsf{ack}})_j$ and t_i, α_i to the server.
11. The server aborts if it receives fewer than $(1-\delta)n$ responses. It verifies all received proofs $\pi^{\mathsf{Enc}(\mathbf{sk}_i, \mathbf{x}_i)}$, $\pi^{0 \le \mathbf{x}_i < t}$, $\pi^{\mathsf{valid}(\mathbf{x}_i)}$ and aborts if any of them fails. Otherwise it forwards all messages $(j, m_{i,j}, \sigma_{i,j}^{\mathsf{ack}})$ to client j.
12. Each remaining client verifies its received signatures using $\mathsf{pk}_{j,2}$, aborting if they fail to verify. Once a client receives p valid signatures from its neighbors, it sends $\{i, h_{i,j}^{\mathbf{seed}}\}_{j \in \mathcal{D}_i}$ and $\{i, h_{i,j}^s\}_{j \in \mathcal{S}_i}$ to the server, which it has obtained by decrypting the ciphertexts $c_{i,j}$ received in Step 6.
13. The server aborts if it receives fewer than $(1-\delta)n$ responses, and otherwise:
    - Collects, for each client $i \in \mathcal{D}$, the set of all received shares in $H_i^{\mathbf{seed}}$, and aborts if there are fewer than t. If not it recovers $\mathbf{seed}_i$ and $\mathbf{s}_i$ using the t shares received from the lowest client IDs.
    - Collects, for each client $i \in \mathcal{S}$, the set of all shares in H_i^s, and aborts if there are fewer than t. If not it recovers $sk_{i,1}$ and $\mathbf{s}_{ij}$ for all $j \in \mathsf{N}(i)$.
    - Computes a decryption key $\mathbf{sk} = \sum_{i \in \mathcal{S}} (\mathbf{s}_i + \sum_{j \in \mathcal{D}_i, j<i} \mathbf{s}_{ij} - \sum_{j \in \mathcal{D}_i, i<j} \mathbf{s}_{ij})$.
    - Using $\mathbf{sk}$, performs the final step of the distributed key correctness protocol and aborts if verification fails.
    - Outputs $\sum_{i \in \mathcal{S}} \mathbf{x}_i$ as $\mathbf{G}^{-1}(\mathsf{Decode}(\mathbf{sk}, \sum_{i \in A_2'} \mathbf{y}_i))$.

Figure 6: Maliciously secure SecAgg from homomorphic encodings with input verification.