

λ -SECAGG: PARTIAL VECTOR FREEZING FOR LIGHTWEIGHT SECURE AGGREGATION IN FEDERATED LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Secure aggregation of user update vectors (e.g. gradients) has become a critical issue in the field of federated learning. Many Secure Aggregation Protocols (SAPs) face exorbitant computation costs, severely constraining their applicability. Given the observation that a considerable portion of SAP’s computation burden stems from processing each entry in the private vectors, we propose **Partial Vector Freezing (PVF)**, a portable module for compressing computation costs without introducing additional communication overhead. **λ -SecAgg**, which integrates SAP with PVF, “freezes” a substantial portion of the private vector through specific transformations, requiring only $\frac{1}{\lambda}$ of the original vector to participate in SAP. Eventually, users can “thaw” the public sum of the “frozen entries” by the result of SAP. To avoid potential privacy leakage, we devise **Disrupting Variables Element** for PVF. We demonstrate that PVF can seamlessly integrate with various SAPs and it poses no threat to user privacy in the semi-honest and active adversary settings. We include 7 baselines, encompassing 5 distinct types of masking schemes, and explore the acceleration effects of PVF on these SAPs. Empirical investigations indicate that when $\lambda = 100$, PVF yields up to $99.5\times$ speedup and up to $32.3\times$ communication reduction.

1 INTRODUCTION

Machine learning technologies are applied in countless fields to improve service performance. However, aggregating large amounts of data for big data mining raises concerns regarding data privacy (Liu et al., 2021). *Federated Learning* (FL) (McMahan et al., 2017) keeps original data on the local devices while only requiring data owners to submit local training updates to a central server. Nonetheless, as Zhu et al. (2019) and Geiping et al. (2020) indicate, attackers can infer a user’s local data by reversing the submitted updates. To address this issue, numerous research efforts have been focusing on *Secure Aggregation Protocols* (SAPs) (Liu et al., 2022b) for aggregating all user’s model information while preserving individual privacy.

The widely discussed SAPs are based on *Secure Multi-Party Computation* (SMPC) (Xu et al., 2022; Sotthiwat et al., 2021), *Mask* (Bonawitz et al., 2017), *Homomorphic Encryption* (HE) (Aono et al., 2017) and *Differential Privacy* (DP) (Wei et al., 2020). For most SAPs, except for DP-based ones, the computation overhead often scales proportionally with the length of the model update vectors since most of these schemes involve masking (encrypting) each entry of the vector sequentially. Therefore the **computation** time for both masking and unmasking always experiences a steep escalation with the increase in vector length, as *PracAgg* (Bonawitz et al., 2017) in Figure 1, significantly constraining real-world applications. Especially in recent applications that utilize FL to fine tune *Large Language Models* (LLMs) (Ye et al., 2024) with billions of parameters, the computational overhead brought by SAP is unbearable.

On the other hand, DP-based solutions although have the best efficiencies, many studies (Stevens et al., 2022; Wang et al., 2021) posit that the minimal noise added by DP is insufficient to thwart attacks such as gradient inversion aimed at stealing users’ local data, where adversaries can recover flawed but recognizable handwritten digit image (Wang et al., 2021). Therefore the security of DP in

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107




Figure 1: Computation time for one round, with 100 users and 10% dropout rate.




Figure 2: Comparison of various aggregation methods, with $n = 100, \eta = 5\%, \lambda = 100$.




Figure 3: Comparison with compression-based techniques, with $n = 100, \eta = 5\%$.

secure aggregation faces challenges, necessitating its combination with masking to bolster privacy. In this work, we primarily explore methods to alleviate **masking-related** overhead.

Our proposal. We think the root cause of computation overhead in SAPs is masking each entry of the original vector. While a few *sparsification*-related approaches (Ergun et al., 2021; Lu et al., 2023) try to reduce the dimensions of uploaded vectors, they raise an inevitable trade-off of discarding some information. In this work, we propose **Partial Vector Freezing (PVF)** to reduce the number of entries processed in SAPs while ensuring intact aggregation of all entries in the original vector. Within the module, each user performs certain transformations on the original vector at a negligible computation cost to selectively *freeze* most entries of the user’s original vector, compressing the length of the vector involved in SAPs to $\frac{1}{\lambda}$ of its original size (λ is the compression factor). The communication overhead and number of communication interactions in SAPs after integrating PVF, which we call **λ -SecAgg**, do not increase. Further, we propose *Disrupting Variables Element* to prevent PVF from leaking the linear relationship between vectors. PVF remains decoupled from SAPs and guarantees individual user privacy under semi-honest and active adversary settings, offering high portability.

Our contributions can be summarized as follows:

- We propose the PVF without incurring additional communication overhead or harming security. It reduces the entries processed in SAP while **ensuring the aggregation of all entries in the original vector**, which means it can compress the computation overhead of SAP to approximately $\frac{1}{\lambda}$ of the original. Moreover, it brings up to $32.3 \times (\lambda = 100)$ additional communication enhancements for HE-based SAPs thanks to the decreased number of ciphertext entries.
- We propose the disrupting variables element to PVF to avoid potential privacy leakage.
- Extensive experiments show the effectiveness of our proposal. We include 7 baselines encompassing 5 types of **masking schemes** for a comprehensive overhead comparison, which is largely unexplored in most research endeavors and reaffirms the high portability of PVF.

2 RELATED WORK

Secure Aggregation Protocols. Various types of SAPs have been proposed, including SMPC-based (Boer & Kramer, 2020; Xu et al., 2022), HE-based (Aono et al., 2017; Ma et al., 2022; Li et al., 2022), DP-based (Geyer et al., 2017; Wei et al., 2020), and Mask-based (Bonawitz et al., 2017) schemes. Most efforts to reduce computation cost focus on enhancing PracAgg (Bonawitz et al., 2017), which is mainly categorized into two types: (i) improving the masking mechanism (Liu et al., 2022a; Stevens et al., 2022; Liu et al., 2022c; Wei et al., 2023) to reduce **masking-related** overhead; (ii) minimizing **interaction-related** overhead, including refining communication structures (Bell et al., 2020; So et al., 2021) and enhancing efficiency in key agreements among users (Kalikinkar et al., 2018; Kadhe et al., 2020; Ma et al., 2023). **Note** that the security of FL remains an open issue. SAPs, though cannot fully guarantee FL security at the moment, remain a promising direction worth exploration. The main objective of our work is to reduce the masking-related overhead of secure aggregation, thereby making it more applicable in practice.

Compression-based Techniques. Rothchild et al. (2020) employs a *Count Sketch* to compress model updates. Additionally, some sparsification-based approaches (Ergun et al., 2021; Lu et al., 2023) can reduce vector dimensions. Our method fundamentally differs from these schemes, as our proposal compresses the entries involved in secure aggregation while retaining the **intact** aggregation result of all original entries.

Defense Against Malicious Server. Several works indicate the malicious server can launch *Model Inconsistency Attacks* (Pasquini et al., 2022), *Multi-round Privacy Stealing Attacks* (So et al., 2023) and *Aggregation Falsification Attacks* (Guo et al., 2021). These studies also propose strategies to counter these attacks accordingly, only requiring minor modifications to the SAP process, as described in Section 3.4.

Input constraints. Several works (Bell et al., 2023; Lycklama et al., 2023) are proposed to mitigate *Poisoning Attacks* in FL. They delineate that the erroneous inputs of malicious users can result in the server obtaining an inaccurate global model, thereby harming the training task. They propose methodologies utilizing *Zero-Knowledge Proofs* to bound user inputs. However, their ability to prevent poisoning attacks is limited (Ma et al., 2023). Establishing strong constraints against malicious inputs remains an unresolved challenge, and it falls beyond the scope of this work. What’s more, Mozaffari et al. (2023) propose *Federated Rank Learning* (FRL), where the server aggregates the parameter rankings instead of the model parameter updates. It can effectively resist poisoning attacks, and enable direct aggregations without any constraints on user submissions. Therefore, FRL can be combined with SAP, and we do not have to worry about whether PVF can be integrated with input constraints in this work.

3 PARTIAL VECTOR FREEZING

Scenario. In the t -th round of FL, the user set $\mathcal{U} = \{u_1, \dots, u_n\}$ conduct local model training and submit model updates $\{\mathbf{x}^{i(t)}\}_{i \in \mathcal{U}} = \{(x_1^{i(t)}; \dots; x_m^{i(t)})\}_{i \in \mathcal{U}}$ (*Original Vectors*) to the server \mathcal{S} . There might be η ($\leq 30\%$) users that drop out during the aggregation due to network instability or other reasons and \mathcal{U}' is the surviving user set. \mathcal{S} aggregates the model updates to compute $\sum_{i \in \mathcal{U}'} \mathbf{x}^{i(t)}$ and redistributes the result to all users (*Plain Aggregation*). This iterative process continues until the completion of model training. SAPs can help obtain $\sum_{i \in \mathcal{U}'} \mathbf{x}^{i(t)}$ while ensuring the privacy of each individual $\mathbf{x}^{i(t)}$. Similar to other SAPs (Bonawitz et al., 2017; Stevens et al., 2022), we define the elements of $\mathbf{x}^{i(t)}$ ($i \in [1, n]$) within \mathbb{Z}_p for some large public prime p and assume there is a secure communication channel between each user and \mathcal{S} . In this section, our emphasis lies in introducing the computation methodology of PVF, specifically considering a single-round aggregation process with superscript “ (t) ” omitted. To ensure multi-round privacy, employing a specialized user selection mechanism for each round is sufficient (Liu et al., 2023; So et al., 2023). For the summary of notations, please refer to Appendix A.

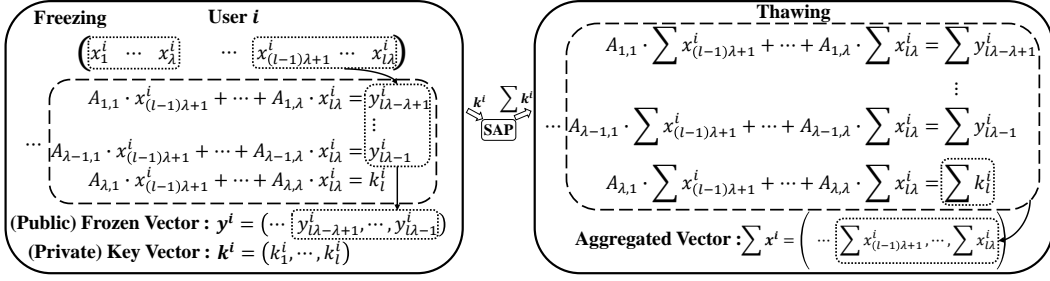
Threat model: Corrupt participants endeavor to infer the privacy of honest parties based on the messages they receive, i.e., the *Semi-honest Model*, and can fabricate messages, i.e., the *Active Adversary Model*. We assume malicious users do not exceed one-third of the total users, aligned with PracAgg (Bonawitz et al., 2017) and EffiAgg (Liu et al., 2022a). The lenient security assumptions of PVF allow its easy integration with secure aggregation protocols. When integrating with a SAP, the security assumptions originally employed in the corresponding protocol are adopted and we assume the integrated SAP can **reliably ensure the privacy of inputs**.

3.1 MOTIVATION

Within SAP, every minor operation on an entry accumulates m times, ultimately imposing significant computational burdens on devices. For example, u_i in PracAgg needs to perform the following calculations on each entry x_j^i of \mathbf{x}^i (b and s are user secret keys, PRG is a pseudorandom number generator):

$$y_j^i = x_j^i + \text{PRG}(b_i) + \sum_{h \in \mathcal{U}: i < h} \text{PRG}(s_{i,h}) - \sum_{h \in \mathcal{U}: i > h} \text{PRG}(s_{h,i}). \quad (1)$$

Based on these observations, we try to reduce the number of entries processed in secure aggregation while ensuring users receive all entries of the aggregated vector. To accomplish this objective, our

Figure 4: Workflow of λ -SecAgg with Main PVF.

mindset is to devise a module that can freeze certain entries (*Frozen Entries*) of the user’s original vector, and only perform secure aggregation on the other entries (*SecAgg Entries*). Upon SAP completion, this module thaws the frozen entries, ensuring that their privacy is well protected throughout the entire process.

Definition 1. Given an invertible $\lambda \times \lambda$ matrix \mathbf{A} and a segment of original vector $\mathbf{x} = (x_1; \dots; x_\lambda)$, we define Incomplete Matrix $\check{\mathbf{A}} = \mathbf{A}_{:\lambda-1,:}$ and Residual Vector $\boldsymbol{\alpha} = \mathbf{A}_{\lambda,:}$. The function for finding the solution set of a system of linear equations is:

$$SLE_{AK}(\mathbf{A}\mathbf{x}) \rightarrow \mathbf{x}, \quad (2)$$

where AK denotes the additional knowledge. We use $\text{rank}(\cdot)$ represents the rank of a matrix. Since $\text{rank}(\mathbf{A}) = \text{rank}(\mathbf{A}, \mathbf{A}\mathbf{x}) = \lambda$, the system of linear equations has a unique solution \mathbf{x} . However, due to $\text{rank}(\check{\mathbf{A}}) = \text{rank}(\check{\mathbf{A}}, \check{\mathbf{A}}\mathbf{x}) < \lambda$, the system has an infinitude of solutions (Suetin et al., 1989), also called an Under-determined System of Linear Equations.

It can be seen that when $\check{\mathbf{A}}$ and $\check{\mathbf{A}}\mathbf{x}$ are known, in the absence of knowledge about $\boldsymbol{\alpha}\mathbf{x}$, \mathbf{x} presents an infinite set of possibilities, rendering it impossible to determine that specific confidential vector. Motivated by this, we propose PVF.

3.2 MAIN METHOD

In this section, we present the computation process of the Main PVF module during a single aggregation round, depicted by the workflow shown in Figure 4.

Phase 0: Main.Setup(\cdot). Generate an invertible matrix $\mathbf{A} \in \mathbb{Z}_p^{\lambda \times \lambda}$ randomly, and obtain \mathbf{A}^{-1} , $\check{\mathbf{A}}$, $\boldsymbol{\alpha}$, which are all **public** parameters. We refer to λ as the *Compression Factor*.

Phase 1: Main.Freeze(\cdot). For $i \in [1, n]$, randomly pad \mathbf{x}^i to ensure the length of the padded vector is a multiple of λ . The padded vector is $\mathbf{x}_{pad}^i = (x_1^i; x_2^i; \dots; x_{m'}^i)$, where $m' = l\lambda$. Then divide \mathbf{x}_{pad}^i into l groups: $\mathbf{d}_j^i = (d_{(j-1)\lambda+1}^i; d_{(j-1)\lambda+2}^i; \dots; d_{j\lambda}^i)$. Use the incomplete matrix $\check{\mathbf{A}}$ to compute *Frozen Vector*:

$$\begin{aligned} \mathbf{y}^i &= (\mathbf{y}_1^i, \dots, \mathbf{y}_l^i) = \left((y_1^i, \dots, y_{\lambda-1}^i), \dots, (y_{(l-1)\lambda+1}^i, \dots, y_{(l-1)\lambda+\lambda-1}^i) \right) \\ &= (\check{\mathbf{A}}\mathbf{d}_1^i, \dots, \check{\mathbf{A}}\mathbf{d}_l^i), \end{aligned} \quad (3)$$

and use the residual vector to compute *Key Vector*:

$$\mathbf{k}^i = (k_1^i, \dots, k_l^i) = (\boldsymbol{\alpha}\mathbf{d}_1^i, \dots, \boldsymbol{\alpha}\mathbf{d}_l^i). \quad (4)$$

Phase 2: Main.SecAgg(\cdot). Users and \mathcal{S} execute SAP, where the vector to be aggregated of user i is $\mathbf{k}^i = (k_1^i, k_2^i, \dots, k_l^i) \in \mathbb{Z}_p^l$. We require all users to send their respective \mathbf{y}^i to \mathcal{S} during SAP, thus **eliminating the need for additional interactions**. Upon completion of SAP, the surviving user set \mathcal{U}' and \mathcal{S} can receive the aggregated result of the key vectors: $\sum_{i \in \mathcal{U}'} \mathbf{k}^i$ or $\text{Enc}(\sum_{i \in \mathcal{U}'} \mathbf{k}^i)$ (in HE-based SAPs). Then \mathcal{S} computes $\sum_{i \in \mathcal{U}'} \mathbf{y}^i$, immune to the impact of dropout users.

Phase 3: Main.Thaw(\cdot). Thawing can be executed either at the server or user side, without compromising privacy and incurring any additional communication:

(1) *Thawing on the server side.* Based on the acquired $\sum_{i \in \mathcal{U}'} \mathbf{k}^i$ and $\sum_{i \in \mathcal{U}'} \mathbf{y}^i$, \mathcal{S} can derive the aggregated result \mathbf{sum} . The correctness stems from the linearity of \mathbf{A} and α :

$$\begin{aligned} \mathbf{z} &= \left(\left(\sum_{i \in \mathcal{U}'} \mathbf{y}_1^i, \sum_{i \in \mathcal{U}'} k_1^i \right), \dots, \left(\sum_{i \in \mathcal{U}'} \mathbf{y}_l^i, \sum_{i \in \mathcal{U}'} k_l^i \right) \right) = \left(\left(\sum_{i \in \mathcal{U}'} \mathbf{A} \mathbf{d}_1^i, \sum_{i \in \mathcal{U}'} k_1^i \right), \dots, \left(\sum_{i \in \mathcal{U}'} \mathbf{A} \mathbf{d}_l^i, \sum_{i \in \mathcal{U}'} k_l^i \right) \right) \\ &= \left(\sum_{i \in \mathcal{U}'} \mathbf{A} \mathbf{d}_1^i, \dots, \sum_{i \in \mathcal{U}'} \mathbf{A} \mathbf{d}_l^i \right) = \left(\mathbf{A} \sum_{i \in \mathcal{U}'} \mathbf{d}_1^i, \dots, \mathbf{A} \sum_{i \in \mathcal{U}'} \mathbf{d}_l^i \right) \end{aligned} \quad (5)$$

Since \mathbf{A}^{-1} and \mathbf{z} are public, \mathcal{S} can thaw frozen vectors and compute the aggregated results of all entries in the original vector by:

$$\begin{aligned} \sum_{i \in \mathcal{U}'} \mathbf{d}_1^i &= \mathbf{A}^{-1} \mathbf{z}_1 = \mathbf{A}^{-1} \mathbf{A} \sum_{i \in \mathcal{U}'} (x_1^i; x_2^i; \dots; x_\lambda^i), \\ &\vdots \\ \sum_{i \in \mathcal{U}'} \mathbf{d}_l^i &= \mathbf{A}^{-1} \mathbf{z}_l = \mathbf{A}^{-1} \mathbf{A} \sum_{i \in \mathcal{U}'} (x_{(l-1)\lambda+1}^i; x_{(l-1)\lambda+2}^i; \dots; x_{l\lambda}^i). \end{aligned} \quad (6)$$

At this point, \mathcal{S} completes the thawing phase and obtain:

$$\mathbf{sum} = \left(\sum_{i \in \mathcal{U}'} \mathbf{d}_1^i; \dots; \sum_{i \in \mathcal{U}'} \mathbf{d}_l^i \right) = \left(\sum_{i \in \mathcal{U}'} x_1^i; \dots; \sum_{i \in \mathcal{U}'} x_{l\lambda}^i \right) = \sum_{i \in \mathcal{U}'} \mathbf{x}_{pad}^i. \quad (7)$$

Subsequently, it transmits \mathbf{sum} to all online users. Upon receiving \mathbf{sum} , users can obtain the final aggregated result by removing the padding.

(2) *Thawing on the user side.* In certain SAPs (Aono et al., 2017; Xu et al., 2022), the aggregated results remain invisible to \mathcal{S} to ensure the protection of users' intellectual property, among other goals. In such situations, \mathcal{S} cannot perform the thawing. Instead, \mathcal{S} sends $\sum_{i \in \mathcal{U}'} \mathbf{y}^i$ and $Enc(\sum_{i \in \mathcal{U}'} \mathbf{k}^i)$ to all surviving users. Users locally decrypt $Enc(\sum_{i \in \mathcal{U}'} \mathbf{k}^i)$ and perform the thawing process to obtain \mathbf{sum} . By removing the padding, users attain the final aggregated result.

3.3 DISRUPTING VARIABLES ELEMENT

While the server cannot obtain any individual element of \mathbf{x}^i within Main PVF, it still obtains certain **linear relationships** involving the private entries. In this section, we present improvements to the Main PVF to ensure that the server cannot obtain any information about \mathbf{x}^i from \mathbf{y}^i .

Firstly, we improve the process of generating \mathbf{y}_j (ignoring the superscript “ i ” for clarity) in Equation (3) as follows:

$$\begin{cases} A_{1,1}(x_{(j-1)\lambda+1} + k_1 + \dots + k_{\lfloor \frac{l}{\lambda} \rfloor}) + \dots + A_{1,\lambda}(x_{j\lambda} + k_{l - \lfloor \frac{l}{\lambda} \rfloor} + \dots + k_l) = y_{(j-1)\lambda+1} \\ \vdots \\ A_{\lambda-1,1}(x_{(j-1)\lambda+1} + k_1 + \dots + k_{\lfloor \frac{l}{\lambda} \rfloor}) + \dots + A_{\lambda-1,\lambda}(x_{j\lambda} + k_{l - \lfloor \frac{l}{\lambda} \rfloor} + \dots + k_l) = y_{j\lambda} \end{cases}, \quad (8)$$

where $j \in [1, l]$ and \mathbf{k} is the same as in Equation (4). Similar to Equation (5), upon thawing, the following can be derived:

$$\begin{cases} A_{1,1}(\sum x_{(j-1)\lambda+1}) + \dots + A_{1,\lambda}(\sum x_{j\lambda}) = \sum y_{(j-1)\lambda+1} - \frac{\sum_{o \in [1, \lambda]} A_{1,o} \sum_{r \in [(o-1)\lfloor \frac{l}{\lambda} \rfloor + 1, o\lfloor \frac{l}{\lambda} \rfloor]} k_r}{\sum_{o \in [1, \lambda]} A_{1,o} \sum_{r \in [(o-1)\lfloor \frac{l}{\lambda} \rfloor + 1, o\lfloor \frac{l}{\lambda} \rfloor]} k_r} \\ \vdots \\ A_{\lambda-1,1}(\sum x_{(j-1)\lambda+1}) + \dots + A_{\lambda-1,\lambda}(\sum x_{j\lambda}) = \sum y_{j\lambda-1} - \frac{\sum_{o \in [1, \lambda]} A_{\lambda-1,o} \sum_{r \in [(o-1)\lfloor \frac{l}{\lambda} \rfloor + 1, o\lfloor \frac{l}{\lambda} \rfloor]} k_r}{\sum_{o \in [1, \lambda]} A_{\lambda-1,o} \sum_{r \in [(o-1)\lfloor \frac{l}{\lambda} \rfloor + 1, o\lfloor \frac{l}{\lambda} \rfloor]} k_r} \\ A_{\lambda,1}(\sum x_{(j-1)\lambda+1}) + \dots + A_{\lambda,\lambda}(\sum x_{j\lambda}) = \sum k_j \end{cases}, \quad (9)$$

where the right side can be obtained given $\sum \mathbf{k}^i$ is known. Then, $\sum \mathbf{x}^i$ can be determined and the thawing phase is successfully completed. This improvement aims at **complicating the relationships among entries**. It can be seen that this additional operation only adds some vector-vector additions, which brings negligible computational overhead.

Secondly, similar to many hybrid schemes combining mask and DP (Bonawitz et al., 2017; Stevens et al., 2022; Liu et al., 2022c) (or encryption and DP (Wang et al., 2021)), DVE adds noise to \mathbf{x}^i before aggregation to enhance privacy by:

$$\mathbf{x}^i = \mathbf{x}^i + \mathbf{e}, \quad (10)$$

where the Gaussian noise $e \sim N(0, \sigma^2)$. And $\mathbf{y}^i = \check{\mathbf{A}}(\mathbf{x}^i + e) = \check{\mathbf{A}}\mathbf{x}^i + e'$ ($\check{\mathbf{A}}$ is public). Therefore, given a uniformly random vector \mathbf{w}^i , Lemma 3 in Appendix D.3 ensures that $(\check{\mathbf{A}}, \mathbf{y}^i)$ and $(\check{\mathbf{A}}, \mathbf{w}^i)$ are indistinguishable, which guarantees **\mathcal{S} does not obtain private information from honest users through frozen vectors**. For clarity and conciseness, we do not differentiate the symbols of the original vector before and after adding noise.

3.4 INTEGRATING PVF WITH DIFFERENT SAPS

To resist **model inconsistency attacks**, appending the hash of the received model to the pseudo-random generator seed is sufficient, without incurring additional overhead (Ma et al., 2023). And utilizing an innovative user selection mechanism (So et al., 2023; Liu et al., 2023) is able to achieve **multi-round privacy**. To resist **aggregation falsification attacks**, verifiable protocols (Hahn et al., 2023) are able to verify the aggregation results through commitments sent by \mathcal{S} , and we provide *Result Verification Extension* in Appendix B.2 to enable PVF to integrate with such SAP. The pipeline of λ -SecAgg is shown in Figure 13. For the detailed portability analysis, please refer to Appendix C.

4 SECURITY ANALYSIS

Evidently, the information that adversaries can obtain about an honest participant only includes *sum* and under-determined systems of linear equations (\mathbf{y}^i). Randomly generating \mathbf{A} and performing certain **pre-checks** (as Section 4.1), the under-determined systems of linear equations can effectively preserve the privacy of each entry, which is also adopted by Liu et al. (2023). And in Section 4.2, we demonstrate adversaries cannot access \mathbf{x}^i throughout λ -SecAgg process.

4.1 PRIVACY OF EACH ELEMENT

In cases of improper selection of \mathbf{A} , \mathcal{S} can access the privacy of some specific elements within an original vector (if without DVE), as illustrated in Example D.1. In the implementation, we initially generate \mathbf{A} randomly, and we transform $\check{\mathbf{A}}$ into *Reduced Row Echelon Form* and verify that no element can be deduced. This ensures privacy of every element. Unless specified otherwise, all subsequent references to \mathbf{A} in the following text are designed to guarantee each element privacy.

4.2 SECURITY ANALYSIS OF ENTIRE VECTORS

Protocol security requires that adversaries can not obtain private information of any **individual** honest participant. The view of a participant consists of its internal information and received messages. Given a SAP that can maintain security in the active adversary setting, Theorem 1 guarantees its security when integrated with PVF. Theorem 1 demonstrates the information revealed during a real execution is **indistinguishable** from that obtained through a random simulation. The proof of Theorem 1 is carried out in a *Random Oracle model*. In this model, we define a trapdoor function to inform *SIM* of the **sum** of existing honest users' private information. During one execution of the protocol, *SIM* can only access it once to obtain necessary information. Let \mathcal{C} denote the set of malicious participants, which is a subset of $\mathcal{U} \cup \{\mathcal{S}\}$. The ideal function is defined as follows:

$$Ideal_{\{\mathbf{x}^i\}_{i \in \mathcal{U} \setminus \mathcal{C}}}(L) = \begin{cases} \sum_{i \in L} \mathbf{x}^i, & L \subseteq (\mathcal{U} \setminus \mathcal{C}) \text{ and } |L| > \left\lceil \frac{n}{3} \right\rceil \\ \perp, & \text{otherwise} \end{cases}. \quad (11)$$

Theorem 1 (Security against malicious participants). *Let $REAL_{\mathcal{C}}^{\mathcal{U}, \lambda}(\{\mathbf{x}^i\}_{i \in \mathcal{U}}, \mathcal{U}')$ denote a random variable representing the joint view of adversaries in an actual protocol execution and $SIM_{\mathcal{C}}^{\mathcal{U}, \lambda}(\{\mathbf{x}^i\}_{i \in \mathcal{U}}, \mathcal{U}')$ denote the joint view of adversaries in a simulated protocol execution. For all $\lambda > 2, \mathcal{U}, \mathbf{x}^i_{i \in \mathcal{U}}, \mathcal{U}', \mathcal{C} \subseteq \mathcal{U} \cup \{\mathcal{S}\}$ and SAP that can ensure privacy in the active adversary setting, there exists a PPT simulator *SIM* such that:*

$$SIM_{\mathcal{C}}^{\mathcal{U}, \lambda}(\{\mathbf{x}^i\}_{i \in \mathcal{U}}, \mathcal{U}') \equiv REAL_{\mathcal{C}}^{\mathcal{U}, \lambda}(\{\mathbf{x}^i\}_{i \in \mathcal{U}}, \mathcal{U}'), \quad (12)$$

where “ \equiv ” denotes the distributions are identical.

The proof of the theorem is provided in Appendix D.3.

5 EVALUATION

5.1 THEORETICAL COMPLEXITY ANALYSIS

Communication. PVF does not increase interaction-related overhead. The vectors sent from each user are $\mathbf{y}^i \in \mathbb{Z}_p^{(\lambda-1)\lceil \frac{m}{\lambda} \rceil}$ and $\mathbf{k}^i \in \mathbb{Z}_p^{\lceil \frac{m}{\lambda} \rceil}$, which contain the same number of entries as $\mathbf{x}_{pad}^i \in \mathbb{Z}_p^{m'}$. Hence, the theoretical communication complexity of SAP remains unchanged.

Computation. The additional computation operations required by PVF involve conducting $\lceil \frac{m}{\lambda} \rceil$ matrix-vector multiplications by both users and \mathcal{S} , incurring a computation cost of $O(\lambda m)$.

The theoretical complexity for one aggregation of users and \mathcal{S} , before and after PVF integration, is summarized in Table 4. For the method to **avoid padding**, please refer to Appendix E.4.

5.2 EXPERIMENTAL SETTINGS

Baselines. We include 7 baselines that encompass 5 types of **masking (encryption) schemes**:

- *PPDL* (Aono et al., 2017), utilizing *Single-private-key HE (Type 1)*, safeguards the global model, all while necessitating only a single interaction round between the users and \mathcal{S} .
- *EPPFL* (Li et al., 2022), relying on *Multi-private-key HE (Type 2)*, requires two interaction round.
- *NIVP-DS* (Xu et al., 2022), based on *SMPC (Type 3)*, requires two non-colluding servers, while involving only a single interaction round.
- *PracAgg* (Bonawitz et al., 2017), based on *Pair-wise Masking (Type 4)*, is widely regarded as the state of the art, involving multiple interactions rounds.
- *PracAgg+* (Bell et al., 2020), a type-4 scheme, exhibits a more efficient communication structure.
- *EffiAgg* (Liu et al., 2022a), based on *Non-pair-wise Masking (Type 5)*, requires server computation of discrete logarithms alongside multiple interactions between users and \mathcal{S} .
- *LPPFedL* (Wei et al., 2023), based on non-pair-wise masking, necessitates users to transmit multiple high-dimensional vectors along with multiple interactions.

To underscore the focal point of our work, our attention is exclusively directed towards secure-aggregation-related operations within the baselines, omitting other parts like “weight aggregation” in EPPFL. And masking schemes of other protocols that reduce **interaction-related** can be classified into the above five types, like Flamingo (Ma et al., 2023) (type 4), LTPA (Liu et al., 2023) (type 4).

Experimental settings. We run on a Linux workstation with 32GB of RAM and an AMD Ryzen 5 5600G. We use 1 NVIDIA GeForce RTX 3090 GPU only for model training, excluding the aggregation process. In all experimental settings, the space of the elements in the input vectors is 32-bit. Same with PracAgg, our main experiments are conducted under a semi-honest setting, with a specific focus on assessing the speed enhancement brought by PVF. We disregard operations such as digital signatures and public key infrastructure in the active adversary setting, which do not impact the asymptotics of the results (Bonawitz et al., 2017).

Evaluation Metrics. We use *Improvement Factor* = $\frac{\text{Time of (un)masking w/o PVF}}{\text{Time of (un)masking w/ PVF}}$ (speedup) to describe the efficacy of PVF. When evaluating communication overhead, we track costs for a single user, with the costs of \mathcal{S} being n times that of each user, a convention widely adopted in prior works (Liu et al., 2022a). The experimental results provided are the average outcomes from 5 repeated executions.

5.3 PERFORMANCE COMPARISON

Effectiveness of compressing costs. Results in Table 1 indicate that integrating PVF can yield a speedup ranging from $70\times$ to $99.5\times$ for the majority of baselines when $\lambda = 100$. The inability to achieve a $\lambda\times$ speedup is attributed to interaction-related overhead within SAP such as key agreements and secret sharing. Overall, the gain from PVF is the weakest for LPPFedL. This is because LPPFedL increases user communication overhead to achieve highly lightweight masking and unmasking. Consequently, interaction-related overhead constitutes a more significant portion of its computation time. As the original vector length increases from $100k$ to $500k$, the proportion

Table 1: Comparison of secure aggregation protocols computation time (in milliseconds) and communication cost (in KB) before and after integrating PVF for a single round. The number of users n is set to 100, with λ set as 100. **w/** denotes the corresponding protocol integrated with PVF.

Vector Length (m)	100k					500k					
	Operation Dropout rate (η)	Each user 0%	Server 0%	Server 10%	Server 30%	Comm. Cost 10%	Each user 0%	Server 0%	Server 10%	Server 30%	Comm. Cost 10%
PPDL w/ PVF	139010 1401	29759 304	27418 283	21789 225	27761 859	27761 859	708579 7128	149293 1571	136155 1446	107998 1133	138679 4291
Improvement Factor	99.2\times	97.9\times	96.8\times	96.8\times	32.3\times	32.3\times	99.4\times	95.0\times	94.1\times	95.3\times	32.3\times
EPPFL w/ PVF	3464 44	755163 7686	754261 7642	749279 7711	9962 681	9962 681	17582 219	3763769 38833	3732358 38924	3745663 39612	49798 3400
Improvement Factor	78.7\times	98.2\times	98.7\times	97.2\times	14.6\times	14.6\times	80.3\times	97.0\times	95.9\times	94.6\times	14.6\times
NIVP-DS w/ PVF	15 13	386 10	372 9	331 12	586 586	586 586	72 63	1902 49	1834 44	1789 39	2932 2931
Improvement Factor	1.2\times	38.6\times	41.3\times	27.6\times	\	\	1.1\times	38.9\times	41.7\times	45.9\times	\
PracAgg w/ PVF	13702 177	14249 187	139936 1472	307031 3207	785 785	785 785	73335 779	71607 794	697950 7063	1515347 15686	3910 3910
Improvement Factor	77.4\times	76.2\times	95.1\times	95.7\times	\	\	94.1\times	90.2\times	98.8\times	96.6\times	\
PracAgg+ w/ PVF	2773 38	13973 159	39735 410	69623 724	782 782	782 782	14487 179	70359 787	197540 2026	347055 4319	3908 3908
Improvement Factor	72.9\times	87.9\times	97.0\times	96.2\times	\	\	80.9\times	89.4\times	97.5\times	80.4\times	\
EffiAgg w/ PVF	1227 32	537771 5677	537329 5702	539814 6135	783 783	783 783	6014 99	2662161 27175	2730404 27440	2637768 28068	3908 3908
Improvement Factor	38.3\times	94.7\times	94.2\times	88.0\times	\	\	60.7\times	98.0\times	99.5\times	94.0\times	\
LPPFedL w/ PVF	176 22	63 21	59 20	46 17	1564 1564	1564 1564	846 53	319 77	291 70	218 61	7814 7814
Improvement Factor	8.0\times	3.0\times	3.0\times	2.7\times	\	\	16.0\times	4.1\times	4.2\times	3.6\times	\

of time spent on masking or encryption rises, so PVF exhibits a more noticeable acceleration effect for these SAPs in general. For PPDL and EPPFL (HE-based SAPs), only $\frac{1}{\lambda}$ of the original vector is transformed into **ciphertext**. Therefore, PVF brings them communication improvements of about 32.3 \times and 14.6 \times , respectively.

End-to-end comparison. In Figure 2, the model employed is LeNet (LeCun et al., 1998), consisting of 61,706 parameters. The dataset is MNIST (Deng, 2012) with each user possessing local data for only two labels. Across different aggregation methods, the model eventually achieves commendable training outcomes, and the integration of PVF showcases a significant acceleration effect, greatly reducing the training speed gap between using SAPs and plain aggregation. In Figure 3, consistent with FetchSGD (Rothchild et al., 2020), we use the ResNet9 (with 6.5M parameters), CIFAR10 dataset and PracAgg, with the optimal setting for FetchSGD and $k = 50,000$ for both FetchSGD and Top-k (Lu et al., 2023). It shows our method takes the **least** time and achieves the **best** accuracy because both FetchSGD and Top-k require more overhead for compression and PVF can ensure the intact aggregation of all entries in the original vector.

5.4 ABLATION STUDY

λ . We conduct experiments to analyze the variation in the acceleration effect of PVF on all baselines under different λ values, as depicted in Figure 5. PVF exhibits a more pronounced acceleration effect for participants with substantial original computational overhead, such as PPDL users. For certain SAPs like EffiAgg, the improvement factor of S can reach up to 1000 \times . As λ increases, the improvement factor of PVF gradually stabilizes and may experience a slight decline in certain SAPs. This trend emerges because when λ becomes sufficiently large, the primary computational overhead in secure aggregation shifts from **masking-related** overhead to **interaction-related** overhead like secret sharing. It’s worth noting that when λ reaches a certain threshold, for certain entities with relatively low primary computational load, such as users in NIVP-DS and servers in LPPFedL, integrating PVF may increase their computational overhead (improvement factor less than 1). This occurs because the computation cost incurred by the linear transformation in PVF becomes notable. But the gains from PVF for the servers in NIVP-DS and users in LPPFedL remain enticing, rendering the cost of the PVF linear transformation negligible.

Disrupting Variables Element. We focus on the impact of DVE on the model’s performance in this part. In the 32-bit input space, the σ is set to 8783 ($> \frac{2 \times 1000}{\sqrt{2\pi}}$), which is equivalent to adding noise with a standard deviation of 0.0409 (consistent with (Stevens et al., 2022)) to the aggregation




Figure 5: Computation acceleration for secure aggregation brought by different λ in PVF. $n = 100$, $m = 100k$, and $\eta = 10\%$.




Figure 6: The impact of DVE on model performance in practical applications. In these two tasks, $n = 100$, $\eta = 100$, and $\lambda=100$.

sum (float). And we explore the impact of DVE on model performance in two practical tasks: **(i) Image classification**, with LeNet model and MNIST dataset; **(ii) Movie recommendation**, using FedMF (Chai et al., 2020) model with the item embedding size of 32, i.e. 118,592 parameters in total, and ML-1M (Harper & Konstan, 2015) dataset. As shown in Figure 6, the impact of using PracAgg integrated with DVE on model performance is negligible.

n . Figure 7 displays the speedup effect of PVF across different n . The speedup remains consistently high for PracAgg+ and PPD across various n . However, as n rises, the speedup for PracAgg diminishes. This occurs because PracAgg necessitates multiple secret sharings for each user across the entire user set, along with multiple secret reconstructions by \mathcal{S} . These interaction-related overhead increases more significantly as n increases. Nonetheless, even when n reaches 1000, the gain brought by PVF for PracAgg servers remains above $85\times$, with user gain still exceeding $55\times$.




Figure 7: The gain introduced by PVF with respect to the number of users, with $\eta = 10\%$. $m = 100k$, and $\lambda = 100$.




Figure 8: Additional computational overhead incurred by PVF itself under different λ and vector lengths.

Transformation. The additional overhead of integrating PVF primarily arises from Equations (3), (4) and (6). Figure 8 illustrates the **overall** additional computation costs of PVF concerning different m and λ . It is evident that when $\lambda = 100$ and the original vector length is substantial (reaching $200k$), the computation time is less than 15 ms, nearly negligible compared to the computational time saved by PVF. As λ gradually increases while vector length remains fixed, the number of loops during freezing and thawing computations decreases in the beginning because of the reduced number of groups. When λ exceeds 60 and continues to increase, the scale of matrix-vector multiplication gradually rises, consequently increasing the computation cost. However, this cost remains consistently within the range of tens of milliseconds, posing a negligible computation burden.

6 CONCLUSION

We present a new perspective aimed at mitigating the formidable computation overhead of SAPs by reducing the number of involved entries while ensuring intact secure aggregation of original vectors. Based on this, we propose PVF, a concrete portable solution. After integrating with PVF, λ -SecAgg involves only $\frac{1}{\lambda}$ of original vectors. Moreover, we introduce the disrupting variables element to improve security. Extensive experiments showcase the remarkable improvements in acceleration and communication brought by PVF and its portability. Consequently, our method undoubtedly renders SAPs genuinely feasible, promising inspiration for future research endeavors.

REFERENCES

- 486
487
488 Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shiho Moriai, et al. Privacy-preserving deep learn-
489 ing via additively homomorphic encryption. *IEEE Transactions on Information Forensics and*
490 *Security*, 13(5):1333–1345, 2017.
- 491 James Bell, Kallista A. Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. Se-
492 cure single-server aggregation with (poly)logarithmic overhead. *Proceedings of the 2020 ACM*
493 *SIGSAC Conference on Computer and Communications Security*, 2020.
- 494
495 James Bell, Adrià Gascón, Tancrede Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and
496 Cathie Yun. {ACORN}: Input validation for secure aggregation. In *32nd USENIX Security*
497 *Symposium (USENIX Security 23)*, pp. 4805–4822, 2023.
- 498
499 Mihir Bellare and Chanathip Namprempre. Authenticated encryption: Relations among notions
500 and analysis of the generic composition paradigm. In *Advances in Cryptology—ASIACRYPT*
501 *2000: 6th International Conference on the Theory and Application of Cryptology and Information*
502 *Security Kyoto, Japan, December 3–7, 2000 Proceedings 6*, pp. 531–545. Springer, 2000.
- 503
504 Derian Boer and Stefan Kramer. Secure sum outperforms homomorphic encryption in (current)
505 collaborative deep learning. *arXiv preprint arXiv:2006.02894*, 2020.
- 506
507 Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. B. McMahan, Sarvar Patel,
508 Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving
509 machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Commu-
510 nications Security*, 2017.
- 511
512 Di Chai, Leye Wang, Kai Chen, and Qiang Yang. Secure federated matrix factorization. *IEEE*
513 *Intelligent Systems*, 36(5):11–20, 2020.
- 514
515 Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE*
516 *Signal Processing Magazine*, 29(6):141–142, 2012.
- 517
518 Tamer Ahmed Eltaras, Farida Sabry, Wadha Labda, Khawla Alzoubi, and Qutaibah Ahmedeltaras.
519 Efficient verifiable protocol for privacy-preserving aggregation in federated learning. *IEEE Trans-
520 actions on Information Forensics and Security*, 18:2977–2990, 2023.
- 521
522 Irem Ergun, Hasin Us Sami, and Basak Guler. Sparsified secure aggregation for privacy-preserving
523 federated learning. *ArXiv*, abs/2112.12872, 2021.
- 524
525 Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-
526 how easy is it to break privacy in federated learning? *Advances in Neural Information Processing*
527 *Systems*, 33:16937–16947, 2020.
- 528
529 Robin C. Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client
530 level perspective. *ArXiv*, abs/1712.07557, 2017.
- 531
532 Xiaojie Guo, Zheli Liu, Jin Li, Jiqiang Gao, Boyu Hou, Changyu Dong, and Thar Baker. Verifi-
533 communication-efficient and fast verifiable aggregation for federated learning. *IEEE Transac-
534 tions on Information Forensics and Security*, 16:1736–1751, 2021.
- 535
536 Changhee Hahn, Hodong Kim, Minjae Kim, and Junbeom Hur. Versa: Verifiable secure aggregation
537 for cross-device federated learning. *IEEE Transactions on Dependable and Secure Computing*,
538 20:36–52, 2023.
- 539
540 F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm*
541 *transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
- 542
543 Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural col-
544 laborative filtering. In *Proceedings of the 26th international conference on world wide web*, pp.
545 173–182, 2017.

- 540 Agrin Hilmkil, Sebastian Callh, Matteo Barbieri, Leon René Sütfield, Edvin Listo Zec, and Olof
541 Mogren. Scaling federated learning for fine-tuning of large language models. In *International*
542 *Conference on Applications of Natural Language to Information Systems*, pp. 15–23. Springer,
543 2021.
- 544 Swanand Kadhe, Nived Rajaraman, O Ozan Koyluoglu, and Kannan Ramchandran. Fast-
545 secagg: Scalable secure aggregation for privacy-preserving federated learning. *arXiv preprint*
546 *arXiv:2009.11248*, 2020.
- 547 Mandal Kalikinkar, Gong Guang, and Liu Chuyi. Nike-based fast privacy-preserving high-
548 dimensional data aggregation for mobile devices. *University of Waterloo, Waterloo, ON, Canada,*
549 *Tech. Rep. CACR*, 10:2018, 2018.
- 550 Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to
551 document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- 552 Yiran Li, Hongwei Li, Guowen Xu, Xiaoming Huang, and Rongxing Lu. Efficient privacy-
553 preserving federated learning with unreliable users. *IEEE Internet of Things Journal*, 9:11590–
554 11603, 2022.
- 555 Bo Liu, Ming Ding, Sina Shaham, Wenny Rahayu, Farhad Farokhi, and Zihuai Lin. When machine
556 learning meets privacy: A survey and outlook. *ACM Computing Surveys (CSUR)*, 54(2):1–36,
557 2021.
- 558 Ziyao Liu, Jiale Guo, Kwok-Yan Lam, and Jun Zhao. Efficient dropout-resilient aggregation for
559 privacy-preserving machine learning. *IEEE Transactions on Information Forensics and Security*,
560 18:1839–1854, 2022a.
- 561 Ziyao Liu, Jiale Guo, Wenzhuo Yang, Jiani Fan, Kwok-Yan Lam, and Jun Zhao. Privacy-preserving
562 aggregation in federated learning: A survey. *IEEE Transactions on Big Data*, 2022b.
- 563 Ziyao Liu, Hsiao-Ying Lin, and Yamin Liu. Long-term privacy-preserving aggregation with user-
564 dynamics for federated learning. *IEEE Transactions on Information Forensics and Security*, 18:
565 2398–2412, 2023.
- 566 Zizhen Liu, Si Chen, Jing Ye, Junfeng Fan, Huawei Li, and Xiaowei Li. Sash: Efficient secure
567 aggregation based on shprg for federated learning. In *Uncertainty in Artificial Intelligence*, pp.
568 1243–1252. PMLR, 2022c.
- 569 Shiwei Lu, Ruihu Li, Wenbin Liu, Chaofeng Guan, and Xiaopeng Yang. Top-k sparsification with
570 secure aggregation for privacy-preserving federated learning. *Computers & Security*, 124:102993,
571 2023.
- 572 Hidde Lycklama, Lukas Burkhalter, Alexander Viand, Nicolas Küchler, and Anwar Hithnawi. Rofl:
573 Robustness of secure federated learning. In *2023 IEEE Symposium on Security and Privacy (SP)*,
574 pp. 453–476. IEEE, 2023.
- 575 Jing Ma, Si-Ahmed Naas, Stephan Sigg, and Xixiang Lyu. Privacy-preserving federated learning
576 based on multi-key homomorphic encryption. *International Journal of Intelligent Systems*, 37(9):
577 5880–5901, 2022.
- 578 Yiping Ma, Jess Woods, Sebastian Angel, Antigoni Polychroniadou, and Tal Rabin. Flamingo:
579 Multi-round single-server secure aggregation with applications to private federated learning. *2023*
580 *IEEE Symposium on Security and Privacy (SP)*, pp. 477–496, 2023.
- 581 Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
582 Communication-efficient learning of deep networks from decentralized data. In *Artificial intelli-*
583 *gence and statistics*, pp. 1273–1282. PMLR, 2017.
- 584 Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev,
585 Matthew Rocklin, Amit Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rath-
586 nayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta,
587 Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán

- 594 Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrmam, and Anthony Sco-
595 patz. SymPy: symbolic computing in Python. *PeerJ Computer Science*, 3, January 2017. doi:
596 10.7717/peerj-cs.103. URL <https://doi.org/10.7717/peerj-cs.103>.
597
- 598 Hamid Mozaffari, Virat Shejwalkar, and Amir Houmansadr. Every vote counts: Ranking-based
599 training of federated learning to resist poisoning attacks. In *USENIX Security Symposium*, 2023.
- 600 Dario Pasquini, Danilo Francati, and Giuseppe Ateniese. Eluding secure aggregation in federated
601 learning via model inconsistency. In *Proceedings of the 2022 ACM SIGSAC Conference on Com-
602 puter and Communications Security*, pp. 2429–2443, 2022.
- 603 Torben Pryds Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing.
604 In *Annual international cryptology conference*, pp. 129–140. Springer, 1991.
- 605
- 606 Mayank Rathee, Conghao Shen, Sameer Wagh, and Raluca Ada Popa. Elsa: Secure aggregation for
607 federated learning with malicious actors. In *2023 IEEE Symposium on Security and Privacy (SP)*,
608 pp. 1961–1979. IEEE, 2023.
- 609 Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *Journal of
610 the ACM (JACM)*, 56(6):1–40, 2009.
- 611
- 612 Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman,
613 Joseph Gonzalez, and Raman Arora. Fetchsgd: Communication-efficient federated learning with
614 sketching. In *International Conference on Machine Learning*, pp. 8253–8265. PMLR, 2020.
- 615 Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- 616
- 617 Jinhyun So, Başak Güler, and A Salman Avestimehr. Turbo-aggregate: Breaking the quadratic
618 aggregation barrier in secure federated learning. *IEEE Journal on Selected Areas in Information
619 Theory*, 2(1):479–489, 2021.
- 620 Jinhyun So, Chaoyang He, Chien-Sheng Yang, Songze Li, Qian Yu, Ramy E Ali, Basak Guler, and
621 Salman Avestimehr. Lightsecagg: a lightweight and versatile design for secure aggregation in
622 federated learning. *Proceedings of Machine Learning and Systems*, 4:694–720, 2022.
- 623 Jinhyun So, Ramy E Ali, Başak Güler, Jiantao Jiao, and A Salman Avestimehr. Securing secure
624 aggregation: Mitigating multi-round privacy leakage in federated learning. In *Proceedings of the
625 AAAI Conference on Artificial Intelligence*, volume 37, pp. 9864–9873, 2023.
- 626
- 627 Ekanut Sotthiwat, Liangli Zhen, Zengxiang Li, and Chi Zhang. Partially encrypted multi-party
628 computation for federated learning. *2021 IEEE/ACM 21st International Symposium on Cluster,
629 Cloud and Internet Computing (CCGrid)*, pp. 828–835, 2021.
- 630 Timothy Stevens, Christian Skalka, Christelle Vincent, John Ring, Samuel Clark, and Joseph Near.
631 Efficient differentially private secure aggregation for federated learning via hardness of learning
632 with errors. In *31st USENIX Security Symposium (USENIX Security 22)*, pp. 1379–1395, 2022.
- 633
- 634 PK Suetin, Alexandra I Kostrikin, and Yu I Manin. *Linear algebra and geometry*. CRC Press, 1989.
- 635 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
636 lay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open founda-
637 tion and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 638 Chuanyin Wang, Cunqing Ma, Min Li, Neng Gao, Yifei Zhang, and Zhuoxiang Shen. Protecting
639 data privacy in federated learning combining differential privacy and weak encryption. In *Science
640 of Cyber Security: Third International Conference, SciSec 2021, Virtual Event, August 13–15,
641 2021, Revised Selected Papers 4*, pp. 95–109. Springer, 2021.
- 642
- 643 Yong Wang, Aiqing Zhang, Shu-Lin Wu, and Shui Yu. Vosa: Verifiable and oblivious secure aggre-
644 gation for privacy-preserving federated learning. *IEEE Transactions on Dependable and Secure
645 Computing*, 20:3601–3616, 2023.
- 646 Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek,
647 and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance
analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.

Zhaohui Wei, Qingqi Pei, Ning Zhang, Xuefeng Liu, Celimuge Wu, and Amirhosein Taherkordi. Lightweight federated learning for large-scale iot devices with privacy guarantee. *IEEE Internet of Things Journal*, 10:3179–3191, 2023.

Yi Xu, Changgen Peng, Weijie Tan, Youliang Tian, Minyao Ma, and Kun Niu. Non-interactive verifiable privacy-preserving federated learning. *Future Generation Computer Systems*, 128:365–380, 2022.

Rui Ye, Wenhao Wang, Jingyi Chai, Dihan Li, Zexi Li, Yinda Xu, Yaxin Du, Yanfeng Wang, and Siheng Chen. Openfedllm: Training large language models on decentralized private data via federated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6137–6147, 2024.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.

Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019.

A NOTATIONS

The primary notations used in this work are listed in Table 2.

Table 2: Notations.

Symbol	Definition
n	Number of users
m	Length of original vectors
η	Dropout rate
λ	Compression factor
μ	The number of broken elements within every λ elements
\mathcal{U}	The user set
\mathcal{U}'	The surviving user set
\mathcal{S}	The server
$\mathbf{x}^{i(t)}$	Original vector of i in the t -th round (with DP-based noise)
\mathbf{A}	A $\lambda \times \lambda$ invertible matrix
\mathbf{A}^{-1}	The inverse matrix of \mathbf{A}
$\check{\mathbf{A}}$	The matrix composed only of the first $\lambda - 1$ rows of \mathbf{A}
α	The vector composed of the λ -th row of \mathbf{A}
$\check{\mathbf{A}}^{\mu+1}$	$\check{\mathbf{A}}$ with μ -security
$\alpha^{\mu+1}$	α with μ -security
$SLE_{AK}(\mathbf{Ax})$	Function for computing \mathbf{x} of \mathbf{Ax} with the knowledge AK
$\mathbf{x}_{pad}^{i(t)}$	$\mathbf{x}^{i(t)}$ after padding
m'	Length of $\mathbf{x}_{pad}^{i(t)}$
$\mathbf{d}_j^{i(t)}$	The j -th group of $\mathbf{x}_{pad}^{i(t)}$
$\mathbf{y}^{i(t)}$	Frozen vector of $\mathbf{x}_{pad}^{i(t)}$
\mathbf{GN}	DP-based noise generated by the Gaussian mechanism
$\mathbf{k}^{i(t)}$	Key vector of $\mathbf{x}_{pad}^{i(t)}$
$\mathbf{c}^{i(t)}$	Commitment vector of $\mathbf{x}^{i(t)}$
$\zeta^{i(t)}$	Corresponding random vector of $\mathbf{c}^{i(t)}$

B OPTIONAL EXTENSIONS

We introduce 3 optional extensions augmenting functionality and security for better portability: (i) μ -security *Extension*, which resists adversaries that possess prior knowledge of partial private

vectors; (ii) *Result Verification Extension* (RVE), which ensures the correctness of performing transformations to prevent malicious server-side computations, and (iii) *User Commitment Extension* (UCE), which ensures the uniqueness of users’ original vectors to prevent malicious attempts at incorrect computations. The detailed pipeline for integrating PVF with all 4 extensions into SAP is depicted in Figure 13.

B.1 μ -SECURITY EXTENSION

Threat. Consider an extreme scenario where an extremely powerful malicious server \mathcal{S} somehow obtains an element x_e^i of \mathbf{x}^i , although this is generally unrealistic in most cases. In this case, \mathcal{S} , during the PVF (without DVE) execution, would acquire $\lambda - 1$ equations of the j -th group where x_e^i lies, namely $\mathbf{y}_j^i = \tilde{\mathbf{A}}\mathbf{d}_j^i$ in Equation (3), along with x_e^i . This allows \mathcal{S} to solve for all elements in the j -th group. Even though only a very small fraction of original vectors, $\frac{\lambda}{m}$, is leaked, we still aim to fortify against this. Assuming the adversary’s capability is to steal μ ($< \lambda - 1$) elements from each group of \mathbf{x}^i .

We expand the definitions of incomplete matrix $\tilde{\mathbf{A}}$, and residual vector $\boldsymbol{\alpha}$ of PVF. We define *Incomplete Matrix with μ -security* $\tilde{\mathbf{A}}^{\mu+1} = \mathbf{A}_{:\lambda-\mu-1,:}$ and *Residual Matrix with μ -security* $\boldsymbol{\alpha}^{\mu+1} = \mathbf{A}_{\lambda-\mu,:}$. Then *Frozen Vector with μ -security* is:

$$\mathbf{y}^i = (\tilde{\mathbf{A}}^{\mu+1}\mathbf{d}_1^i, \tilde{\mathbf{A}}^{\mu+1}\mathbf{d}_2^i, \dots, \tilde{\mathbf{A}}^{\mu+1}\mathbf{d}_l^i), \quad (13)$$

and *Key Vector with μ -security* is:

$$\mathbf{k}^i = (\boldsymbol{\alpha}^{\mu+1}\mathbf{d}_1^i, \boldsymbol{\alpha}^{\mu+1}\mathbf{d}_2^i, \dots, \boldsymbol{\alpha}^{\mu+1}\mathbf{d}_l^i). \quad (14)$$

Clearly, the aforementioned changes won’t affect the execution process and security of PVF. However, the compression factor of PVF will decrease from λ to $\frac{\lambda}{\mu+1}$ ($0 \leq \mu < \lambda - 1$). With these modifications, the adversary, apart from the known private elements, cannot obtain additional elements from the execution of PVF. Particularly, in the methodology outlined in Section 3.2, we consider $\mu = 0$.

B.2 RESULT VERIFICATION EXTENSION

Threat. The malicious server might intentionally provide incorrect aggregated results to disrupt training, even if it **doesn’t compromise user privacy**. When integrating PVF with SAPs featuring result verification capabilities, we contemplate incorporating RVE.

Most recent SAPs that can verify aggregated results typically necessitate honest and non-dropping participants, such as *Collectors* (Wang et al., 2023) and *Auxiliary Nodes* (Eltaras et al., 2023), introducing additional assumptions. Consequently, by utilizing this extension, we unavoidably add security assumptions. Without loss of generality, we adopt the approach and security assumptions from VerSA (Hahn et al., 2023) for PVF, which requires the server and users not to collude but does not need other trusted third parties. Most SAPs (Aono et al., 2017; Bonawitz et al., 2017; Bell et al., 2020; Li et al., 2022; Xu et al., 2022; Liu et al., 2022a; Stevens et al., 2022; Wei et al., 2023) lack the functionality to verify aggregated results. This extension merely constrains the server within the PVF module to obtain the correct $\sum_{i \in \mathcal{U}'} \mathbf{y}^i$, and it allows secure integration of PVF and SAPs in the face of aggregation falsification attacks.

Phase 0: MainRV.Setup(\cdot). Execute Main.Setup(\cdot).

Phase 1: MainRV.Freeze(\cdot). Execute Main.Freeze(\cdot).

Phase 2: MainRV.SecAgg(\cdot). During the execution of VerSA, all users obtain a set of random vectors (κ_1, κ_2) derived from shared keys. In the verification step, users send $\hat{\mathbf{y}}^i = \kappa_1 \mathbf{y}^i$ and $\hat{\mathbf{y}}^i = \mathbf{y}^i + \kappa_2$ to the server, which simultaneously aggregates $\hat{\mathbf{y}}^i$ and $\hat{\mathbf{y}}^i$. In this process, the verification of the sum of frozen vectors and the sum of key vectors is combined, avoiding additional interaction. VerSA ensures users obtain a consistent \mathcal{U}' and the correct $\sum_{i \in \mathcal{U}'} \mathbf{k}^i$. Users also receive $\sum_{i \in \mathcal{U}'} \hat{\mathbf{y}}^i$ and $\sum_{i \in \mathcal{U}'} \hat{\mathbf{y}}^i$.

Phase 3: MainRV.Thaw(\cdot). After receiving the results, users proceed with verification:

$$\sum_{i \in \mathcal{U}'} \hat{\mathbf{y}}^i \setminus \kappa_1 \stackrel{?}{=} \sum_{i \in \mathcal{U}'} \hat{\mathbf{y}}^i - |\mathcal{U}'| \kappa_2. \quad (15)$$

\mathbf{y}^i is unknown to \mathcal{S} throughout the process. If the verification passes, execute `Main.Thaw(\cdot)` for thawing on the user side. Otherwise, conclude that the server deviates from the protocol and terminate the execution.

B.3 USER COMMITMENT EXTENSION

Threat. In PVF, if a malicious user i uses inconsistent \mathbf{d}_j^i and \mathbf{d}'_j^i for freezing, resulting in obtaining $\check{\mathbf{A}}\mathbf{d}_j^i$ and $\mathbf{a}\mathbf{d}'_j^i$, or applies mismatched $\check{\mathbf{A}}$ and α' to \mathbf{d}_j^i for linear transformations, the frozen vector and key vector no longer correspond. Although this remains **harmless to the privacy** of honest users, it will lead to inaccuracies in the final aggregated result.

This malicious tampering is unrelated to SAPs. In the freezing phase of PVF, we use *Pedersen Commitment* (Pedersen, 1991) to ensure users' vector consistency. **Note** that, to enable verification, it's crucial to ensure that during the validation phase, the server is semi-honest and cannot collude with users. Otherwise, a malicious server could illegitimately approve malicious user actions, making user verification invalid, which is widely adopted in prior works (Rathee et al., 2023). Below, we outline UCE to modify the Main method to achieve the aforementioned functionality. UCE supports consistency verification of the user inputs in scenarios with advanced needs. Note that UCE is only applicable to SAPs where the server has access to the aggregated results.

Phase 0: MainUC.Setup(\cdot). Execute `Main.Setup(\cdot)`. Given the security parameter ρ , it generates the group (\mathbb{G}_p, p, g) , where p is the order of \mathbb{G}_p and g is its generator. h is an element of \mathbb{G}_p . \mathbb{G}_p, p, g, h are public.

Phase 1: MainUC.Freeze(\cdot). The user $i \in [1, n]$ generates the *Random Vector* $\zeta^i = (\zeta_1^i, \zeta_2^i, \dots, \zeta_{i\lambda}^i)$ and calculates the *Commitment Vector*:

$$\begin{aligned} \mathbf{c}^i &= \left((c_1^i, \dots, c_\lambda^i), \dots, (c_{(i-1)\lambda+1}^i, \dots, c_{i\lambda}^i) \right) \\ &= ((g^{x_1^i} h^{\zeta_1^i}, g^{x_2^i} h^{\zeta_2^i}, \dots, g^{x_\lambda^i} h^{\zeta_\lambda^i}), \dots, (g^{x_{(i-1)\lambda+1}^i} h^{\zeta_{(i-1)\lambda+1}^i}, g^{x_{(i-1)\lambda+2}^i} h^{\zeta_{(i-1)\lambda+2}^i}, \dots, g^{x_{i\lambda}^i} h^{\zeta_{i\lambda}^i})). \end{aligned} \quad (16)$$

Execute `Main.Freeze(\cdot)`.

Phase 2: MainUC.SecAgg(\cdot). Execute `Main.SecAgg(\cdot)`, once completed, all users obtain the aggregated result *sum*. Each user i sends \mathbf{c}^i and ζ^i to the server. For $j \in [1, l]$, $r \in [1, \lambda]$, \mathcal{S} validates:

$$\prod_{i \in \mathcal{U}'} c_{(j-1)\lambda+r}^i \stackrel{?}{=} g^{\text{sum}_{(j-1)\lambda+r}} h^{\sum_{i \in \mathcal{U}'} \zeta_{(j-1)\lambda+r}^i}. \quad (17)$$

If the validation fails, it implies the user deviates from the protocol, and the protocol is terminated. Otherwise, it signifies that the user employs the consistent original vector when generating the frozen vector and key vector, and the utilized public parameters are accurate.

Phase 3: MainUC.Thaw(\cdot). Execute `Main.Thaw(\cdot)`.

Clearly, this extension will inevitably introduce a performance decrease. It's worth noting that in the majority of SAPs, many malicious user behaviors, such as incorrectly executing key agreements or submitting counterfeited secret shares, result in aggregation errors. And strictly enforcing user behavior remains a pending issue (Ma et al., 2023).

C PORTABILITY ANALYSIS

PVF does not attempt to alter SAP, and the decoupling greatly enhances the portability. PDDL (Aono et al., 2017) and EPPFL (Li et al., 2022) employ homomorphic encryption, eliminating the need for secret sharing. NIVP-DS (Xu et al., 2022) constitutes a dual-server secure multi-party computation scheme, requiring users to share secrets between two servers. PracAgg (Bonawitz et al., 2017) and PracAgg+ (Bell et al., 2020) represent classic mask-based solutions, with PracAgg+ necessitating an additional user grouping process. EffiAgg (Liu et al., 2022a) and LPPFedL (Wei et al., 2023) respectively introduce specialized masking mechanisms to lightweight PracAgg. VerSA (Hahn et al., 2023) enables users to verify the aggregated result at the conclusion, and its integration with PVF requires RVE. LTPA (Liu et al., 2023) and MRSA (So et al., 2023) individually design specific user selection mechanisms to ensure privacy in **multi-round aggregation**. Resistance against **model**

Table 3: Symbolic representation of SAP execution process without and with PVF. The symbols’ meanings are as follows: US : User Selection; UG : User Grouping; KA : Key Agreements among Users; SS : Secret Sharing; E/M : Encryption or Masking user’s original vectors; UA : Upload and Aggregation; D/U : Decryption or Unmasking vectors; Ver : Verification of aggregated results; F : Freezing process in PVF Main Method, generating frozen vectors and key vectors; T : Thawing process in PVF Main Method, deriving all entries based on aggregated entries; RV : Result Verification Extension of PVF; X' : Specialized design for process X in the corresponding protocol; $SA(\cdot)$: Operations involved in the aggregation process; \square : Protocol can execute in the semi-honest setting; \blacksquare : Protocol can execute in the active adversary setting.

Type	Scheme	w/o PVF	w/ PVF
\	PlainAgg	$[US, UA]$	\
HE-based	PPDL	$[US, SA(KA, E/M, UA, D/U)]_{\square}$	$[US, F, SA(KA, E/M, UA, D/U), T]_{\square}$
	EPPFL	$[US, SA(KA, E/M, UA, D/U)]_{\square}$	$[US, F, SA(KA, E/M, UA, D/U), T]_{\square}$
SMPC-based	NIVP-DS	$[US, SA(KA, SS, E/M, UA, D/U)]_{\square}$	$[US, F, SA(KA, SS, E/M, UA, D/U), T]_{\square}$
Mask-based	PracAgg	$[US, SA(KA, SS, E/M, UA, D/U)]_{\square\blacksquare}$	$[US, F, SA(KA, SS, E/M, UA, D/U), T]_{\square\blacksquare}$
	PracAgg+	$[US, SA(UG, KA, SS, E/M, UA, D/U)]_{\square\blacksquare}$	$[US, F, SA(UG, KA, SS, E/M, UA, D/U), T]_{\square\blacksquare}$
	EffiAgg	$[US, SA(KA, SS, E/M', UA, D/U')]_{\square\blacksquare}$	$[US, F, SA(KA, SS, E/M', UA, D/U'), T]_{\square\blacksquare}$
	LPPFedL	$[US, SA(KA, SS, E/M', UA, D/U')]_{\square}$	$[US, F, SA(KA, SS, E/M', UA, D/U'), T]_{\square}$
Result Veri.	VerSA	$[US, SA(KA, SS, E/M, UA, D/U, Ver)]_{\square}$	$[US, F, SA(KA, SS, E/M, UA, D/U, Ver), RV, T]_{\square}$
Multi. Privacy	LTPA	$[US', SA(KA, SS, E/M, UA, D/U)]_{\square}$	$[US', F, SA(KA, SS, E/M, UA, D/U), T]_{\square}$
	MRSA	$[US', SA(KA, SS, E/M, UA, D/U)]_{\square}$	$[US', F, SA(KA, SS, E/M, UA, D/U), T]_{\square}$
Resist. M. Incon.	\	$[US, SA(KA, SS, E/M', UA, D/U)]_{\square}$	$[US, F, SA(KA, SS, E/M', UA, D/U), T]_{\square}$

inconsistency attacks can be achieved by making slight modifications to PRG without incurring additional overhead (Ma et al., 2023). PVF also fits the **asynchronous** setting. Since PVF itself is one-shot and decoupled from specific SAP, it does not affect the one-shot masking or recovery in asynchronous SAP (such as LightSecAgg (So et al., 2022)). We symbolically represent the entire process of federated learning aggregation in Table 3. It is evident that PVF is decoupled from SAPs, not interfering with the internal execution process of SAP.

D DETAILED SECURITY ANALYSIS

D.1 EXAMPLE OF IMPROPER MATRIX

Example 1. Consider a 3×3 matrix:

$$\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 3 \\ 1 & 2 & 4 \end{bmatrix}, \quad (18)$$

which is an invertible matrix. The corresponding incomplete matrix is:

$$\check{\mathbf{A}} = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 3 \end{bmatrix}. \quad (19)$$

Assume the original vector is $\mathbf{x} = (x_1; x_2; x_3) = (1; 2; 3)$, and the frozen vector is $\check{\mathbf{A}}\mathbf{x} = (14, 16)$. \mathcal{S} obtains the under-determined system of linear equations as follows:

$$\begin{cases} x_1 + 2x_2 + 3x_3 = 14 & (i) \\ x_1 + 3x_2 + 3x_3 = 16 & (ii) \end{cases} \quad (20)$$

While \mathcal{S} cannot obtain the complete \mathbf{x} , it can deduce $x_2 = 2$ by $(ii) - (i)$. Unlike \mathcal{S} obtaining prior knowledge of elements through attacks as discussed in Section B.1, here, \mathcal{S} deduces $x_2 = 2$ through the computation process within PVF. However, since \mathbf{A} is public, any maliciously constructed \mathbf{A} can be easily detected by honest users.

864 D.2 SUPPLEMENTARY CRYPTOGRAPHIC PRIMITIVES

865 Here, we supplement the symmetric authenticated encryption and the digital signature we require in
866 the active adversary model.

867 D.2.1 SYMMETRIC AUTHENTICATED ENCRYPTION

868 Symmetric authenticated encryption can ensure the confidentiality of a message, including:

- 869 • $AE.gen(k) \rightarrow (sk)$, where k is the security parameter. It outputs a secret key sk .
- 870 • $AE.enc(sk, m) \rightarrow (c)$. It encrypts the message m using sk and outputs the ciphertext c .
- 871 • $AE.dec(sk, c) \rightarrow m$ or \perp . If sk is the correct key corresponding to the ciphertext c and c passes
872 integrity verification, it outputs the plaintext m . Otherwise, it outputs an error symbol.

873 We need the encryption scheme to be indistinguishable under chosen plaintext attacks (IND-CPA)
874 and ciphertext integrity (IND-CTXT) (Bellare & Namprempre, 2000). In Figure 13, we omit the
875 encryption of messages before transmission and the decryption after reception by each participant.
876 If any error occurs during encryption or decryption process, the protocol will be immediately termi-
877 nated.

878 D.2.2 DIGITAL SIGNATURE

879 Digital signature can ensure the authenticity and integrity of a message. We use the signature scheme
880 that achieves security against universal forgery under chosen message attack (UF-CMA). The digital
881 signature scheme consists of:

- 882 • $DS.gen(k) \rightarrow (sk, pk)$, where k is the security parameter. It outputs a secret key sk and a public
883 key pk .
- 884 • $DS.sign(sk, m) \rightarrow (sig)$. It outputs a digital signature sig on the message m .
- 885 • $DS.verify(sig, pk, m) \rightarrow True$ or $False$. It verifies whether the signature sig is valid on m .

886 D.3 DETAILED EXPLANATION AND PROOF OF THEOREM 1

887 First and foremost, it is evident that in PVF, the user only transmits \mathbf{y}^i to the server, and the server
888 only sends $\sum \mathbf{y}^i$ back to the users after SAP ends. Notably, \mathbf{k}^i and $\sum \mathbf{k}^i$ are transmitted through
889 the SAP, independent of PVF. In the active adversary model, we obviously cannot guarantee the
890 correctness of the aggregation result because the malicious server can arbitrarily modify the result.
891 However, we can guarantee the privacy of honest users' inputs. We provide a detailed explanation
892 of the active attacks that malicious participants can launch within PVF and how PVF leverages the
893 cryptographic primitives to defend against them.

- 894 • Forging fake users to participate in PVF. This type of attack, also known as a *Sybil Attack*, involves
895 fake users reporting received information to the server. Such attacks primarily target scenarios
896 where users share secret keys among themselves but keep the keys secret from the server, like
897 PDDL. Alternatively, an attacker may attempt to forge a large number of fake users (more than
898 $\frac{1}{3}|\mathcal{U}|$) to reconstruct users' private keys in the secret-sharing scheme. Since PVF does not involve
899 information that is kept secret from the server but shared among all users, and consistent with the
900 assumption in PracAgg that the number of malicious users does not exceed $\frac{1}{3}|\mathcal{U}|$, PVF is resistant
901 to this type of attack.
- 902 • Attempting to forge or tamper with honest users' messages. Such attacks may occur in PVF in the
903 following situations: malicious participants forging or tampering with an honest user's \mathbf{y}^i . This
904 can be avoided by the digital signature σ_1^i employed in PVF. Similarly, malicious participants may
905 attempt to forge or tamper with $\sum \mathbf{y}^i$ sent by the server, which is prevented by the use of σ_3 .
- 906 • Sending malformed messages. In PVF, such attacks include malicious users sending malformed
907 ciphertexts of \mathbf{y}^i or the malicious server sending malformed ciphertexts of $\sum \mathbf{y}^i$. Such attacks are
908 prevented by the IND-CPA and IND-CTXT security of the symmetric authenticated encryption
909 used in PVF. If decryption fails, the protocol is immediately terminated.

- Intercepting and stealing private information. Malicious adversaries may intercept messages sent by honest users to extract private information. This is effectively avoided by the symmetric authenticated encryption employed in PVF.

The use of symmetric authenticated encryption and digital signatures to ensure privacy under the active adversary model is a relatively mature application in the field of secure aggregation, and our design follows these existing works. Then we present the following lemmas:

Lemma 1 (Privacy during the freezing phase). Fix $p, \mathcal{U}, m, \lambda, \mathbf{A}$ and a private vector $\mathbf{x}^i = (\mathbf{d}_1^i, \dots, \mathbf{d}_{\lceil \frac{m}{\lambda} \rceil}^i)$ (with noise) of an honest user $i \in \mathcal{U}$. For any probabilistic polynomial-time (PPT) adversary M who is given $\{\mathbf{y}^i\}_{i \in \mathcal{U}}$ and \mathcal{C} , the advantage of M to obtain any unbroken element a_j is defined as:

$$Adv_M^y(\lambda) := Pr[SLE_C(\tilde{\mathbf{A}}\mathbf{d}_j^i) \rightarrow a_j]_{j \in [1, \lceil \frac{m}{\lambda} \rceil]}. \quad (21)$$

There exists a negligible function ε such that $Adv_M^y(\lambda) \leq \varepsilon$.

Remark 1. The adversary can obtain $l(\lambda - 1)$ independent equations, with $l\lambda$ variables. Hence there are infinite solutions, and the probability of M determining the unique (\mathbf{x}^i) is $\frac{1}{\infty}$.

Lemma 2 (Privacy during the thawing phase). Fix $p, \mathcal{U}, m, \lambda, \mathbf{A}$ and the sum of private vectors $\sum_{i \in \mathcal{U}'} \mathbf{x}^i$. For any PPT adversary M who is given $\{\mathbf{y}^i\}_{i \in \mathcal{U}'}$, \mathcal{C} and $\sum_{i \in \mathcal{U}'} \mathbf{k}^i$, the advantage of M to obtain any unbroken element a_j is defined as:

$$Adv_M^{y,k}(\lambda) := Pr[SLE_C(\mathbf{A} \sum_{j \in \mathcal{U}'} \mathbf{d}_j^i) \rightarrow a_j]_{j \in [1, \lceil \frac{m}{\lambda} \rceil]}. \quad (22)$$

There exists a negligible function ε such that $Adv_M^{y,k}(\lambda) \leq \varepsilon$.

Remark 2. \mathcal{S} can obtain $\sum_{j \in \mathcal{U}'} \mathbf{d}_j^i$ by Equation (6) or Equation (9). For any individual \mathbf{d}_j^i (with noise), the information known to M is $\sum_{j \in \mathcal{U}'} \mathbf{d}_j^i$ and $\tilde{\mathbf{A}}\mathbf{d}_j^i$. So there are still infinite solutions, and $Adv_M^{y,k}(\lambda) = \frac{1}{\infty}$.

Lemma 3 (The hardness of the Learning With Errors decision problem). Given a finite field \mathbb{F}_q and a discrete probability distribution \mathcal{X} over \mathbb{F}_q . Let $\mathbf{s} \in \mathbb{F}_q^v$ be a secret vector, $\mathbf{A} \in \mathbb{F}_q^{u \times v}$ be a matrix that is chosen uniformly at random and $\mathbf{e} \in \mathbb{F}_q^u$ be the error vector that is sampled from \mathcal{X} . (v, q, σ) parameterize an LWE instance, where σ is the standard deviation of \mathcal{X} . The Learning With Errors (LWE) (search) problem is to find \mathbf{s} , given the pair (\mathbf{A}, \mathbf{b}) , where $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e}$. And the LWE decision problem is to distinguish between two uniformly randomly generated pairs.

Remark 3. Regev (2009) shows that if the size of q is polynomial in v and \mathcal{X} is a discrete Gaussian distribution on \mathbb{F}_q with standard deviation $\sigma > \frac{2\sqrt{v}}{\sqrt{2\pi}}$, the LWE decision problem is at least as hard as the LWE search problem and solving the LWE search problem can be reduced to solving the Shortest Vector Problem. In DVE, $v = \lambda$, and we use \mathbb{Z}_p as \mathbb{F}_q .

We use a standard hybrid argument to prove the theorem.

Proof. We define a sequence of hybrid distributions H_0, H_1, \dots to denote a series of modifications to *REAL*, which can finally get *SIM*. We prove *SIM* and *REAL* are indistinguishable by proving two adjacent hybrids are indistinguishable.

H_0 In this hybrid, *SIM* is exactly the same as *REAL*.

H_1 This hybrid is distributed similarly to the previous one, except for the following modifications. *SIM* obtains $\sum_{i \in \mathcal{U}' \setminus \mathcal{C}} \mathbf{x}^i$ by calling $Ideal_{\{\mathbf{x}^i\}_{i \in \mathcal{U}' \setminus \mathcal{C}}}(\mathcal{U}' \setminus \mathcal{C})$. *SIM* aborts if there is an illegal request. We replace the ciphertexts of $\{\mathbf{y}^i\}_{i \in \mathcal{U}}$ with the ciphertexts of uniformly random vectors $\{\mathbf{w}^i\}_{i \in \mathcal{U}}$ that satisfy $\sum_{i \in \mathcal{U}' \setminus \mathcal{C}} \mathbf{w}^i = \sum_{i \in \mathcal{U}' \setminus \mathcal{C}} \mathbf{y}^i$. $\sum_{i \in \mathcal{U}' \setminus \mathcal{C}} \mathbf{y}^i$ can be computed from Equation (9) based on $\sum_{i \in \mathcal{U}' \setminus \mathcal{C}} \mathbf{x}^i$. The IND-CPA and IND-CTXT security of symmetric authenticated encryption guarantees the distribution of this hybrid is indistinguishable from the previous one.

H_2 This hybrid is distributed exactly as the previous one, except *SIM* aborts if there is an invalid signature $(\sigma_1^i, \sigma_2^i$ or $\sigma_3)$. The UF-CMA security of the digital signature scheme can ensure \mathcal{C} cannot

forge any valid signature of an honest user, so the distribution of this hybrid is indistinguishable from the previous one.

H_3 This hybrid is distributed similarly to the previous one, except for the following modifications. **Firstly**, according to the security analysis process of the integrated SAP, we replace the corresponding messages conveyed in SAP with random strings of equal length. **Secondly**, we replace the frozen vectors $\{\mathbf{y}^i\}_{i \in \mathcal{U}}$ received by \mathcal{S} with uniformly random vectors $\{\mathbf{w}^i\}_{i \in \mathcal{U}}$. \mathbf{x}^i is added with noise ($\tilde{\mathbf{A}}\mathbf{e}$) through Equation (10). Therefore, Lemma 3 ensures that $(\tilde{\mathbf{A}}, \mathbf{y}^i)$ and $(\tilde{\mathbf{A}}, \mathbf{w}^i)$ are indistinguishable, which guarantees \mathcal{S} does not obtain private information from honest users through frozen vectors. Therefore, the security of SAP and the lemmas ensure the distribution of this hybrid is indistinguishable from the previous one.

Therefore, the distribution of SIM which is the same as H_3 is indistinguishable from $REAL$. SIM does not depend on the inputs of honest parties, and \mathcal{C} can only learn about the sum of original vectors. If too many users drop out, SAP will abort and still guarantee the above conclusion. Clearly, the security of λ -SecAgg still holds in the semi-honest setting. \square

E ADDITIONAL COMPARISON AND EXPERIMENT

E.1 THEORETICAL COMPLEXITY COMPARISON

Table 4 indicates for some SAPs, the theoretical computation complexity increases after integrating PVF. This is attributed to the transformation required for the entire original vector within PVF. However, intuitively, the computation time for secure aggregation per entry (e.g., homomorphic encryption, PRG expansions, or modular exponentiations) tends to be **significantly greater compared to the computation time per entry in linear transformations**. This suggests that PVF still manages to compress computation overhead, which is evident in the experiment outcomes.

In practice, the choice of λ mainly considers: (i) Security requirements in DVE (see Lemma 3). (ii) SAP. For schemes with more masking-related overhead, a larger λ performs better. For schemes with more interaction-related overhead, a smaller λ performs better (as in Section 5.3). (iii) m . For larger m , masking-related overhead is greater, so a larger λ performs better.

Table 4: Theoretical complexity of SAP without and with PVF of single-round aggregation in the semi-honest setting. $O(\cdot)(\cdot)$ in the rightmost column indicates the communication complexity pertains to users, followed by the number of interactions between users and the server. ‘‘P. of G. M.’’ stands for privacy of the global model. In the real world, $p \gg m \gg n$. \uparrow indicates the integration of PVF into the protocol would increase its theoretical complexity.

SAP	Each user		Server		P. of G. M.	Communi. (Inter.)
	w/o PVF	w/ PVF	w/o PVF	w/ PVF		
PPDL	$O(m)$	$O(\lambda m) \uparrow$	$O(mn)$	$O(\frac{1}{\lambda}mn + \lambda m)$	\checkmark	$O(m)$ (1)
EPPFL	$O(m)$	$O(\lambda m) \uparrow$	$O(mn)$	$O(\frac{1}{\lambda}mn + \lambda m)$	\times	$O(m)$ (2)
NIVP-DS	$O(m)$	$O(\lambda m) \uparrow$	$O(mn)$	$O(\frac{1}{\lambda}mn + \lambda m)$	\checkmark	$O(m)$ (1)
PracAgg	$O(mn + n^2)$	$O(\frac{1}{\lambda}mn + \lambda m + n^2)$	$O(mn^2)$	$O(\frac{1}{\lambda}mn^2 + \lambda m)$	\times	$O(m + n)$ (4)
PracAgg+	$O(m \log n + \log^2 n)$	$O(\frac{1}{\lambda}m \log n + \lambda m + \log^2 n)$	$O(mn \log n + n \log^2 n)$	$O(\frac{1}{\lambda}mn \log n + \lambda m + n \log^2 n)$	\times	$O(m + \log n)$ (4)
EffiAgg	$O(m + n^2)$	$O(\lambda m + n^2) \uparrow$	$O(n\sqrt{p} + n)$	$O(\frac{1}{\lambda}m\sqrt{p} + \lambda m + n)$	\times	$O(m + n)$ (4)
LPPFedL	$O(m + n^2)$	$O(\lambda m + n^2) \uparrow$	$O(m + n)$	$O(\lambda m + n) \uparrow$	\times	$O(m + n)$ (4)

E.2 MORE IMPLEMENTATION DETAILS

We implement the baselines using Python. Specifically, we utilize *AES-GCM* with 128-bit keys for the symmetric authenticated encryption, standard (t, n) *Shamir Secret Sharing* (Shamir, 1979), *AES* in counter mode for the pseudorandom generator, *SHA-256* hash to implement a homomorphic pseudorandom generator for EffiAgg, *Sympy* library (Meurer et al., 2017) to compute discrete logarithms for EffiAgg, and *Paillier Encryption* with 1024-bit keys for PPDL. In Section 3.3, experimental settings of the image classification task is the same as Figure 2, and for the movie recommendation task, the dataset is split using a *Leave-One-Out* (He et al., 2017) approach for training and testing, where users with fewer than 10 records are excluded, using *Hit Ratio* (HR) as metrics to assess the performance of the recommendations, with higher values indicating superior effectiveness.

In Figure 2, we use three widely discussed SAPs: PracAgg+, PracAgg, and PDDL. NIVP-DS falls behind due to the requirement of two non-colluding servers, while LPPFedL demands increased communication overhead. Remarkably, PDDL, previously considered impractical due to its use of HE, has its computational overhead significantly mitigated by integrating PVF. And its advantages of operating with a single server, single-interaction communication, and preserving the global model make it more appealing.

E.3 FOR FL OF LLM

LLMs have a profound impact on the entire AI research community due to their excellent contextual learning and instruction following ability (Zhao et al., 2023). Given privacy concerns, training (or fine-tuning) LLM in a federated setting has been explored (Hilmkil et al., 2021; Ye et al., 2024). In the context of LLMs, during the aggregation process, the length of user original vectors reaches billions. Taking Llama2-7B (Touvron et al., 2023) as an example, assuming 1% of the parameters need to be updated during the fine-tuning, which is 700M, the computational cost of using a general secure aggregation scheme is unimaginable, making PVF particularly important. As shown in Table 5, the time required to train one round without using PVF can basically meet the requirement of training 100 rounds with PVF. At present, FL for LLM among a multitude of lightweight clients is unrealistic. The SOTA scheme OpenFedLLM (Ye et al., 2024) involves only $n \in \{2, 4, 5\}$ clients per round (using Llama2-7B). And PVF can contribute to future FedLLMs and involving more clients.

Table 5: Estimated computational overhead per round with and without PVF for fine-tuning Llama2-7B. $n = 100$, $m = 700M$, $\lambda = 100$, and $\eta = 0\%$.

Scheme	Each user		Server	
	w/o	w/	w/o	w/
PDDL	~300h	~3h	~60h	~0.6h
PracAgg	~27h	~0.27h	~27h	~0.27h
PracAgg+	~5h	~0.05h	~27h	~0.27h

E.4 THE IMPACT OF PADDING

PVF necessitates padding the original vector, thereby increasing the vector size. In theory, the additional cost that padding introduces in computation and communication does not exceed $\frac{\lambda}{m}$ of the original, as **the maximum value of padding length** is λ . Table 6 showcases the impact of no padding versus padding λ entries on the computation and communication overhead, where $\frac{\lambda}{m} = 0.001$. It’s evident that the overhead induced by padding is almost negligible. To strictly adhere to the principle of “**not increasing any communication overhead**”, we present a method to avoid padding. We extract the first $\lfloor \frac{m}{\lambda} \rfloor \lambda$ entries from the original vector and apply PVF to them. The remaining entries, which are fewer than λ , are appended to the key vector \mathbf{k} for participation in the secure aggregation. This approach allows us to obtain the aggregate result of the entire original vector while eliminating the need for padding.

Table 6: Comparison of overhead with and without padding. $n = 100$, $m = 100k$, $\lambda = 100$, and $\eta = 10\%$.

Scheme	User comp. (ms)		Server comp. (ms)		Comm. cost (KB)	
	No pad	Pad	No pad	Pad	No pad	Pad
PDDL	1402	1403, $\uparrow 1$	303	304, $\uparrow 1$	859	860, $\uparrow 1$
EPPFL	41	42, $\uparrow 1$	7337	7351, $\uparrow 14$	681	681, $\uparrow 0$
NIVP-DS	13	13, $\uparrow 0$	10	10, $\uparrow 0$	587	587, $\uparrow 0$
PracAgg	162	164, $\uparrow 2$	1456	1463, $\uparrow 7$	785	785, $\uparrow 0$
PracAgg+	38	39, $\uparrow 1$	405	409, $\uparrow 4$	783	783, $\uparrow 0$
EffiAgg	29	32, $\uparrow 3$	5404	5482, $\uparrow 78$	783	783, $\uparrow 0$
LPPFedL	19	19, $\uparrow 0$	20	20, $\uparrow 0$	1564	1564, $\uparrow 0$

E.5 MEMORY USAGE

In Figure 9, we evaluate the memory usage of each user and \mathcal{S} during PracAgg without PVF ($\lambda = 0$) and with PVF ($\lambda = 100, 1000$).

For each user, the increased memory usage for $\lambda = 1000$ compared to $\lambda = 100$ is due to the need to store a larger transformation matrix ($\lambda \times \lambda$) required by PVF. The increase in memory usage for $\lambda = 0$ compared to $\lambda = 100$ is because PVF reduces the number of random numbers generated ($nm \rightarrow \frac{nm}{\lambda}$) and decreases the scale of vector addition computations ($m \rightarrow \frac{m}{\lambda}$) during the masking process.

For \mathcal{S} , all three methods require summing up the vectors uploaded by all users, with a memory usage of approximately $O(mn)$, so the additional $O(\lambda^2)$ overhead introduced by PVF is negligible. The increased memory usage for $\lambda = 0$ compared to $\lambda = 100$ is due to PVF reducing the number of random numbers generated ($(\eta(1-\eta)n^2 + (1-\eta)n)m \rightarrow \frac{(\eta(1-\eta)n^2 + (1-\eta)n)m}{\lambda}$) and decreasing the scale of vector addition computations during the unmasking process.

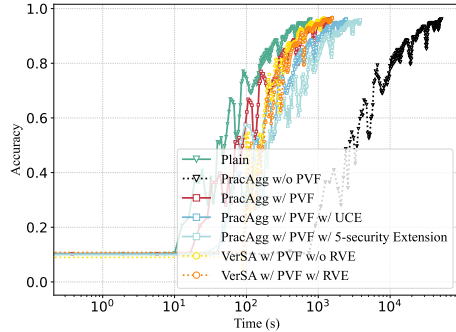



Figure 9: Memory usage with different λ . $n = 100$, $m = 100k$, and $\eta = 10\%$. Figure 10: Comparison of various aggregation methods, with $n = 100$, $\eta = 5\%$, $\lambda = 100$.

E.6 EVALUATION OF EXTENSIONS

End-to-end comparison. In accordance with the experimental settings of Figure 2, we evaluate the impact of RVE, UCE, and μ -security on the overall model training time (as illustrated in the Figure 10). The acceleration effect and communication expansion when integrating different extensions are $34.5\times$ and $1.4\times$ (RVE), $21.1\times$ and $2.9\times$ (UCE), $13.5\times$ and $1.0\times$ (5-security).

(λ, μ) . We evaluate the speedup with different $(\lambda, \mu) \in \{100, 300, 500, 700, 1000\} \times [1, 10]$. The integration of PVF with the μ -security extension reduces the entries of vectors involved in secure aggregation to $\frac{\mu+1}{\lambda}$ ($\mu < \lambda - 1$) of their original size. Figure 11 showcases incorporating the μ -security extension does diminish the improvement factor. However, even when $\mu = 0.1\lambda$, PVF still yields an acceleration gain of $10\times$ along with communication improvements exceeding $5\times$.

RVE and UCE. UCE necessitates users to commit to each dimension of the original vectors, while \mathcal{S} needs to validate each dimension. RVE requires users to submit frozen vectors twice and \mathcal{S} to perform summations twice. Thereby, both UCE and RVE impose on each participant an additional communication complexity of $O(m)$ and computation complexity of $O(m)$.

Figure 12 presents the overhead required for SAPs with different extensions. Compared to not integrating PVF, the computation overhead after adding extensions is still nearly **an order of magnitude lower**. Communication costs from extensions primarily stem from the transmission of additional vectors, such as commitment vectors. We leave mitigating these overheads to future work.

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

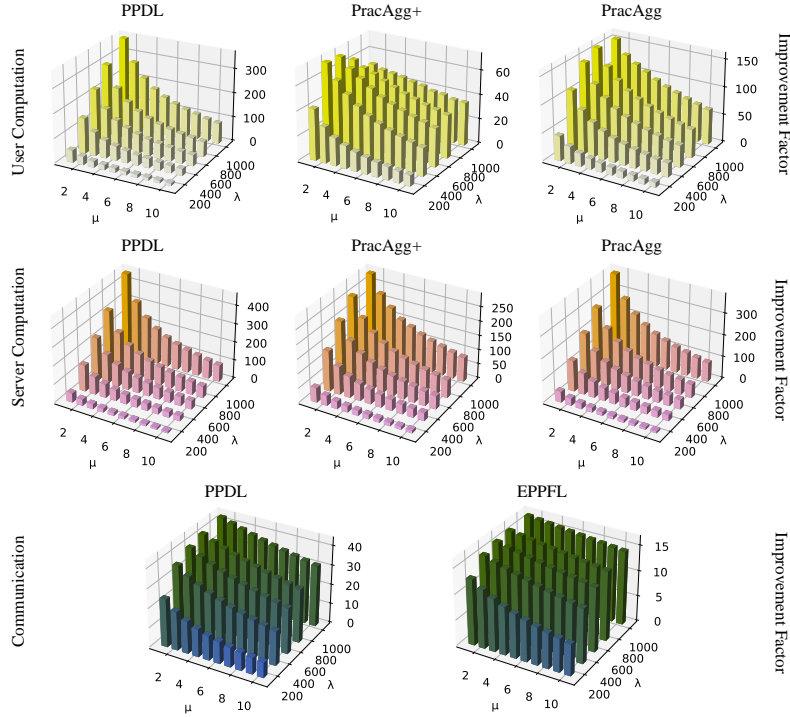


Figure 11: The impact of different (λ, μ) on improvement factor. $n = 100$, $m = 100k$, $\lambda = 100$, and $\eta = 10\%$.




Figure 12: Comparison of costs across different extensions. $n = 100$, $m = 100k$, $\lambda = 100$, and $\eta = 10\%$.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

Participants: \mathcal{S} and User set $\mathcal{U} = \{u_1, u_2, \dots, u_n\}$.
Public Inputs: $A, \mu, A^{\mu+1}, \alpha^{\mu+1}, \lambda, \mathbb{Z}_p, g$ and h . Users' public keys for signatures $\{sig_i^{pk}\}_{i \in \mathcal{U}}$ and the server's public key for signatures sig_S^{pk} .
Private Inputs: Original vectors $\{x^{i(t)}\}_{i \in \mathcal{U}}$ of t -th iteration. Users' secret keys for signatures $\{sig_i^{sk}\}_{i \in \mathcal{U}}$ and the server's secret key for signatures sig_S^{sk} .
Outputs: Surviving user set \mathcal{U}' , $\sum_{i \in \mathcal{U}'} x^{i(t)}$.
• Phase 1 Freezing
 User $i \in \mathcal{U}$:
 - pad $x^{i(t)}$ randomly and group the entries.
 - add noise to $x^{i(t)}$ via Equation (10).
 - calculate key vector $k^{i(t)}$ via Equation (14).
 - calculate frozen vector $y^{i(t)}$ via Equation (8) and Equation (13).
 - generate random vector $\zeta^{i(t)} = (\zeta_1^{i(t)}, \zeta_2^{i(t)}, \dots, \zeta_\lambda^{i(t)})$ and calculate commitment vector $c^{i(t)}$ via Equation (16).
 - obtain $m_1^i = y^{i(t)} \parallel (\text{Do not send } y^{i(t)} \text{ when there is RVE}) \parallel \zeta^{i(t)} \parallel c^{i(t)}$, send $\sigma_1^i \rightarrow DS.sign(sig_i^{sk}, m_1^i)$ to \mathcal{S} .
• Phase 2 SecAgg
 \mathcal{S} and Users:
 - execute SAP for $\{k^{i(t)}\}_{i \in \mathcal{U}}$.
 * ...
 * users get (κ_1, κ_2) and obtain $m_2^i = \hat{y}^{i(t)}(\kappa_1, y^{i(t)}) \parallel \hat{y}^{i(t)}(y^{i(t)} + \kappa_2)$, send $\sigma_2^i \rightarrow DS.sign(sig_i^{sk}, m_2^i)$ to \mathcal{S} .
 * ...
 - all participants receive \mathcal{U}' and $\sum_{i \in \mathcal{U}'} k^{i(t)}$ (or $Enc(\sum_{i \in \mathcal{U}'} k^{i(t)})$).
 \mathcal{S} :
 - if $DS.verify(\sigma_1^i, sig_i^{pk}, m_1^i) \rightarrow False$, abort. Otherwise, calculate $\sum_{i \in \mathcal{U}'} y^{i(t)}$ (can not get $\sum_{i \in \mathcal{U}'} y^{i(t)}$ calculate $\sum_{i \in \mathcal{U}'} \hat{y}^{i(t)}$ and $\sum_{i \in \mathcal{U}'} \hat{y}^{i(t)}$).
 - for $j \in [1, \lambda]$, $r \in [1, \lambda]$, reveal the commitments via Equation (17). If verification fails, abort.
• Phase 3 Thawing
 Thawing on the server side
 \mathcal{S} :
 - calculate $sum = \sum_{i \in \mathcal{U}'} x_{pad}^{i(t)}$ via Equation (9) and send sum and $\sigma_3 \rightarrow DS.sign(sig_S^{sk}, sum)$ to $i \in \mathcal{U}'$.
 User $i \in \mathcal{U}'$:
 - receive $\sum_{i \in \mathcal{U}'} x^{i(t)}$, if $DS.verify(\sigma_3, sig_S^{pk}, sum) \rightarrow False$, abort. Otherwise, unpad and output.
 Thawing on the user side
 \mathcal{S} :
 - send $\sum_{i \in \mathcal{U}'} y^{i(t)}$, $Enc(\sum_{i \in \mathcal{U}'} k^{i(t)})$ and $\sigma_3 \rightarrow DS.sign(sig_S^{sk}, \sum_{i \in \mathcal{U}'} y^{i(t)} \parallel Enc(\sum_{i \in \mathcal{U}'} k^{i(t)}))$ to $i \in \mathcal{U}'$.
 User $i \in \mathcal{U}'$:
 - if $DS.verify(\sigma_3, sig_S^{pk}, \sum_{i \in \mathcal{U}'} y^{i(t)} \parallel Enc(\sum_{i \in \mathcal{U}'} k^{i(t)})) \rightarrow False$, abort. Otherwise, verify $\sum_{i \in \mathcal{U}'} y^{i(t)}$ via Equation (15). If verification fails, abort.
 - decrypt $Enc(\sum_{i \in \mathcal{U}'} k^{i(t)})$.
 - calculate $sum = \sum_{i \in \mathcal{U}'} x_{pad}^{i(t)}$ via Equation (9), unpad and output.

Figure 13: The pipeline of λ -SecAgg with all 4 extensions for one aggregation. The red and underlined parts are required in the user commitment extension. The blue and underlined parts are required in the result verification extension.