

Non-Interactive Verifiable Aggregation

Ojaswi Acharya
University of Massachusetts Amherst
oacharya@umass.edu

Suvasree Biswas
The George Washington University
suvasree@gwu.edu

Weiqi Feng
University of Massachusetts Amherst
weiqifeng@umass.edu

Adam O'Neill
University of Massachusetts Amherst
adamoneill@umass.edu

Arkady Yerukhimovich
The George Washington University
arkady@gwu.edu

Abstract

Consider a weak analyst that wishes to outsource data collection and computation of aggregate statistics over a potentially large population of (also weak) clients to a powerful server. For flexibility and efficiency, we consider public-key and non-interactive protocols, meaning the clients know the analyst's public key but do not share secrets, and each client sends at most one message. Furthermore, the final step should be silent, whereby the analyst simply downloads the (encrypted) result from the server when needed. To capture this setting, we define a new primitive we call Non-Interactive Verifiable Aggregation (NIVA).

We require both privacy and robustness for a NIVA protocol to be deemed secure. Namely, our security notion for NIVA ensures that the clients' data remains private to both the server and the analyst, while also ensuring that malicious clients cannot skew the results by providing faulty data.

We propose a secure NIVA protocol, which we call PEAR (for Private, Efficient, Accurate, Robust), which can validate inputs according to any NP validity rule. PEAR is based on a novel combination of functional encryption for inner-products (Abdalla et al., PKC 2015) and fully-linear probabilistically-checkable proofs (Boneh et al., Crypto 2019). We emphasize that PEAR is non-interactive, public-key, and makes black-box use of the underlying cryptographic primitives.

Additionally, we devise substantial optimizations of PEAR for practically-relevant validity rules. Finally, we implement PEAR to show feasibility for such validity rules, conducting a thorough performance evaluation. In particular, we compare PEAR to two more straightforward or "off-the-shelf" NIVA protocols and show performance gains, demonstrating the merit of our new approach. The bottleneck in our protocol comes from the fact that we require the underlying IPFE scheme to be "unrestricted" over a large field. As more efficient such schemes are developed, they can be immediately plugged into PEAR for further gains.

Keywords

Secure aggregation, Input validation, Inner-product functional encryption, Fully-linear PCPs

1 Introduction

1.1 Background and Motivation

MOTIVATING APPLICATIONS. There are numerous applications in which an 'analyst' wishes to conduct a study on a large population of 'clients' to learn aggregate statistics, but input collection from the clients is too burdensome for the analyst. Therefore, the analyst wishes to *outsource* the collection and aggregation of the data to a (potentially untrusted) 'server' who aggregates the inputs and returns the result to the analyst. Here are some examples:

- **Population studies:** Suppose that think tanks or policymakers wish to collect statistics such as average age, *etc.* about a certain segment of the population. The think tank may not have the intensive computational resources to collect all the survey data, and may choose to outsource this to the cloud. The cloud server can then collect and aggregate the data and give the result to the think tank.
- **Sensor networks:** Suppose a researcher wishes to deploy sensors to measure the environment (*e.g.*, temperature readings) in some remote setting. Rather than set up their own network to collect the readings, the researcher can use a cloud server to collect and aggregate the readings before providing a summary. Note that as the sensors may be placed in sensitive locations, privacy of the collected readings may need to be protected.
- **Pharmaceutical research:** Pharmaceutical companies may wish to collect statistics about genomic sequences of certain cohorts of hospital patients. Rather than sending the data directly, the hospitals may instead upload data to a server and want to ensure that the pharmaceutical company only gets aggregate information rather than data about any individual patient.
- **Disease outbreak detection:** To help detect disease outbreaks, disease control centers would like to collect aggregate statistics about symptoms and diagnoses from hospitals' patient records. Similar to the pharmaceutical research use-case, aggregating this data requires a server to produce the aggregate without leaking the individual patient data.

Abstracting away the details of the above examples, we are concerned with outsourcing the learning of aggregate statistics about *sensitive* data. Thus, to ensure privacy, we want the clients' inputs to remain private from other clients and from the server, as well as from the analyst (to the extent possible given that the analyst learns the final result of the computation). In tandem, it is crucial to defend against *data poisoning*, where some clients' data is corrupted

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies YYYY(X), 1–20

© YYYY Copyright held by the owner/author(s).

<https://doi.org/XXXXXXXXXXXXXX>



either due to unintentional error or intentional malice, which is a well-known concern. To address this, we want the server to be able to verify that the clients' inputs satisfy some analyst-specified validity rule before aggregating them. Of course, this validation check must still preserve the privacy of the inputs.

Crucially for the applications we consider, we want anybody to be able to provide input to the aggregation – i.e., to serve as a client. To enable this, we operate in a *public key* setting where we assume the clients can obtain the authentic public keys of the analyst and server, but do not share any secret information with either. To achieve all of these goals, we propose a new primitive called (public-key, outsourced) *Non-Interactive Validated Aggregation* (NIVA). An overview of the NIVA functionality is given in Figure 1.

Before outlining the performance and security requirements of NIVA, we wish to highlight one critical design goal.

ON THE IMPORTANCE OF BEING NON-INTERACTIVE. We require that all NIVA schemes be non-interactive meaning that each client sends *a single message* to the server to provide input for the aggregation and after validating and aggregating the inputs, the server sends one message to the analyst. No additional rounds of communication are allowed either among the clients or between the clients and the server or analyst.

For the case of client-server communication, such non-interactive protocols are crucial in settings of low connectivity such as sensor networks. While client devices may be able to perform local computation, they can only communicate with the server when online and need to do so in short bursts without waiting for a response from the server.

Moreover, by not requiring any interaction NIVA allows for clients to remain unknown to the server (and analyst) prior to submitting input. This is critical in settings such as public population surveys allowing anybody to participate.

We note that this restriction to unknown clients does place some limitation on the security that can be achieved against a malicious server. Since the analyst does not know the set of clients, a malicious server can pretend to be a client adding inputs of its choice into the aggregation or drop clients who have supplied valid inputs. Similar to prior work (e.g., [18, 33]), we thus only aim for privacy, but not integrity, against a malicious server.

For the case of server-analyst interaction, prohibiting communication between the server and analyst ensures that the analyst does not need to be online during data collection and only needs to receive the results, which can be stored for him.

1.2 NIVA Performance Requirements

To achieve the goals outlined above, NIVA has a unique set of performance requirements that we detail below.

- The analysts' cost must be independent of the number of clients. In fact, we require that the analyst not even know the identity of the clients, thus precluding any sort of secret-key setup between them.
- We require that each client's computation and communication be independent of the number of clients. Moreover, a client should only send a *single message* to the server and have no communication directly with the analyst or other

clients. Thus, clients cannot pre-register before providing input.

- The server should run the most computationally intensive component of our protocol, namely validation and aggregation of the clients' inputs. Thus, the server's computation and communication grows linearly in the number of clients. But, we require that both of these steps be non-interactive (i.e., performed locally by the server).
- The validity rule must not depend on the underlying cryptographic tools. This means that when proving the validity of their inputs, the clients must prove the statement "my input satisfies the validity condition" rather than "my ciphertext is an encryption of a valid input."

1.3 NIVA Security Requirements

To meet the needed security for a secure (outsourced) aggregation protocol, NIVA must satisfy the following security conditions.

- A malicious server must not learn the content of the clients' inputs except for whether or not they satisfy the validity constraint. Additionally, she should learn nothing about the final aggregate result unless it is made public by the analyst. This is true even if she colludes with a subset of the clients. As discussed previously, we only guarantee integrity (i.e., that the analyst gets the correct output) against a semi-honest server.
- Malicious clients who do not collude with the server must not be able to get the server to accept any inputs that do not satisfy the validity rule. This must be true even if all clients are malicious. That is, we guarantee correctness against malicious clients. Moreover, malicious clients should not be able to block or tamper with the input of an honest client.
- A semi-honest analyst, who does not collude with the server and thus is assumed to not see the clients' (plaintext) communication with the server, should learn nothing about individual client inputs other than what is revealed by the final aggregated information. In particular, he should not learn which of the clients provided valid or invalid inputs. In fact, we achieve a somewhat stronger guarantee that as long as the analyst runs the Setup procedure honestly¹, then security is maintained even against an otherwise malicious analyst.

1.4 Technical Overview

We propose a NIVA protocol supporting summation, which we call PEAR (for Private, Efficient, Accurate, Robust). Summation allows for a wide range of applications, in particular all those supported by PRIO [19] using their "affine aggregatable encodings." Moreover, PEAR achieves all of our desiderata described above. In particular, to our knowledge it is the first system for non-interactive aggregation with a single server and input validation where the validity rule does *not* depend on the underlying cryptographic primitives.

INTUITION. The design of PEAR is based on a novel combination of additively-homomorphic inner-product functional encryption

¹Note that this procedure is run only once before any clients submit their inputs.




Figure 1: An overview of NIVA functionality.

(IPFE) [1] and fully-linear probabilistically-checkable proofs (FL-PCPs) [12]. In IPFE, anyone with mpk can encrypt a vector x . Furthermore, a secret key sk_y associated to a vector y can be generated using msk . While msk can decrypt ciphertexts as usual, an evaluation algorithm run on ct_x and sk_y yields $\langle x, y \rangle$ in the clear. Our key observation is that the verification procedure of an FL-PCP consists of a secret linear step on the instance x and proof π followed by a non-linear step on the result. As proven by [12], the output of the linear step is in fact zero-knowledge, allowing for the non-linear step to be performed in the clear. Thus, to preserve privacy, it is sufficient to perform the secret linear step under the hood of the IPFE. Moreover, existing IPFE schemes are additively homomorphic (although this property is not usually exploited in prior work), allowing to then homomorphically sum up the valid input vectors.

However, for this to work with the FL-PCP, we need IPFE supporting “unrestricted” inner-product, without a polynomial upper-bound on its magnitude. While initial schemes did not have this property as they recover inner-product in the exponent, we use a scheme based on class groups by Castagnos *et al.* [17], which exploits a subgroup where computing discrete log becomes easy.

OUR PROTOCOL. In more detail, PEAR works as follows for input-validity specified by some language $\mathcal{L} \in \text{NP}$:

- (1) The analyst generates keypair (mpk, msk) for IPFE and publishes mpk . Additionally, he chooses ℓ random vectors $q_1, \dots, q_\ell$ (which will serve as FL-PCP query points) and generates IPFE secret keys $sk_1, \dots, sk_\ell$ for each of these vectors, then sends them privately to the server.
- (2) On input $x_i \in \mathcal{L}$ and a witness w_i for x_i , client C_i generates an FL-PCP proof π_i to prove that $x_i \in \mathcal{L}$ and encrypts $x_i || \pi_i$ under IPFE public-key mpk to get ciphertext ct_i , which it sends to the server.
- (3) Upon receiving ciphertext ct_i , the server computes IPFE evaluation on ct_i with each of $sk_1, \dots, sk_\ell$ to perform the linear step of the FL-PCP verification procedure. Following this, the server can run the remainder of the procedure locally to determine whether $x \in \mathcal{L}$. If so, the server stores ct_i and otherwise discards it.

- (4) Upon receiving an *a priori* agreed-upon number (n) of ciphertexts, the server uses the additive homomorphism of the IPFE to aggregate all the valid ciphertexts producing a single aggregate ciphertext ct_{agg} encrypting the corresponding sum, which it then sends to the analyst.

- (5) The analyst uses msk to decrypt ct_{agg} and recover the result.

We remark that the above can easily be extended to *weighted* sum aggregation assuming the IPFE supports homomorphic multiplication-by-constant, as is the case for the scheme of [17].

1.5 Comparison to Other Approaches for NIVA

Here we discuss other more straightforward or “off-the-shelf” approaches to building a secure NIVA protocol and why our new approach is superior.

NIVA BASED ON HE+NIZK. A secure NIVA protocol can also be constructed in a generic manner from additively homomorphic encryption (HE) and non-interactive zero-knowledge proofs (NIZKs) for languages in NP [22]. The protocol works as follows. Fix an analyst-specified NP language $\mathcal{L}$. First, the analyst generates a keypair (pk, sk) for HE and broadcasts pk . Given input $x_i \in \mathcal{L}$ with witness w_i , client C_i encrypts x_i under pk to get ct'_i and attaches to it a proof π_i that ct'_i encrypts a member of $\mathcal{L}$. Upon receiving $ct_i = ct'_i || \pi_i$, the server verifies π_i and discards ct'_i if this check fails. The server then homomorphically aggregates the remaining ct'_i and sends the result to the analyst, who decrypts it under sk .

However, this protocol is *non-blackbox* because the statement proven in ZK depends on the encryption circuit. While it is certainly possible to design optimized encryption schemes and proofs to reduce the overhead of these proofs, as done in some prior work — e.g., [7] — this limits the flexibility of the resulting NIVA protocol as it is not possible to replace the encryption scheme by a more efficient one developed down the line without reworking the proof. In particular, since optimizing performance is not necessarily the same as optimizing the “proof-cost” of a cryptographic primitive, improvements in building blocks may not yield improvements in the resulting construction. Optimizations to the protocol notwithstanding, to better demonstrate the overhead of the above non-blackbox solution we implement “an off-the-shelf” version of it based on Bulletproofs [15] and (the additively-homomorphic version of) El

Scheme	PEAR(Our Work)	ACORN [7]	RoFL [33]	EiFFeL [18]
Number of rounds	1	≥ 5	≥ 5	≥ 5
Malicious clients threshold	$n - 1$	$n/3$	t neighbors	$(n - 1)/3$
Client computation	$O(1)$	$O(\log n)$	$O(\log n)$	$O(\log n)$
Server computation	$O(n)$	$O(n \log n)$	$O(n)$	$O(n^2)$
Proof system	FL-PCP	Bulletproofs	Bulletproofs	SNIP
Key setup	public	pairwise	pairwise	pairwise

Table 1: Comparing our NIVA protocol PEAR to SA protocols with input validation.

Gamal encryption [21]. We then compare its performance to PEAR, demonstrating significant gains. See Section 6 for details.

NIVA BASED ON FHE+FL-PCP. As another point of comparison, we consider a secure NIVA protocol based around *fully*-homomorphic encryption (FHE) [23]. The idea is that we use an FL-PCP as before but the entire validation of the ciphertexts and aggregation of the valid ones is done completely “under the hood” of FHE. Note that while this approach is black-box in the underlying cryptography, it does *not* take advantage of the fact that the outputs of the inner products in the validation circuit is ZK and the remainder of the computation can be done in the clear. Indeed, we demonstrate impracticality of this FHE-based NIVA protocol for an “off-the-shelf” implementation of it using the SEAL library [40]. Specifically, verifying the FL-PCP validation result under the hood of FHE, without client assistance, is highly expensive, and the server must repeat this task for each client ciphertext. Our approach is to exponentiate the FL-PCP validation bit to the size of the plaintext space, ensuring that zero remains zero while all other values become one, thereby filtering out invalid ciphertexts by flipping the bit. See Section 6 for further details.

1.6 Adding Differential Privacy

As our protocol is described above, the analyst learns nothing about clients’ inputs except what is revealed by the output. Of course, this does not necessarily suffice to guarantee the privacy of clients’ inputs since the output may reveal information about such inputs.

A standard way to guarantee that the output of the aggregation does not reveal too much about the individual inputs is to guarantee differential privacy (DP) [20]. In the typical case of input-validity checks for bounded norm, DP can be added to PEAR. The simplest way to add DP to PEAR is to have the server add DP noise by choosing the noise based on the global sensitivity of the aggregation (e.g., if inputs are bounded then so is the sensitivity). The server can then encrypt the noise using the IPFE key and add it to the final aggregation result using linear homomorphism. The analyst then only gets noisy results, thus preserving DP. Since we see DP as being largely orthogonal to the design of our protocol, we do not elaborate on it further in the remainder of the paper.

1.7 Performance Evaluation

We implement PEAR using the unrestricted IPFE scheme by Castagnos et al. [17] and the FL-PCP by Boneh et al. [12]. We construct two validation circuits: one for verifying whether an input is a binary number and another for ensuring that the L_2 norm of an input vector is bounded. We benchmark PEAR and compare against a naive

approach that combines off-the-shelf additively-homomorphic encryption with non-interactive zero-knowledge proofs of ciphertext validity as described earlier. In terms of client runtime, PEAR achieves a speedup of $2.8\times$ to $11\times$, while for server runtime, it achieves a speedup of $1.2\times$ to $3.6\times$. The server-side speedup could be significantly improved if the performance of the unrestricted IPFE scheme were comparable to that of restricted IPFE schemes, such as the optimized one proposed by Kim *et al.* [31]. We leave the development of a more efficient unrestricted IPFE scheme as an intriguing open problem. We also compare PEAR with the approach of using FHE with FL-PCP and observe that its server runtime for the binary number validation is $20\times$ to over $100\times$ slower.

2 Related Work

SECURE AGGREGATION. Secure aggregation, introduced by [11], allows a server to privately sum the inputs of a potentially large number of clients. These protocols use secret-shared masks between clients to enable efficient aggregation, *without* verifying validity of inputs. In SA, the server learns the sum; there is no separate analyst. Since then, various works have also considered how to extend these protocols to build *verifiable* secure aggregation as in our setting. In the single-server setting, various protocols have been proposed that require several rounds of communication between the server and clients [7, 18, 33]. Another line of work shows how to achieve secure aggregation in a multi-server setting relying on (interactive) secure computation between the servers [3, 6, 13, 19, 38, 41]. There is also another line of work that adds additional roles to some distinguished clients and necessitates additional communication between any client and these distinguished ones. An example is the committee-based setting of [8, 29, 32]. In comparison, PEAR does not require that any of the clients be known to anyone else. In Table 1, we compare PEAR against popular single-server, verifiable SA protocols. Here n is number of clients. Of course, compared to existing single-server, verifiable SA protocols, PEAR operates in a setting where there is an additional analyst, who is separate from the server. Besides being natural from an application standpoint, we chose this modification because there are fundamental limitations to the security of such non-interactive protocols when the server learns the aggregation result herself. Even without public keys, the server can perform a so-called ‘residual attack’ if it corrupts clients. In a residual attack, the server locally re-runs the protocol many times, changing the corrupted clients’ inputs each time, to learn more and more about other clients’ inputs. In our public-key setting the limitations are even more severe, because the server does not need to corrupt any clients to run such attacks. By having separate server and analyst we simply avoid such issues. Note that

[29] (which does not consider input validation) avoids them by having a committee in addition to a server.

EFFICIENT ZERO-KNOWLEDGE PROOFS. Our protocol crucially relies on efficient non-interactive zero-knowledge (NIZK) proofs to prove validity of the clients' inputs. Zero-knowledge proofs were originally introduced by Goldwasser *et al.* [24] with the non-interactive variant introduced by Feige *et al.* [22]. Since then, a number of works (e.g. [5, 15, 25, 30, 34, 37]) have investigated how to make such proofs shorter and more efficient. In PEAR, we require proofs that allow proving validity of an input unknown to the verifier. For this purpose, we build off of the work of Boneh *et al.* [12], who consider proofs over secret-shared inputs.

INNER PRODUCT FUNCTIONAL ENCRYPTION (IPFE). Another crucial tool for our construction is inner-product functional encryption (IPFE) which allows computing the inner product between an encrypted vector and a vector in the (functional) evaluation key. IPFE was first introduced by Abdalla *et al.* [1], who showed simple constructions based on DDH and LWE. Further variants of IPFE were given by e.g. [2, 4, 10]. In PEAR, we use the IPFE scheme due to Castagnos *et al.* [17] based on the hard subgroup membership (HSM) assumption in certain class groups. Importantly for us, this IPFE scheme is *unrestricted* in that it allows efficient evaluation when an inner product may output an *arbitrary* field element, rather than restricting the output to small set that can be brute-forced. This is needed for compatibility with the FL-PCP.

3 Preliminaries

3.1 Notation and Conventions

If $\mathbf{v}$ is a vector then $|\mathbf{v}|$ is its length (the number of its coordinates) and $\mathbf{v}[i]$ is its i -th coordinate. Strings are identified with vectors over $\{0, 1\}$, so that $|Z|$ denotes the length of a string Z and Z_i denotes its i -th bit. By ε we denote the empty string or vector. By $x\|y$ we denote the concatenation of strings x, y . For two vectors $\mathbf{x}, \mathbf{y}$ of the same length over the same field, we denote their inner-product by $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i \in [|\mathbf{x}|]} \mathbf{x}_i \cdot \mathbf{y}_i$.

If S is a finite set, then $|S|$ denotes its size. If X is a finite set, we let $x \leftarrow X$ denote picking an element of X uniformly at random and assigning it to x . Algorithms may be randomized unless otherwise indicated. Running time is the worst case, which for an algorithm with access to oracles means across all possible replies from the oracles. We use $\perp$ (bot) as a special symbol to denote rejection, and it is assumed to not be in $\{0, 1\}^*$.

A function $\nu: \mathbb{N} \rightarrow \mathbb{N}$ is *negligible* if for every positive polynomial $p: \mathbb{N} \rightarrow \mathbb{R}$ there is a $\lambda_p \in \mathbb{N}$ such that $\nu(\lambda) \leq 1/p(\lambda)$ for all $\lambda \geq \lambda_p$. "PT" stands for "polynomial time," whether for randomized or deterministic algorithms. By 1^λ we denote the unary representation of the integer security parameter $\lambda \in \mathbb{N}$.

GAMES. We use the code-based game-playing framework of BR [9]. By $\Pr[G \Rightarrow y]$ we denote the probability that the execution of game G results in this output being y . In games, integer variables, set variables, boolean variables, and string variables are assumed initialized, respectively, to 0, $\emptyset$, false, and $\perp$. For an adversary $\mathcal{A}$ playing game G , we may write another adversary $\mathcal{B}$ in the same format as G , with the understanding that $\mathcal{B}$ runs this game with $\mathcal{A}$.

Variable name	Meaning
λ	security parameter
$\mathbf{x}$	input instance vector
d	input instance vector length
$\mathbf{w}$	witness vector
h	witness vector length
π	proof vector
m	proof length
$\mathbf{q}$	FL-PCP query
ℓ	number of FL-PCP queries

Table 2: Notation used throughout the paper.

3.2 Fully-Linear PCP Proof Systems

BINARY RELATIONS. For $d, h \in \mathbb{N}$, let $\mathcal{R} \subseteq \mathbb{F}^d \times \mathbb{F}^h$ be a binary relation over a finite field $\mathbb{F}$. Additionally, let $\mathcal{R}_p \subseteq \mathbb{F}_p^d \times \mathbb{F}_p^h$ be a family of binary relations over finite fields $\mathbb{F}_p$ parameterized by the field size p . For simplification, we drop the subscript p once it is fixed. We define the corresponding language $\mathcal{L}_{\mathcal{R}}$ such that $\mathbf{x} \in \mathcal{L}_{\mathcal{R}}$ if and only if there exists an $\mathbf{w} \in \mathbb{F}^h$ such that $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$. We also define an associated arithmetic circuit $C_{\mathcal{R}}$ over the field $\mathbb{F}$ such that $C_{\mathcal{R}}(\mathbf{x}, \mathbf{w}) = 0$ if and only if $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$. A message $\mathbf{x} \in \mathcal{L}_{\mathcal{R}}$ is a valid message. We stress that in our notation if $C_{\mathcal{R}}$ outputs 0 that means the input is *valid*, and any non-zero value means the input is *invalid*.

SYNTAX. The following definition is adapted from [12]. A fully linear probabilistically checkable proof system (FL-PCP) for $\mathcal{R}$ over $\mathbb{F}$ with proof length m and query complexity ℓ is a tuple of algorithms $\Pi_{\mathcal{R}} = (\text{ProofGen}, \text{Query}, \text{Decision})$ where the prover runs the ProofGen algorithm and the verifier runs the Query and Decision algorithms described below. Additionally, all algorithms also have an implicit input aux which contains relevant information about the corresponding relation $\mathcal{R}$.

- **Setup**($1^\lambda, \mathbb{F}$): On input a unary encoding of the security parameter 1^λ and a field $\mathbb{F}$ st $|\mathbb{F}| \in \Omega(2^\lambda)$, it updates the internal state.
- **ProofGen**($1^\lambda, \mathbf{x}, \mathbf{w}$): On input a unary encoding of the security parameter 1^λ , an instance $\mathbf{x}$, and its witness $\mathbf{w}$, the proof generation algorithm outputs a proof $\pi \in \mathbb{F}^m$.
- **Query**($1^\lambda, 1^d$): On input a unary encoding of the security parameter 1^λ and the length of an input instance 1^d , the query algorithm outputs ℓ queries $\mathbf{q}_1, \dots, \mathbf{q}_\ell \in \mathbb{F}^{d+m}$ and state information st .
- **Decision**($\text{st}, a_1, \dots, a_\ell$): On input state information st and ℓ scalar values $a_1, \dots, a_\ell \in \mathbb{F}$, the decision algorithm outputs a bit b .

Without loss of generality, we assume that the state of the query algorithm consists of the coins of the algorithm. Thus, the query algorithm can also be written as $(\mathbf{q}_1, \dots, \mathbf{q}_\ell) \leftarrow \text{Query}_{\mathcal{R}}(1^\lambda, 1^d; \text{st}_{\Pi})$.

The above algorithms satisfy the following properties:

COMPLETENESS. For all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$, the verifier accepts a valid proof. Formally, an FL-PCP $\Pi_{\mathcal{R}}$ satisfies completeness if for all $\lambda, d \in \mathbb{N}$

and all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$:

$$\Pr[\text{Decision}(\text{st}, \langle (\mathbf{x} || \boldsymbol{\pi}), \mathbf{q}_1 \rangle, \dots, \langle (\mathbf{x} || \boldsymbol{\pi}), \mathbf{q}_\ell \rangle) \Rightarrow 1] = 1,$$

where the probability is over

$$\boldsymbol{\pi} \leftarrow \text{ProofGen}(1^\lambda, \mathbf{x}, \mathbf{w}) ; (\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}) \leftarrow \text{Query}(1^\lambda, 1^d).$$

SOUNDNESS. For all $(\mathbf{x}, \mathbf{w}) \notin \mathcal{R}$, the verifier does not accept any proof. Formally, a FL-PCP $\Pi_{\mathcal{R}}$ satisfies γ -soundness for a function $\gamma(\cdot)$ if for all $\lambda, d \in \mathbb{N}$ and all $\mathbf{x}^* \notin \mathcal{L}_{\mathcal{R}}$ and a $\boldsymbol{\pi}^* \in \mathbb{F}^m$:

$$\Pr[\text{Decision}(\text{st}, \langle (\mathbf{x}^* || \boldsymbol{\pi}^*), \mathbf{q}_1 \rangle, \dots, \langle (\mathbf{x}^* || \boldsymbol{\pi}^*), \mathbf{q}_\ell \rangle) \Rightarrow 1] \leq \gamma(\lambda)$$

where the probability is over $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}) \leftarrow \text{Query}(1^\lambda, 1^d)$.

STRONG HVZK. Intuitively, the strong honest-verifier zero knowledge (HVZK) property states that the verifier does not learn anything about the instance $\mathbf{x}$ other than if $\mathbf{x} \in \mathcal{L}_{\mathcal{R}}$. A fully-linear PCP $\Pi_{\mathcal{R}}$ satisfies strong HVZK if there exists a simulator $S_{\Pi, \mathcal{R}}$ such that for all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ and $\lambda, d \in \mathbb{N}$, the distribution of $S_{\Pi, \mathcal{R}}(1^\lambda, 1^d)$ is identical to the following:

$$(\text{st}, \langle (\mathbf{x} || \boldsymbol{\pi}), \mathbf{q}_1 \rangle, \dots, \langle (\mathbf{x} || \boldsymbol{\pi}), \mathbf{q}_\ell \rangle), (\mathbf{q}_1, \dots, \mathbf{q}_\ell)$$

where $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}) \leftarrow \text{Query}(1^\lambda, 1^d)$ and $\boldsymbol{\pi} \leftarrow \text{ProofGen}(1^\lambda, \mathbf{x}, \mathbf{w})$.

3.3 Inner-Product Functional Encryption

SYNTAX. A public-key inner product functional encryption IPFE scheme is a tuple of algorithms $\text{IPFE} = (\text{Setup}, \text{KeyGen}, \text{Encrypt}, \text{Eval}, \text{Decrypt})$ described as follows:

- **Setup**($1^\lambda, 1^d$): On input unary encodings of security parameter λ and dimension d , the setup algorithm outputs a master public key mpk (in particular containing d and description of $\mathbb{F}$), a master secret key msk .
- **KeyGen**($\text{msk}, \mathbf{y}$): On input a master secret key msk and a function $\mathbf{y} \in \mathbb{F}^d$, the key generation algorithm outputs a functional decryption key $\text{sk}_{\mathbf{y}}$.
- **Encrypt**($\text{mpk}, \mathbf{x}$): On input a master public key mpk and a message $\mathbf{x} \in \mathbb{F}^d$, the encryption algorithm outputs a ciphertext $\text{ct}_{\mathbf{x}}$.
- **Eval**($\text{sk}_{\mathbf{y}}, \text{ct}_{\mathbf{x}}$): On input a functional decryption key $\text{sk}_{\mathbf{y}}$, and a ciphertext $\text{ct}_{\mathbf{x}}$, the evaluation algorithm outputs a string $z \in \mathbb{F}$ or a special symbol $\perp$.
- **Decrypt**($\text{msk}, \text{ct}_{\mathbf{x}}$): On input a master secret key msk and a ciphertext $\text{ct}_{\mathbf{x}}$, the decryption algorithm outputs a message $\mathbf{x} \in \mathbb{F}^d$ or a special symbol $\perp$.

CORRECTNESS We require IPFE to be correct meaning it has:

- (1) Evaluation correctness: For all $\lambda, d \in \mathbb{N}$, $\mathbf{x}, \mathbf{y} \in \mathbb{F}^d$:

$$\Pr[\text{IPFE.Eval}(\text{IPFE.KeyGen}(\text{msk}, \mathbf{y}), \text{IPFE.Encrypt}(\text{mpk}, \mathbf{x})) = \langle \mathbf{x}, \mathbf{y} \rangle] = 1$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{IPFE.Setup}(1^\lambda, 1^d)$,

- (2) Decryption correctness:

$$\Pr[\text{IPFE.Decrypt}(\text{msk}, \text{IPFE.Encrypt}(\text{mpk}, \mathbf{x})) = \mathbf{x}] = 1$$

for all $\lambda, d \in \mathbb{N}$, $(\text{pk}, \text{msk}, \text{pp}) \leftarrow \text{IPFE.Setup}(1^\lambda, 1^d)$, $\mathbf{x} \in \mathbb{F}^d$. The probability is taken over the coins of the encryption algorithm.

SECURITY The ipfe-ind game is defined in Figure 2.




Figure 2: Indistinguishability-based security game for IPFE.

We consider a multi-challenge setting. For an adversary $\mathcal{A}$ interacting with the experiment in Figure 2, the advantage is defined as follows:

$$\text{Adv}_{\text{IPFE}, \mathcal{A}}^{\text{ind-multi}}(\lambda) = 2 \cdot \Pr[G_{\text{IPFE}, \mathcal{A}}^{\text{ind}}(\lambda) \Rightarrow 1] - 1.$$

ADDITIVE HOMOMORPHISM. We use the notion of additive homomorphism given in [14]. We say that IPFE is *additively homomorphic* if it has an additional algorithm IPFE.Add as follows such that $\text{Add}(\text{mpk}, \text{ct}_{\mathbf{x}_1}, \text{ct}_{\mathbf{x}_2})$: On inputs a master public key mpk and two ciphertexts $\text{ct}_{\mathbf{x}_1}, \text{ct}_{\mathbf{x}_2}$, the Add algorithm outputs a ciphertext ct . Moreover, the following ciphertexts are identically distributed for all $\lambda, d \in \mathbb{N}$, and all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{F}^d$: $\text{Add}(\text{mpk}, \text{Enc}(\text{pk}, \mathbf{x}_1), \text{Enc}(\text{pk}, \mathbf{x}_2))$ and $\text{Enc}(\text{pk}, \mathbf{x}_1 + \mathbf{x}_2)$, where $(\text{pk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^d)$. Note that this means that the ciphertexts output by Add are compact in size. For notational convenience, by $\text{Add}(\text{mpk}, \text{ct}_1, \dots, \text{ct}_n)$ we mean that the Add algorithm is run repeatedly to generate a single ciphertext. It is easy to check that existing IPFE schemes are additively homomorphic.

UNRESTRICTED. We say that IPFE is *unrestricted* if its evaluation algorithm runs in polynomial time (PT). The IPFE schemes proposed in [1, 4] are not PT because they recover the inner product in the exponent. As a result, they restrict the message space so that the number of possible inner product values remains sufficiently small. The lattice-based IPFE scheme proposed in [4] (though over a small field) and the hard subgroup membership (HSM)-based scheme in [17] (which we use in our implementation) are unrestricted.

Henceforth, we assume that IPFE is always unrestricted.

PREIMAGE SAMPLEABILITY. The inner product function is *preimage sampleable* as defined in [35]. This means that given a set of n vectors $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ and their inner product values $a_1, \dots, a_n$ with some unknown vector $\mathbf{y}$, we can compute the vector $\mathbf{y}$ by solving the system of linear equations given that at least one such vector

y exists. Free variables in the system can be assigned arbitrarily. This is used in our security games as an algorithm VecSamp that inputs $((x_1, a_1), \dots, (x_n, a_n))$ and outputs a consistent vector v . Note that this is a property of the inner product function and, therefore, applies to any inner-product functional encryption scheme.

4 Non-Interactive Validated Aggregation

4.1 Syntax and Correctness

SYNTAX. A *non-interactive validated aggregation (NIVA)* protocol for a relation $\mathcal{R} \subseteq \mathbb{F}^d \times \mathbb{F}^h$ for a field $\mathbb{F}$ and $d, h, n \in \mathbb{N}$, is a protocol consisting of three kinds of participants: an analyst, a server and n clients and a tuple of algorithms $\text{NIVA}_{\mathcal{R}} = (\text{Setup}, \text{Encrypt}, \text{Validate}, \text{Aggregate}, \text{Decrypt})$ defined as follows:

- $\text{Setup}(1^\lambda, 1^d)$: On input unary encodings of the security parameter λ and an input length d , the setup algorithm outputs a public key mpk , a master secret key msk , and a validation key vk .
- $\text{Encrypt}(\text{mpk}, \mathbf{x}, \mathbf{w})$: On input a public key mpk , an input $\mathbf{x}$, and a witness $\mathbf{w}$, the encryption algorithm outputs a ciphertext ct .
- $\text{Validate}(\text{vk}, \text{ct})$: On input a validation key vk and a ciphertext ct , the validation algorithm outputs a bit b .
- $\text{Aggregate}(\text{mpk}, \{\text{ct}_j\}_{j \in [n]})$: On input a master public key mpk and a set of ciphertexts $\{\text{ct}_j\}_{j \in [n]}$ for some $n \in \mathbb{N}$, the aggregation algorithm outputs an aggregated ciphertext ct_{agg} .
- $\text{Decrypt}(\text{msk}, \text{ct})$: On input a master secret key msk and a ciphertext ct , the decryption algorithm outputs $\mathbf{y} \in \mathbb{F}^d$.

The protocol, which is also depicted in Fig. 1, is described next.

- The analyst runs the Setup algorithm and sends (mpk, vk) to the server and keeps msk to itself. In addition, the server sends mpk to each client when they join the system.
- Each client runs the Encrypt algorithm using the mpk and its input-witness pair $(\mathbf{x}, \mathbf{w})$ and sends the resulting ct to the server.
- The server upon receiving some number of ciphertexts from clients, runs the Validate algorithm on each ciphertext and discards invalid ones.
- The server uses the Aggregate algorithm to aggregates the n valid ciphertexts where n is less than or equal to the number of clients. Then, the server sends the aggregate ciphertext ct_{agg} to the analyst when requested.
- The analyst decrypts the aggregate ciphertext using the Decrypt algorithm and its msk .

CORRECTNESS. We require *validation correctness* meaning that for all $\lambda, d, h \in \mathbb{N}$ and all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$

$$\Pr[\text{Validate}(\text{vk}, \text{Encrypt}(\text{mpk}, \mathbf{x}, \mathbf{w})) \Rightarrow 1] = 1$$

where $(\text{mpk}, \text{msk}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, 1^d)$ and the probability is over the coins of the Setup and Encrypt algorithms. We say that $\text{NIVA}_{\mathcal{R}}$ has *aggregation error* γ if for all $\lambda, n \in \mathbb{N}$, all $(\mathbf{x}_1, \mathbf{w}_1), \dots, (\mathbf{x}_n, \mathbf{w}_n) \in \mathcal{R}$:

$$\Pr \left[\text{Decrypt} \left(\text{msk}, \text{Aggregate}(\text{mpk}, \{\text{ct}_j\}_{j \in [n]}) \right) = \sum_{j \in [n]} \mathbf{x}_j \right] \geq 1 - \gamma(\lambda)$$

where $(\text{mpk}, \text{msk}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, 1^d)$, $\text{ct}_j \leftarrow \text{Encrypt}(\text{mpk}, \mathbf{x}_j, \mathbf{w}_j)$ for all $j \in [n]$, and the probability is over the coins of the Setup and Encrypt algorithms.

4.2 Security

PRIVACY AGAINST A MALICIOUS SERVER. We define an indistinguishability-based notion of privacy for $\text{NIVA}_{\mathcal{R}}$ via the game in Figure 3. Here, an adversary receives a master public key and a validation key and has access to an encryption oracle. The adversary is modeled after a malicious server. She must make encryption queries with *valid* message and witness pairs. Then, she receives ciphertexts for the left or the right tuple and wins the game if she can determine whether the left or the right tuples are being encrypted. Note that we do not consider privacy of invalid inputs as that is equivalent to a malicious client. Formally, for an adversary $\mathcal{A}$, for every $\lambda, d \in \mathbb{N}$ we let its ind-advantage be:

$$\text{Adv}_{\text{NIVA}, \mathcal{A}}^{\text{ind}}(\lambda, d) = 2 \cdot \Pr[\text{G}_{\text{NIVA}, \mathcal{A}}^{\text{ind}}(\lambda, d) \Rightarrow 1] - 1.$$

We say that $\mathcal{A}$ is *valid* if every encryption query $((\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1))$ comprises of valid message-witness pair i.e., $(\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1) \in \mathcal{R}$. Intuitively, a malicious server should not be able to get any information about honest client data except for their validity status even when deviating from the protocol.




Figure 3: Game defining NIVA privacy against the server.

PRIVACY AGAINST THE ANALYST. We also define a notion of privacy against the analyst for $\text{NIVA}_{\mathcal{R}}$ via the game in Figure 4. An adversary submits two sets of inputs such that the valid inputs in each set have the same sum and receives an aggregated ciphertext and a master secret key. The adversary is modeled after a semi-honest analyst. The adversary wins the game if he can determine whether the left or the right set of inputs are validated and aggregated. Formally, for an adversary $\mathcal{A}$, for every $\lambda, d \in \mathbb{N}$, we let its ind-advantage be:

$$\text{Adv}_{\text{NIVA}, \mathcal{A}}^{\text{analyst-ind}}(\lambda, d) = 2 \cdot \Pr[\text{G}_{\text{NIVA}, \mathcal{A}}^{\text{analyst-ind}}(\lambda, d) \Rightarrow 1] - 1.$$

We say that $\mathcal{A}$ is *valid* if her submission $((\mathbf{x}_i^0, \mathbf{w}_i^0)_{i \in [n]}, (\mathbf{x}_i^1, \mathbf{w}_i^1)_{i \in [n]})$ is such that the following holds:

$$\sum_{i \in n: \mathbf{x}_i^0 \in \mathcal{L}_{\mathcal{R}}} \mathbf{x}_i^0 = \sum_{i \in n: \mathbf{x}_i^1 \in \mathcal{L}_{\mathcal{R}}} \mathbf{x}_i^1.$$

We say that $\text{NIVA}_{\mathcal{R}}$ is *analyst-private* if for all $d \in \mathbb{N}$, $\text{Adv}_{\text{NIVA}, \mathcal{A}}^{\text{ind}}(\cdot, d)$ is negligible for all valid PT $\mathcal{A}$.

We want to capture the idea that an analyst does not get any information about the individual client data including the number of valid inputs or validity status of individual data points and only receives the aggregated result of valid inputs.

Game $G_{NIVA, \mathcal{A}}^{\text{analyst-ind}}(\lambda, d)$

INITIALIZE($(\mathbf{x}_i^0, \mathbf{w}_i^0)_{i \in [n]}, (\mathbf{x}_i^1, \mathbf{w}_i^1)_{i \in [n]}$):

- 1 $b \leftarrow \{0, 1\}; S \leftarrow \emptyset$
- 2 $(\text{mpk}, \text{msk}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, 1^d)$
- 3 For $i \in [n]$:
- 4 $\text{ct}_i \leftarrow \text{Encrypt}(\text{mpk}, \mathbf{x}_i^b, \mathbf{w}_i^b)$
- 5 If $\text{Validate}(\text{vk}, \text{ct}_i) : S \leftarrow S \cup \{\text{ct}_i\}$
- 6 $\text{ct}_b \leftarrow \text{Aggregate}(\text{mpk}, S)$
- 7 Return $(\text{msk}, \text{ct}_b)$

FINALIZE(b'):

- 8 Return $(b = b')$

Figure 4: Game defining NIVA privacy against the analyst.

INTEGRITY. We define a notion of integrity for $NIVA_{\mathcal{R}}$ via the game in Figure 5. An adversary on receiving a master public key and has access to a validate oracle. The adversary is modeled after a malicious client. She submits ciphertexts to the validate oracle and receives a bit indicating if the ciphertext verified or not. The adversary submits a final ciphertext and wins the game if this ciphertext encrypts an invalid instance but passes the validity check. Concretely, for an adversary $\mathcal{A}$, for every $\lambda, d \in \mathbb{N}$ we let its int-advantage be:

$$\text{Adv}_{NIVA, \mathcal{A}}^{\text{int}}(\lambda, d) = \Pr[G_{NIVA, \mathcal{A}}^{\text{int}}(\lambda, d) \Rightarrow 1].$$

A PT adversary is *valid* if she submits a ct such that size of ct is $|\text{Encrypt}(\text{mpk}, \mathbf{v})|$ where $|\mathbf{v}| = d + m$.

We say that a $NIVA_{\mathcal{R}}$ scheme has integrity if the advantage of any valid PT adversary is negligible in the security parameter λ .

Intuitively, we want to capture the fact that it is difficult for malicious clients to get their invalid ciphertexts accepted and aggregated with valid ciphertexts even if they know the validity status of previous ciphertexts.

Game $G_{NIVA, \mathcal{A}}^{\text{int}}(\lambda, d)$

INITIALIZE:

- 1 $(\text{mpk}, \text{msk}, \text{vk}) \leftarrow \text{Setup}(1^\lambda, 1^d)$
- 2 Return mpk

VALIDATEO(ct):

- 3 Return $\text{Validate}(\text{vk}, \text{ct})$

FINALIZE(ct*):

- 4 If $\text{Validate}(\text{vk}, \text{ct}^*) :$
- 5 $\mathbf{y}^* \leftarrow \text{Decrypt}(\text{msk}, \text{ct}^*)$
- 6 If $C_{\mathcal{R}}(\mathbf{y}^*) \neq 0$: Return 1
- 7 Return 0

Figure 5: Game defining NIVA integrity.

4.3 Reusability in FL-PCPs

REUSABLE SOUNDNESS. First, we define an extension of soundness for FL-PCPs that we call *reusable* soundness, which involves reusing the verifier's queries across multiple proofs. Intuitively, this property states that reusing the state and queries of the FL-PCP across multiple proofs does not significantly degrade the soundness of the protocol, which will be important in our application.

For $\lambda, d \in \mathbb{N}$ and an adversary $\mathcal{A}$ interacting with the game given in Figure 6, we let its reusable soundness advantage be as follows:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{reuse-sound}}(\lambda, d) = \Pr[G_{NIVA, \mathcal{A}}^{\text{reuse-sound}}(\lambda, d) \Rightarrow 1].$$

We say that $\Pi_{\mathcal{R}}$ satisfies γ -reusable soundness if for all adversaries making Q REPEATO queries, and all $\lambda, d \in \mathbb{N}$ its reusable-soundness advantage is bounded as follows:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{reuse-sound}}(\lambda, d) \leq \gamma(\lambda, Q).$$

The FL-PCP of Boneh *et al.* [12], BBG⁺, and our optimized FL-PCP in Section 5 fulfill this property, which we prove in the appendix A.1.

Game $G_{\Pi, \mathcal{A}}^{\text{reuse-sound}}(\lambda, d)$

INITIALIZE:

- 1 $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}) \leftarrow \text{Query}(1^\lambda, 1^d)$

REPEATO($\mathbf{x} \parallel \pi$):

- 2 For $i \in [\ell] : a_i \leftarrow \langle \mathbf{x} \parallel \pi, \mathbf{q}_i \rangle$
- 3 $b \leftarrow \text{Decision}(\text{st}, a_1, \dots, a_\ell)$
- 4 Return b

FINALIZE($\mathbf{x}^* \parallel \pi^*$):

- 5 If $\mathbf{x}^* \in \mathcal{R}$: Return $\perp$
- 6 For $i \in [\ell] : a_i^* \leftarrow \langle \mathbf{x}^* \parallel \pi^*, \mathbf{q}_i \rangle$
- 7 $b^* \leftarrow \text{Decision}(\text{st}, a_1^*, \dots, a_\ell^*)$
- 8 Return b^*

Figure 6: Game defining reusable soundness property for FL-PCP.

REUSABLE ZERO-KNOWLEDGE. We also define a corresponding extension of strong HVZK for FL-PCP that we call *reusable* strong HVZK using the games in Figure 7. Intuitively, the property states that reusing the state and queries of the FL-PCP across multiple instances of proofs does not reveal any information about valid input instances.

For $\lambda, d \in \mathbb{N}$ and an adversary $\mathcal{A}$ interacting with the game given in Figure 7, we let its reusable strong HVZK advantage be as follows:

$$\text{Adv}_{\Pi, \mathcal{A}}^{\text{rshvzk}}(\lambda, d) = \Pr[G_{\Pi, \mathcal{A}}^{\text{rshvzk-1}}(\lambda, d) \Rightarrow 1] - \Pr[G_{\Pi, \mathcal{A}, S}^{\text{rshvzk-0}}(\lambda, d) \Rightarrow 1].$$

We say that $\Pi_{\mathcal{R}}$ satisfies reusable strong HVZK if for all adversaries and all $\lambda, d \in \mathbb{N}$ its reusable strong HVZK advantage is negligible in λ .

The FL-PCPs BBG⁺ and our optimized FL-PCP in Section 5 fulfill this property, which we prove in the appendix A.2.

4.4 PEAR: Our Black-Box NIVA Protocol

As discussed in the Introduction, a NIVA protocol that is non-blackbox in the underlying cryptography can be constructed from

Game $G_{\Pi, \mathcal{A}}^{\text{rshvzk-1}}(\lambda, d)$	Game $G_{\Pi, \mathcal{A}, S}^{\text{rshvzk-0}}(\lambda, d)$
INITIALIZE: 1 $(q_1, \dots, q_\ell, st) \leftarrow \text{Query}(1^\lambda, 1^d)$ 2 Return st REPEAT O(x, w): 3 $\pi \leftarrow \text{ProofGen}(1^\lambda, x, w)$ 4 For $i \in [\ell] : a_i \leftarrow \langle x \ \pi, q_i \rangle$ 5 Return $(a_1, \dots, a_\ell)$ FINALIZE(b): 6 Return b	INITIALIZE: 1 $st \leftarrow S_1(1^\lambda, 1^d)$ 2 Return st REPEAT O(x, w): 3 $(a_1, \dots, a_\ell) \leftarrow S_2(1^\lambda, 1^d, st)$ 4 Return $(a_1, \dots, a_\ell)$ FINALIZE(b): 5 Return b

Figure 7: Games defining reusable strong honest-verifier zero knowledge property for FL-PCP.

homomorphic encryption and NIZK. We now present our black-box solution.

CONSTRUCTION. Let $\Delta = (\text{Setup}, \text{KeyGen}, \text{Encrypt}, \text{Eval}, \text{Decrypt})$ be an additively-homomorphic IPFE scheme. Let $\mathcal{R}_p$ be a family of relations over finite fields $\mathbb{F}_p$ parameterized by the order p , a prime. Let $\Pi_{\mathcal{R}} = (\text{ProofGen}, \text{Query}, \text{Decision})$ be an FL-PCP system for a ² relation $\mathcal{R}_p \subseteq \mathbb{F}_p^d \times \mathbb{F}_p^h$. We define the associated NIVA protocol $\text{NIVA}_{\mathcal{R}}[\Delta, \Pi_{\mathcal{R}}] = (\text{Setup}, \text{Encrypt}, \text{Validate}, \text{Aggregate}, \text{Decrypt})$. Henceforth, we denote it by PEAR and describe it below with formal details given in Figure 8.

- During setup, the analyst runs the IPFE setup algorithm and generates query vectors for the FL-PCP. Next, the analyst generates function keys for the queries and sends the IPFE public key along with the verification key—consisting of the function keys and a state—to the server. The master secret key for IPFE is retained by the analyst.
- During encryption, a client generates a proof for its message and witness with respect to the validation circuit of the relation. The client then encrypts the message, concatenated with its proof, using the IPFE public key that is broadcasted by the server, and sends the resulting ciphertext to the server.
- During validation, the server receives a ciphertext and evaluates the inner product of the ciphertext with each query vector. It then runs the decision algorithm of the FL-PCP using the state of the query algorithm and the computed inner product values.
- During aggregation, the server combines n valid ciphertexts using IPFE’s homomorphic addition algorithm. Furthermore, it homomorphically adds a mask to the second component of the ciphertexts to obscure the sum of the proofs. The server then sends the aggregate ciphertext to the analyst.
- Upon receiving an aggregate ciphertext, the analyst decrypts it using the IPFE master secret key.

CORRECTNESS. PEAR is validation correct which follows from perfect evaluation correctness of the IPFE Δ and completeness of the FL-PCP $\Pi_{\mathcal{R}}$.

²We only support relations over finite fields of prime order because the IPFE scheme we use is over fields of prime order.

³mpk has the description of $\mathbb{F}$

⁴following the syntax of Δ as in section 3.3

Setup($1^\lambda, 1^d$): 1 $(\text{mpk}, \text{msk}) \leftarrow \Delta.\text{Setup}(1^\lambda, 1^d)$ 2 $\Pi_{\mathcal{R}}.\text{Setup}(1^\lambda, \text{mpk})$ ² 3 $(q_1, \dots, q_\ell, \text{st}_{\Pi}) \leftarrow \Pi_{\mathcal{R}}.\text{Query}(1^\lambda, 1^d)$ 4 For $j \in [\ell] :$ 5 $\text{sk}_{q_j} \leftarrow \Delta.\text{KeyGen}(\text{msk}, q_j)$ 6 $\text{vk} \leftarrow (\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_{\Pi})$ 7 Return $(\text{mpk}, \text{msk}, \text{vk})$ Encrypt(mpk, x, w): 8 $\pi \leftarrow \Pi_{\mathcal{R}}.\text{ProofGen}(1^\lambda, x, w)$ 9 $\text{ct} \leftarrow \Delta.\text{Encrypt}(\text{mpk}, (x \ \pi))$ ³ 10 Return ct Validate(ct, vk): 11 $(\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_{\Pi}) \leftarrow \text{vk}$ 12 For $j \in [\ell] :$ 13 $a_j \leftarrow \Delta.\text{Eval}(\text{sk}_{q_j}, \text{ct})$ 14 $b \leftarrow \Pi_{\mathcal{R}}.\text{Decision}(\text{st}_{\Pi}, a_1, \dots, a_\ell)$ 15 Return b Aggregate($\text{mpk}, \text{ct}_1, \dots, \text{ct}_n$): 16 $r \leftarrow \mathbb{F}^m$ 17 $\text{ct}_{n+1} \leftarrow \Delta.\text{Encrypt}(\text{mpk}, (0^d \ r))$ 18 If $n = 0$: Return ct_{n+1} 19 $\text{ct}_{\text{agg}} \leftarrow \Delta.\text{Add}(\text{mpk}, \text{ct}_1, \dots, \text{ct}_{n+1})$ 20 Return ct_{agg} Decrypt(msk, ct): 21 $y \leftarrow \Delta.\text{Decrypt}(\text{msk}, \text{ct})$ 22 Return $y[1 : d]$

Figure 8: PEAR: Our NIVA construction.

PEAR also has aggregation error $\gamma(\lambda) = 0$ which follows from additive homomorphism and perfect evaluation correctness of the IPFE Δ .

PRIVACY AGAINST A MALICIOUS SERVER. We show that PEAR is private against a malicious server under the indistinguishability-based security notion defined in Section 4.2.

THEOREM 1. *Let PEAR be as given in Fig. 8. Then PEAR is an ind-secure NIVA protocol if the IPFE scheme Δ is ind-secure and the FL-PCP protocol $\Pi_{\mathcal{R}}$ satisfies reusable strong HVZK. Namely, for any adversary $\mathcal{A}$, there exists an adversary $\mathcal{B}$ such that:*

$$\text{Adv}_{\text{PEAR}, \mathcal{A}}^{\text{ind}}(\lambda, d) \leq \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d).$$

The running-time of $\mathcal{B}$ is about the same as $\mathcal{A}$.

We prove the above theorem in a game-playing format. The first game is the privacy against the server game of NIVA. In the second game, we swap the vector that is encrypted with IPFE with another vector that has the same inner product values using a vector sampling algorithm. This is possible because inner-product is a pre-image sampleable function. An adversary that can distinguish between these two games can break indistinguishability-based security of the IPFE scheme. After that, we swap the query vector generated during the setup procedure with queries generated by a reusable strong honest-verifier zero-knowledge simulator for the

FL-PCP. The inner product values generated during each encryption are also replaced with outputs from the simulator. Here, these two games are indistinguishable because of the reusable strong honest-verifier zero-knowledge property of the FLPCP. A formal proof of the theorem including the hybrid games and the adversary is provided in Appendix B.

PRIVACY AGAINST THE ANALYST. Privacy of PEAR against the analyst follows directly from the property of the underlying IPFE Δ that homomorphically aggregated ciphertexts are distributed identically to freshly generated ones.

INTEGRITY. The integrity of PEAR follows directly from reusable soundness of the underlying FL-PCP as defined in 3.2. We summarize it in the following theorem.

=

THEOREM 2. *Let PEAR be as constructed in Fig. 8. Then, PEAR is an int-secure NIVA protocol. Formally, for any adversary $\mathcal{A}$ for all $\lambda, d \in \mathbb{N}$:*

$$\text{Adv}_{\text{PEAR}, \mathcal{A}}^{\text{int}}(\lambda, d) \leq \gamma(\lambda)$$

where $\gamma(\cdot)$ is the reusable soundness error of $\Pi_{\mathcal{R}}$.

5 Optimized FL-PCP for Conjunction-of-Predicates

We give an optimization to the FL-PCP construction of Boneh et al. [12] for relations that require multiple predicates to be simultaneously satisfied and where the predicates can be written as arithmetic subcircuits. A similar optimization has been made by Boneh et al. [12] for their Fully Linear Interactive Oracle Proofs (FL-IOP) where the verifier checks multiple identical subcircuits at the expense of a higher number of interaction rounds.

5.1 Conjunction-of-Predicates Relations

A common structure of an input validity check in practice involves verifying that the input satisfies multiple predicates at once. Namely, let $\mathcal{R} \subseteq \mathbb{F}^d \times \mathbb{F}^h$ for $d, h \in \mathbb{N}$. Let $C_1, \dots, C_k$ be arithmetic circuits and define

$$\mathcal{R}_{\text{AND}}[C_1, \dots, C_k](\mathbf{x}, \mathbf{w}) = (C_1(\mathbf{x}, \mathbf{w}) = 0) \wedge \dots \wedge (C_k(\mathbf{x}, \mathbf{w}) = 0)$$

which we call a *conjunction-of-predicates* relation³. For simplicity, we assume that each C_i has input length $d + h$. We give two main examples of such relations before discussing our optimization of the FL-PCP of [12] for them.

BINARY VALIDATION. We first consider the binary validation relation which checks that an input is a binary vector. Formally, $\mathcal{R}_{\text{BV}}(\mathbf{x}) = 1$ if for all $i \in [d]$, $x_i \in \{0, 1\}$ where $h = 0$. This relation does not use the witness. Then there exist arithmetic circuits $C_1, \dots, C_d$ such that $\mathcal{R}_{\text{BV}} = \mathcal{R}_{\text{AND}}[C_1, \dots, C_d]$, namely $C_i(\mathbf{x}) = x_i \cdot (x_i - 1)$ for all $i \in [d]$.

L_2 -NORM BOUND VALIDATION. We next consider the L_2 -norm bound validation relation parameterized by $b \in \mathbb{N}$ such that $2^{2b} < |\mathbb{F}|/d$, which checks that the square of an input vector's L_2 -norm *over the integers* is bounded by 2^{2b} . Note that it would not suffice to simply check the L_2 -norm *over the field* because of potential wrap-around.

³What follows later as $C_{\mathcal{R}}$ is the same as the C_i for some $i \in k$ here.

To ensure that there is no wrap-around we need to check that each component of the input $\mathbf{x}$ is bounded by 2^b . (Thus, the trivial bound on the square of the L_2 -norm is $d \cdot 2^{2b}$, but we check that it is actually much less; furthermore, our approach generalizes to any bound that is a power of two.) A simple way to check this, and the aforementioned bound on the square of the L_2 -norm, is to look at their binary decomposition. Specifically, if the decomposition of x_i has at most b bits, then $x_i < 2^b$. Thus, for our L_2 -norm bound validation we provide these binary decompositions as witnesses. Namely, for input $\mathbf{x} \in \mathbb{F}^d$, let witness $\mathbf{w} = (\mathbf{u}, \mathbf{v}) \in \mathbb{F}^{db} \times \mathbb{F}^{2b}$ the L_2 -norm bound validation relation $\mathcal{R}_{\text{NBV}}[b](\mathbf{x}, \mathbf{w}) = 1$ if:

- (1) $\mathbf{u}$ and $\mathbf{v}$ are binary vectors.
- (2) $\mathbf{u} = (\mathbf{u}_1 \parallel \dots \parallel \mathbf{u}_d)$ is a vector of binary decompositions of the components of $\mathbf{x}$ where each $\mathbf{u}_i$ is of length b .
- (3) $\mathbf{v}$ is the binary representation of $\|\mathbf{x}\|_2^2$

Observe that our witness $\mathbf{w}$ has length $h = db + 2b$. Again, we claim that $\mathcal{R}_{\text{NBV}}[b] = \mathcal{R}_{\text{AND}}[C_1, \dots, C_k]$ for $k = h + d + 1$ for arithmetic circuits $C_1, \dots, C_k$. The circuits are as follows:

- (1) For $j \in [h] : C_j(\mathbf{x}, \mathbf{w}) = w_j \cdot (w_j - 1)$.
- (2) For $i \in [d] : C_{h+i} = x_i - \left(\sum_{j=1}^b \mathbf{u}_{i-b+j} \cdot 2^{j-1} \right)$.
- (3) $C_{h+d+1} = \left(\sum_{i=1}^d x_i^2 \right) - \left(\sum_{j=1}^{2b} v_j \cdot 2^{j-1} \right)$.

5.2 Overview of Our Optimization

EXISTING FL-PCP. To explain our optimization, we first describe some details of BBG^+ , [12]'s FL-PCP and why it does not directly support conjunction-of-predicates relations optimally. Recall that their construction works on arithmetic circuits containing many instances of the same sub-circuit, where each sub-circuit takes L inputs and has one output. These sub-circuits are called “G-gates”. The construction guarantees that there exists an FL-PCP with strong HVZK for the arithmetic circuit $C_{\mathcal{R}}$.

The idea of the FL-PCP is that the prover computes $C(\mathbf{x}, \mathbf{w})$ and enumerates all the inputs and outputs of the G-gates as interpolated polynomials. The verifier checks that the polynomials are correctly generated by evaluating them at a uniform random point, which can be done via inner-product. The construction is provided in full in Fig. 14 for completeness and further details can be found in [12].

SUPPORTING CONJUNCTION-OF-PREDICATES. As BBG^+ works with arithmetic circuits, it does not directly support conjunctions-of-predicates relations in an efficient manner. Although, a boolean AND gate can be simulated by an arithmetic circuit in principle because both computational models are complete, this incurs a significant blowup in the circuit size. A slightly better idea is to modify BBG^+ by letting the output of each C_i be the output of a G-gate, and having the prover reveal the output of the *last* k G-gates, corresponding to the outputs of $C_i(\mathbf{x}, \mathbf{w})$, instead of just the last G-gate.

Note that since the output of $C_i(\mathbf{x}, \mathbf{w})$ on valid inputs is equal to 0 for all $i \in [k]$, revealing the output of these k gates instead of just one is still zero knowledge. On the other hand, if the input is invalid, zero knowledge is not required to hold. We omit details since this is not our final construction. Indeed, the problem is that revealing the output of the last k G-gates requires k additional linear queries

by the verifier, which in our NIVA construction corresponds to k decryptions under the IPFE, a very expensive operation.

OUR OPTIMIZATION. To describe our optimization, we first observe that the verifier computes an affine transformations on the proof, thus it is possible to take linear combinations of these transformations to compute a linear combination of the output of the transformations. Thus, we could compute $\sum_{i=1}^k C_i(\mathbf{x}, \mathbf{w})$ using a single linear query by simply adding up the queries that compute each $C_i(\mathbf{x}, \mathbf{w})$. However, as previously stated it does not suffice to check that $\sum_{i=1}^k C_i(\mathbf{x}, \mathbf{w})$ is equal to 0.

Our main observation is that checking that a *random* linear combination of $C_i(\mathbf{x}, \mathbf{w})$ equals to 0, where the randomness is *chosen by the verifier* allows to use just *one* query while ensuring soundness with overwhelming probability. In more detail, the verifier computes $\sum_{i=1}^k \psi_i \cdot C_i(\mathbf{x}, \mathbf{w})$ and checks that it is equal to 0, where $\psi_i \leftarrow \mathbb{F}$. It is unlikely for this sum to be equal to 0 when $C_i(\mathbf{x}, \mathbf{w})$ are not all equal to 0 individually. Finally, note that this sum can indeed be computed by a linear query equal to a weighted sum of the linear queries that compute $C_i(\mathbf{x}, \mathbf{w})$. We stress that this check only requires opening the output of this one value, instead of having to open all k G -gates as proposed earlier.

We note that for the naive method of opening the last k G -gates, the total query size and decision time increase quadratically with the input length. This quadratic growth happens because more G -gates must be opened as the input length increases⁴, and each query becomes longer. In contrast, our optimization results in a linear increase in both total query size and decision time. It requires only a single query to recover the random linear combination of all the G -gates, with only the query length growing as the input length increases.

5.3 Details of the Optimized FL-PCP

We now give our optimized FL-PCP in detail, which we denote by $\Pi' = (\text{ProofGen}, \text{Query}, \text{Decision})$, in Fig. 9. We show below that it fulfills completeness, soundness, and strong honest-verifier zero-knowledge properties. For simplicity, assume that the output of the last k G -gates are the outputs of $C_1, \dots, C_k$. Here, the verifier wants to check that $C_i(\mathbf{x}, \mathbf{w}) = 0$ for each $i \in [k]$. Since all the coefficients of the polynomial p are a part of the proof vector, the verifier can compute $p(i)$ for any i by taking the inner product of the coefficient vector with the following vector: $(i^0, i^1, \dots, i^{\deg})$ where $\deg$ is the degree of p . To compute $\sum_{i=1}^k \psi_i \cdot C_i(\mathbf{x}, \mathbf{w})$ the verifier computes $\sum_{i=1}^k \psi_i \cdot p(M - i + 1)$ using the query vector on line 19 of Fig. 9.

While we assume that the output of each C_i is the output of a G -gate, we can also accommodate circuits that consist only of affine gates. In particular, in our L_2 -norm bound validation, circuits that check if a binary decomposition is correct are affine circuits. Since the output of this type of circuit is not the output of a G -gate, it cannot be revealed using polynomial evaluation. However, its output is simply a linear query on the instance-proof vector.

We show that this modified version of FL-PCP satisfies Completeness, Soundness, and Strong honest-verifier zero-knowledge in the Appendix C.

⁴increase in input length is inclusive of the increase in depth

```

ProofGen( $1^\lambda, \mathbf{x}, \mathbf{w}$ ):
1 Evaluate  $C_1(\mathbf{x}, \mathbf{w}), \dots, C_k(\mathbf{x}, \mathbf{w})$ 
2 For  $i \in [L]$  :
3    $f[i][0] \leftarrow \mathbb{F}$ 
4   For  $j \in [M]$  :
5      $f[i][j] \leftarrow i^{th}$  input to the  $j^{th}$   $G$ -gate
6 Use polynomial interpolation for  $i \in [L]$  as follows:
7    $\{(j, f[i][j])\}_{j=0}^M$  to compute  $f_i$ 
8    $p \leftarrow G(f_1, \dots, f_L)$ 
9 Let  $\mathbf{c}_p$  be the coefficients of  $p$ 
10  $\pi \leftarrow (\mathbf{w} || f[1][0] \dots f[L][0] || \mathbf{c}_p)$ 
11 Return  $\pi$ 

Query( $1^\lambda, 1^d$ ):
12  $r \leftarrow \mathbb{F} \setminus [M]$ 
13 Compute linear queries for polynomial evaluation as follows:
14   compute  $\lambda$  to evaluate  $p(r)$ 
15   For  $i \in [L]$  : compute  $\lambda_i$  to evaluate  $f_i(r)$ 
16  $\psi \leftarrow \mathbb{F}^k$ 
17  $\mathbf{q}_1 \leftarrow \lambda$ 
18  $\mathbf{q}_2, \dots, \mathbf{q}_L \leftarrow \lambda_1, \dots, \lambda_L$ 
19  $\mathbf{q}_{L+1} \leftarrow (0^{d+h+L} || \sum_{i=1}^k \psi_i \cdot i^0, \dots, \sum_{i=1}^k \psi_i \cdot i^{\deg})$ 
20  $\mathbf{st} \leftarrow (r, \psi)$ 
21 Return  $(\mathbf{q}_1, \dots, \mathbf{q}_{L+1}, \mathbf{st})$ 

Decision( $\mathbf{st}, a_1, \dots, a_{L+1}$ ):
22 If  $a_1 = G(a_2, \dots, a_L)$  and  $a_{L+1} = 0$ : Return 1
23 Return 0

```

Figure 9: Optimized FL-PCP for Conjunction-of-Predicates.

6 Experimental Evaluation

In this section, we present an experimental evaluation of PEAR.⁵ For the underlying unrestricted IPFE scheme, we utilize the construction from [17], which we refer to as CLT. Details of this construction are provided in Appendix D for completeness. We implement circuits for the Binary Validation $\mathcal{R}_{BV}$ and L_2 -Norm Bound Validation $\mathcal{R}_{NBV}$ relations as previously defined in Section 5.1.

We establish a baseline for comparison using the straightforward NIVA protocol based on additively homomorphic encryption (HE) combined with a non-interactive zero-knowledge proof of knowledge (NIZK). Recall that in this approach, the client encrypts their plaintexts with the HE scheme and uses the NIZK to prove that the ciphertext is well-formed and satisfies the required constraints. The client then sends both the ciphertext and the proof to the server, who verifies the proof and aggregates the validated ciphertexts. In the following discussion, we refer to this approach as NIZK+HE.

ON FHE+FL-PCP. We also attempted to compare PEAR with a second “off-the-shelf” NIVA protocol based on FHE+FL-PCP, where FL-PCP validation is performed under the hood of FHE. Recall that the FL-PCP decision algorithm computes a G -gate and compares the output with another value, denoting their difference as t . The decision algorithm determines the proof is valid if $t = 0$. It is straightforward for IPFE to check the value of t as it obtains the value in the clear. However, checking the same for the FHE approach is significantly more challenging, since the server does not interact with the client. We take the approach of exponentiating the FHE ciphertext of t to the order of the plaintext space to obtain

⁵We will open-source our implementation upon publication.

a bit, where a zero bit indicates “valid”. We then flip the bit and use it to multiply the corresponding ciphertext to enable aggregation without decrypting the intermediate result.

We implement this approach using the BFV scheme from the SEAL library [40], written in C++. The FL-PCP scheme should operate over a 128-bit field; however, the BFV scheme in SEAL naively supports a plaintext modulus of at most 2^{60} . Although longer integers could be managed by initializing multiple schemes and concatenating their ciphertexts, we omit this step and instead use a smaller plaintext size of 2^{60} for BFV to provide it with some advantage and demonstrate that PEAR already offers a significant speedup while also uses a much larger field. For the $\mathcal{R}_{BV}$ relation, the server runtime of the FHE+FL-PCP approach is 20× to over 100× slower than that of PEAR in our experiments. Therefore, we exclude the FHE+FL-PCP approach from our detailed analysis below.

IMPLEMENTATION SETUP. We implement PEAR in Sage [39], which uses a C backend, leveraging the PARI/GP library [26] to provide robust support of ideal class groups (for the CLT scheme). It also supports efficient polynomial computations over a field (for the FL-PCP). For the NIZK+HE approach, we use Bulletproofs and ElGamal implementations from the Tongsuo library [36], also written in C. We assume a network latency of 30 ms and a bandwidth of 100 Mbps. All experiments were conducted on a cloud server equipped with an Intel Xeon Platinum 8160 CPU (96 cores, 2.10 GHz) and 64 GB of DDR4 memory, running Ubuntu 22.04.

6.1 Server-Side Optimizations

We introduce server-side optimizations to PEAR tailored to the relations we consider and to the IPFE scheme CLT. We observe that, for both the $\mathcal{R}_{BV}$ and $\mathcal{R}_{NBV}$ relations, the G -gate used is a single multiplication gate. Additionally, to recover the inner product r of a ciphertext ct_x and a function key sk_y , CLT.Eval first computes g^r , where g is a generator of a group. It then uses a trapdoor to take the discrete logarithm and recover the exponent, a process that remains significantly more expensive than basic group operations such as multiplication and exponentiation.

However, to validate the FL-PCP proof, recovering the exact values of $f_1(r)$ and $f_2(r)$ is unnecessary; the goal is simply to verify that their product *equals* $p(r)$. This can be done by instead checking that $g^{f_1(r) \cdot f_2(r)}$ equals $g^{p(r)}$. However, at least in the CLT scheme, there is no direct way to compute $f_1(r) \cdot f_2(r)$ in the exponent. Thus, we first recover the value of $f_1(r)$ as described above. Then during the decryption of $f_2(r)$, we skip the final step of using the trapdoor to take the discrete logarithm and compute $(g^{f_2(r)})^{f_1(r)}$ to obtain $g^{f_1(r) \cdot f_2(r)}$. Similarly, when decrypting $p(r)$, we skip the discrete logarithm step and just retrieve $g^{p(r)}$, then check equality of the resulting group elements. This optimization avoids two trapdoor discrete logarithm computations per input validation.

Another server-side optimization applies to the aggregation step. In the aggregation step of PEAR, we are only interested in the aggregate value of the first part of the input vectors. In the IPFE CLT, a ciphertext ct of a vector $(x||\pi)$ consists of components $(ct_0, ct_1, \dots, ct_{d+m})$ where $d + m$ is the length of $(x||\pi)$. Here, ct_0 is the randomness used during encryption and ct_i is a function of ct_0 and x_i . Thus, instead of adding the entire vectors, we can input




Figure 10: Client runtimes.




Figure 11: Server runtimes.

the components $(ct_0, ct_1, \dots, ct_d)$ to CLT.Add. This optimization reduces both server runtime and bandwidth.

6.2 PEAR Evaluation

We benchmark the performance of PEAR and NIZK+HE with the $\mathcal{R}_{BV}$ and $\mathcal{R}_{NBV}$ validations. For the $\mathcal{R}_{NBV}$ relation, we assume that each x_i is bounded by 2^{10} , and the L_2 -norm is bounded by 2^{20} .

CLIENT PERFORMANCE. In Figure 10, client runtimes are split into two components: proof generation time and encryption time. The proof generation time for PEAR is significantly shorter than for NIZK+HE, primarily because the FL-PCP proof generation operates solely over the field. While the most computationally expensive step is interpolating polynomials over the field, we optimize this process to $O(n \log n)$ using the Fast Fourier Transform [27]. In contrast, Bulletproof’s proof generation involves elliptic curve operations, which are inherently slower. On the other hand, encryption time for PEAR is slower than that of NIZK+HE. This is because in the case of the former the plaintext to be encrypted is longer, as it includes a proof and witness. However, since proof generation time dominates the client’s computation, PEAR still achieves a speedup of 2.8× to 11× in client computation time compared to the NIZK+HE method.




Figure 12: Analyst runtimes.

SERVER PERFORMANCE. In Figure 11, the server’s runtime is split into validation time and aggregation time. We note that validation time for NIZK+HE is larger since this requires verifying correctness of encryption as well as validity of the input. However, the relative cost of this decreases as the validity circuit grows. The reported aggregation time corresponds to adding two ciphertexts. In both approaches, aggregation time is independent of the circuit complexity, depending only on the input length. Overall, PEAR achieves a performance speedup of $1.2\times$ to $3.6\times$ in server computation time.

ANALYST PERFORMANCE. Figure 12 presents the analyst’s runtime which consists of running the setup procedure and decrypting the aggregated result. For PEAR, the Setup time is longer because it also involves executing CLT.KeyGen to generate queries for FL-PCP proof validation. Note that for both of these approaches, however, the runtime of the analyst is independent of the number of clients, which may be the dominating term in practice. Moreover, if more efficient unrestricted IPFE is developed in future work, it can be immediately plugged into our protocol to make the analyst’s runtime much smaller, likely outperforming that for NIZK+HE.

OVERALL PROTOCOL. Figure 13 illustrates the total runtime of the protocol. The client time reflects the duration for a single client to generate a proof for the $\mathcal{R}_{BV}$ relation and encrypt data with an input length of 100. The server time is measured with 20 participating clients, where selectivity denotes the percentage of valid client ciphertexts. As selectivity increases, meaning more validated ciphertexts are aggregated, the total server time increases only slightly. This indicates that the dominant server cost arises from validating each ciphertext rather than from aggregation. Regardless of selectivity, the overall communication time remains constant. Under our network setting of 30 ms latency and 100 Mbps bandwidth, the analyst in both PEAR and NIZK+HE can retrieve the aggregated results within 40 ms. When running the entire protocol, server time dominates the total runtime even with only 10 participating clients. Specifically, for PEAR and NIZK+HE, server time accounts for over 87% and 77% of the total runtime, respectively. As the number of clients increases, this percentage continues to grow, since the analyst’s time and each client’s time remain constant when the input length is fixed. Hence, the overall speedup PEAR




Figure 13: Total runtimes.

achieves is similar to the speedup it achieves in server time when there are many participating clients. We stress that our server runtime can be trivially reduced by parallel processing (as ciphertext validation and aggregation are independent across clients), which is not reflected in our experiments.

7 Conclusion

We have shown the first instantiation of a NIVA protocol based on IPFE and FL-PCP. Our result shows basic feasibility for this primitive and demonstrates that, with further optimizations of the underlying primitives, this can become practical for real-world use.

A number of open questions remain for further extensions of this primitive. We only consider the setting where the server is semi-honest. It is an open question if one can build a non-interactive, black-box protocol where the server is allowed to be malicious. A particularly challenging case is guaranteeing integrity when a malicious server may collude with an input client.

A second open question to consider is whether we can remove the analyst instead providing output directly to the server. This seems difficult to achieve in the non-interactive setting since the server could mix and match ciphertexts to learn more information about the messages, i.e., so called residual attacks. Handling *residual attacks* in a non-interactive setting where server and malicious clients collude remains an interesting open question.

Acknowledgments

Arkady Yerukhimovich is supported in part by NSF grants CNS-2144798 (CAREER) and CNS-1955620.

References

- [1] Michel Abdalla, Florian Bourse, Angelo De Caro, and David Pointcheval. 2015. Simple Functional Encryption Schemes for Inner Products. In *PKC 2015: 18th International Conference on Theory and Practice of Public Key Cryptography (Lecture Notes in Computer Science, Vol. 9020)*, Jonathan Katz (Ed.). Springer Berlin Heidelberg, Germany, Gaithersburg, MD, USA, 733–751. https://doi.org/10.1007/978-3-662-46447-2_33
- [2] Michel Abdalla, Romain Gay, Mariana Raykova, and Hoeteck Wee. 2017. Multi-input Inner-Product Functional Encryption from Pairings. In *Advances in Cryptology – EUROCRYPT 2017, Part I (Lecture Notes in Computer Science, Vol. 10210)*, Jean-Sébastien Coron and Jesper Buus Nielsen (Eds.). Springer, Cham, Switzerland, Paris, France, 601–626. https://doi.org/10.1007/978-3-319-56620-7_21
- [3] Surya Addanki, Kevin Garbe, Eli Jaffe, Rafail Ostrovsky, and Antigoni Polychroniadou. 2022. Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares. In *SCN 22: 13th International Conference on Security in Communication Networks (Lecture Notes in Computer Science, Vol. 13409)*, Clemente Galdi and Stanislaw Jarecki (Eds.). Springer, Cham, Switzerland, Amalfi, Italy, 516–539. https://doi.org/10.1007/978-3-031-14791-3_23
- [4] Shweta Agrawal, Benoît Libert, and Damien Stehlé. 2016. Fully Secure Functional Encryption for Inner Products, from Standard Assumptions. In *Advances in Cryptology – CRYPTO 2016, Part III (Lecture Notes in Computer Science, Vol. 9816)*, Matthew Robshaw and Jonathan Katz (Eds.). Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA, 333–362. https://doi.org/10.1007/978-3-662-53015-3_12
- [5] Scott Ames, Carmit Hazay, Yuval Ishai, and Muthuramakrishnan Venkatasubramanian. 2017. Liger: Lightweight sublinear arguments without a trusted setup. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*. 2087–2104.
- [6] Laasya Bangalore, Mohammad Hossein Faghihi Sereshgi, Carmit Hazay, and Muthuramakrishnan Venkatasubramanian. 2023. Flag: A Framework for Lightweight Robust Secure Aggregation. In *ASIACCS 23: 18th ACM Symposium on Information, Computer and Communications Security*, Joseph K. Liu, Yang Xiang, Surya Nepal, and Gene Tsudik (Eds.). ACM Press, Melbourne, VIC, Australia, 14–28. <https://doi.org/10.1145/3579856.3595805>
- [7] James Bell, Adrià Gascón, Tancrède Lepoint, Baiyu Li, Sarah Meiklejohn, Mariana Raykova, and Cathie Yun. 2023. ACORN: Input Validation for Secure Aggregation. In *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9–11, 2023*, Joseph A. Calandrino and Carmela Troncoso (Eds.). USENIX Association, 4805–4822. <https://www.usenix.org/conference/usenixsecurity23/presentation/bell>
- [8] James Bell-Clark, Adrià Gascón, Baiyu Li, Mariana Raykova, and Philipp Schoppmann. 2024. Willow: Secure Aggregation with One-Shot Clients. *Cryptology ePrint Archive*, Paper 2024/936. <https://eprint.iacr.org/2024/936>
- [9] Mihir Bellare and Phillip Rogaway. 2006. The Security of Triple Encryption and a Framework for Code-Based Game-Playing Proofs. In *Advances in Cryptology – EUROCRYPT 2006 (Lecture Notes in Computer Science, Vol. 4004)*, Serge Vaudenay (Ed.). Springer Berlin Heidelberg, Germany, St. Petersburg, Russia, 409–426. https://doi.org/10.1007/11761679_25
- [10] Allison Bishop, Abhishek Jain, and Lucas Kowalczyk. 2015. Function-Hiding Inner Product Encryption. In *Advances in Cryptology – ASIACRYPT 2015, Part I (Lecture Notes in Computer Science, Vol. 9452)*, Tetsu Iwata and Jung Hee Cheon (Eds.). Springer Berlin Heidelberg, Germany, Auckland, New Zealand, 470–491. https://doi.org/10.1007/978-3-662-48797-6_20
- [11] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *ACM CCS 2017: 24th Conference on Computer and Communications Security*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM Press, Dallas, TX, USA, 1175–1191. <https://doi.org/10.1145/3133956.3133982>
- [12] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2019. Zero-knowledge proofs on secret-shared data via fully linear PCPs. In *Annual International Cryptology Conference*. Springer, 67–97.
- [13] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2021. Lightweight Techniques for Private Heavy Hitters. In *2021 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, San Francisco, CA, USA, 762–776. <https://doi.org/10.1109/SP40001.2021.00048>
- [14] Florian Bourse. 2017. *Functional encryption for inner-product evaluations*. Theses. Université Paris sciences et lettres. <https://theses.hal.science/tel-01665276>
- [15] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. 2018. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 315–334.
- [16] Guilhem Castagnos and Fabien Laguillaumie. 2015. Linearly Homomorphic Encryption from DDH. In *Topics in Cryptology – CT-RSA 2015 (Lecture Notes in Computer Science, Vol. 9048)*, Kaisa Nyberg (Ed.). Springer, Cham, Switzerland, San Francisco, CA, USA, 487–505. https://doi.org/10.1007/978-3-319-16715-2_26
- [17] Guilhem Castagnos, Fabien Laguillaumie, and Ida Tucker. 2018. Practical Fully Secure Unrestricted Inner Product Functional Encryption Modulo p . In *Advances in Cryptology – ASIACRYPT 2018, Part II (Lecture Notes in Computer Science, Vol. 11273)*, Thomas Peyrin and Steven Galbraith (Eds.). Springer, Cham, Switzerland, Brisbane, Queensland, Australia, 733–764. https://doi.org/10.1007/978-3-030-03329-3_25
- [18] Amrita Roy Chowdhury, Chuan Guo, Somesh Jha, and Laurens van der Maaten. 2022. EIFFeL: Ensuring Integrity for Federated Learning. In *ACM CCS 2022: 29th Conference on Computer and Communications Security*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM Press, Los Angeles, CA, USA, 2535–2549. <https://doi.org/10.1145/3548606.3560611>
- [19] Henry Corrigan-Gibbs and Dan Boneh. 2017. Prio: Private, Robust, and Scalable Computation of Aggregate Statistics. In *14th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2017, Boston, MA, USA, March 27–29, 2017*, Aditya Akella and Jon Howell (Eds.). USENIX Association, 259–282. <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/corrigan-gibbs>
- [20] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *TCC 2006: 3rd Theory of Cryptography Conference (Lecture Notes in Computer Science, Vol. 3876)*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Germany, New York, NY, USA, 265–284. https://doi.org/10.1007/11681878_14
- [21] Taher ElGamal. 1984. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. In *Advances in Cryptology – CRYPTO’84 (Lecture Notes in Computer Science, Vol. 196)*, G. R. Blakley and David Chaum (Eds.). Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA, 10–18. https://doi.org/10.1007/3-540-39568-7_2
- [22] Uriel Feige, Dror Lapidot, and Adi Shamir. 1990. Multiple Non-Interactive Zero Knowledge Proofs Based on a Single Random String (Extended Abstract). In *31st Annual Symposium on Foundations of Computer Science*. IEEE Computer Society Press, St. Louis, MO, USA, 308–317. <https://doi.org/10.1109/FSCS.1990.89549>
- [23] Craig Gentry. 2009. Fully homomorphic encryption using ideal lattices. In *41st Annual ACM Symposium on Theory of Computing*, Michael Mitzenmacher (Ed.). ACM Press, Bethesda, MD, USA, 169–178. <https://doi.org/10.1145/1536414.1536440>
- [24] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1985. The Knowledge Complexity of Interactive Proof-Systems (Extended Abstract). In *17th Annual ACM Symposium on Theory of Computing*. ACM Press, Providence, RI, USA, 291–304. <https://doi.org/10.1145/22145.22178>
- [25] Jens Groth. 2016. On the size of pairing-based non-interactive arguments. In *Advances in Cryptology – EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II 35*. Springer, 305–326.
- [26] The PARI Group. 2024. PARI/GP Library. <http://pari.math.u-bordeaux.fr>. 2.17.1.
- [27] Michael Hecht and Ivo F. Sbalzarini. 2019. Fast Interpolation and Fourier Transform in High-Dimensional Spaces. In *Intelligent Computing*, Kohei Arai, Supriya Kapoor, and Rahul Bhatia (Eds.). Springer International Publishing, Cham, 53–75.
- [28] IEEE S&P 2023 2023. *2023 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, San Francisco, CA, USA.
- [29] Harish Karthikeyan and Antigoni Polychroniadou. 2024. OPA: One-shot Private Aggregation with Single Client Interaction and its Applications to Federated Learning. *Cryptology ePrint Archive*, Paper 2024/723. <https://eprint.iacr.org/2024/723>
- [30] Joe Kilian. 1992. A note on efficient zero-knowledge proofs and arguments. In *Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*. 723–732.
- [31] Sam Kim, Kevin Lewi, Avradip Mandal, Hart Montgomery, Arnab Roy, and David J Wu. 2018. Function-hiding inner product encryption is practical. In *Security and Cryptography for Networks: 11th International Conference, SCN 2018, Amalfi, Italy, September 5–7, 2018, Proceedings 11*. Springer, 544–562.
- [32] Hanjun Li, Huijia Lin, Antigoni Polychroniadou, and Stefano Tessaro. 2023. LERNA: Secure Single-Server Aggregation via Key-Homomorphic Masking. In *Advances in Cryptology – ASIACRYPT 2023, Part I (Lecture Notes in Computer Science, Vol. 14438)*, Jian Guo and Ron Steinfeld (Eds.). Springer, Singapore, Singapore, Guangzhou, China, 302–334. https://doi.org/10.1007/978-981-99-8721-4_10
- [33] Hidde Lycklama, Lukas Burkhalter, Alexander Viand, Nicolas Küchler, and Anwar Hithnawi. 2023. RoFL: Robustness of Secure Federated Learning, See [28], 453–476. <https://doi.org/10.1109/SP46215.2023.10179400>
- [34] Silvio Micali. 2000. Computationally sound proofs. *SIAM J. Comput.* 30, 4 (2000), 1253–1298.
- [35] Adam O’Neill. 2010. Definitional Issues in Functional Encryption. *Cryptology ePrint Archive*, Report 2010/556. <https://eprint.iacr.org/2010/556>
- [36] OpenAtom. 2024. Tongsuo Library. <https://www.tongsuo.net/en/docs/>. 8.4.0.
- [37] Bryan Parno, Jon Howell, Craig Gentry, and Mariana Raykova. 2013. Pinocchio: Nearly Practical Verifiable Computation. In *2013 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, Berkeley, CA, USA, 238–252. <https://doi.org/10.1109/SP.2013.47>
- [38] Mayank Rathee, Conghao Shen, Sameer Wagh, and Raluca Ada Popa. 2023. ELISA: Secure Aggregation for Federated Learning with Malicious Actors, See [28], 1961–1979. <https://doi.org/10.1109/SP46215.2023.10179468>

- [39] Sagemath. 2024. Sage Library. <https://www.sagemath.org/development-ack.html>. 10.5.0.
- [40] SEAL 2023. Microsoft SEAL (release 4.1). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA..
- [41] Jianan Su, Laasya Bangalore, Harel Berger, Jason Yi, Alivia Castor, Micah Sherr, and Muthuramakrishnan Venkitasubramaniam. 2024. SCIF: Privacy-Preserving Statistics Collection with Input Validation and Full Security. Cryptology ePrint Archive, Paper 2024/1821. <https://eprint.iacr.org/2024/1821>

A Reusability of BBG⁺

A.1 Soundness

LEMMA 1. *The FL-PCP BBG⁺, satisfies γ -reusable soundness as defined in Section 4.3 where*

$$\gamma \leq \frac{Q(d+1)}{|\mathbb{F}| - (M + Qd)}$$

and Q is the number of queries made to REPEATO, M is the number of G -gates in the validation circuit, and d is the arithmetic degree of the gate G .

PROOF. We include BBG⁺ here for completeness in Figure 14. We omit specifics of the QueryVecs algorithm for simplification. Here $\mathcal{R}$ is a relation such that $C_{\mathcal{R}}$ is the corresponding validation circuit that consists of G -gates with in-degree L . Additionally, let $\ell = L + 2$ be the query complexity.




Figure 14: Fully Linear PCP BBG⁺ [12].

We present the reusable soundness game modified to include a bad flag in Figure 15. The modifications are given in red. First, we show that the bad flag is set to 1 with negligible probability. Then, we show that if the bad flag is not set to 1 when we reach the FINALIZE procedure, any adversary $\mathcal{A}$ has a negligible advantage using arguments similar to the proof of soundness of BBG⁺.

First, we analyze the input $(\mathbf{x}^* \parallel \pi^*)$ to the FINALIZE procedure. A valid input to this procedure must contain an input $\mathbf{x}^*$ such that $\mathbf{x}^* \notin \mathcal{R}$. We parse π^* as $\mathbf{w}^* \parallel z_1^* \dots z_L^* \parallel \mathbf{c}_{p^*}$ where:

- (1) $\mathbf{w}^*$ is the witness for $\mathbf{x}^*$.
- (2) $z_1^*, \dots, z_L^*$ are some constants.
- (3) $\mathbf{c}_{p^*}$ is a vector of coefficients of a polynomial p^* .

For $(\mathbf{x}^* \parallel \pi^*)$, let $f_1^*, \dots, f_{\ell-2}^*$ and p^* be the polynomials purported to be associated with the input and output wires of the evaluation of $C_{\mathcal{R}}(\mathbf{x}^* \parallel \mathbf{w}^*)$. Additionally, let a_i^* be equal to $\langle \mathbf{x}^* \parallel \pi^*, \mathbf{q}_i \rangle$ for each $i \in [\ell]$.

If the FINALIZE procedure outputs 1 when $\mathbf{x}^* \notin \mathcal{R}$, this means that both $G(a_1^*, \dots, a_{\ell-2}^*) = a_{\ell-1}^*$ and $a_\ell^* = 0$. However, since there does not exist a witness $\mathbf{w}^*$ such that $C(\mathbf{x}^*, \mathbf{w}^*) = 0$, it must be




Figure 15: Game defining reusable soundness property for BBG⁺. Modifications to include a bad flag are given in red.

that p^* represented by the coefficient vector $\mathbf{c}_{p^*}$ is not the same as $G(f_1^*, \dots, f_{\ell-2}^*)$. Hence, $p^* - G(f_1^*, \dots, f_{\ell-2}^*)$ is not an identically zero polynomial and the state value st must be a root of this polynomial.

Next, consider the queries made to the REPEATO procedure. Note that in the REPEATO oracle, we do not impose a restriction on whether the submitted $\mathbf{x}_i$ must be in $\mathcal{R}$. The adversary can submit any combination of valid and invalid input $\mathbf{x}$, witness $\mathbf{w}$, and proof π . However, the adversary can only get more information about the state st if the polynomial $p^* - G(f_1^*, \dots, f_{\ell-2}^*)$ is not identically zero similar to the case of the FINALIZE procedure discussed above. Additionally, we can assume that an adversary submits queries such that this polynomial has distinct roots (including roots of polynomials corresponding to previous queries) to maximize information gain.

Claim: The probability of bad flag being set to 1 when we reach the FINALIZE procedure is given as follows:

$$\Pr[\text{bad} = 1 \text{ at FINALIZE}] \leq \frac{Qd}{|\mathbb{F}| - (M + Qd)}$$

where the probability is taken over the coins of the INITIALIZE procedure.

The bad flag is set to 1 when we reach the FINALIZE procedure if for at least one $j \in [Q]$, the state st is a root of the polynomial $p^j - G(f_1^j, \dots, f_{\ell-2}^j)$ corresponding to the j^{th} query $(\mathbf{x}_j \parallel \pi_j)$ made by the adversary $\mathcal{A}$ to the REPEATO oracle. The probability of

this event can be calculated as follows:

$$\begin{aligned}
\Pr[\text{bad} = 1 \text{ at FINALIZE}] &= \sum_{j \in [N]} \Pr[\text{bad} = 1 \text{ for the first time at } j] \\
&= \sum_{j \in [Q]} (\Pr[\text{bad} = 0 \text{ at } j-1] \cdot \Pr[\text{bad} = 1 \text{ at } j]) \\
&= \sum_{j \in [Q]} \left(\prod_{k < j} \left(1 - \frac{d}{|\mathbb{F}| - (M + kd)} \right) \cdot \frac{d}{|\mathbb{F}| - (M + jd)} \right) \\
&\leq \sum_{j \in [Q]} \frac{d}{|\mathbb{F}| - (M + jd)} \\
&\leq \frac{Qd}{|\mathbb{F}| - M - Qd}.
\end{aligned}$$

The third equality follow from the fact that at query j , the probability that the state is a root of the corresponding polynomial $p^j - G(f_1^j, \dots, f_{\ell-2}^j)$ is given by $\frac{d}{|\mathbb{F}| - (M + jd)}$.

Claim: For an adversary $\mathcal{A}$, the following holds:

$$\Pr[G_{\Pi, \mathcal{A}}^{\text{reuse-sound}}(\lambda, d) \Rightarrow 1 | \text{bad} = 0 \text{ at FINALIZE}] \leq \frac{d}{|\mathbb{F}| - (M + Qd)}$$

where Q is the total number of queries made to the REPEAT oracle and the probability is taken over the coins of the INITIALIZE procedure.

In game $G_{\Pi, \mathcal{A}}^{\text{reuse-sound}}$, when the output is 1 the state value st is the root of the polynomial $p^* - G(f_1^*, \dots, f_{\ell-2}^*)$ corresponding to $\mathbf{x}^* \parallel \pi^*$, as previously established. From the construction of the game, we know that st is sampled uniformly at random from the set $\mathbb{F} \setminus [M]$. If the bad flag is 0, then the adversary did not find the state st during one of its REPEAT queries. In more detail, there does not exist a query $(\mathbf{x}_j \parallel \pi_j)$ to the REPEAT oracle such that the state st is a root of the corresponding polynomial $p^j - G(f_1^j, \dots, f_{\ell-2}^j)$. Thus, using similar arguments as the proof of soundness of BBG⁺ and the Schwartz-Zippel lemma, the probability of this event is given by $\frac{d}{|\mathbb{F}| - (M + Qd)}$.

Finally, combining the two scenarios of the value of bad at FINALIZE, we get that the soundness parameter γ is $\frac{Q(d+1)}{|\mathbb{F}| - (M + Qd)}$ as claimed. $\square$

REMARK A.1. The FL-PCP described in Section 5 and Fig. 9 satisfies reusable soundness the proof of which is identical to the above proof.

A.2 Strong Honest-Verifier Zero-Knowledge

LEMMA 2. BBG⁺ satisfies reusable strong HVZK as defined in Section 4.3.

PROOF. In order to prove that BBG⁺ satisfies reusable strong HVZK, we first provide a simulator $\mathcal{S}_{\Pi, \mathcal{R}} = (\mathcal{S}_1, \mathcal{S}_2)$ in Figure 16. Then, we show that when using this simulator, the output of the two experiments are identically distributed.

The output of the Initialize procedure in the game $G_{\Pi, \mathcal{A}}^{\text{rshvzk-1}}$ being distributed identically to that in the game $G_{\Pi, \mathcal{A}}^{\text{rshvzk-0}}$ follows from the fact that the output state st is sampled uniformly at random from the same domain $\mathbb{F} \setminus [M]$ in both games.

Next, we want to show that the responses from the Repeat oracle are also identically distributed in both games. This can be done

Simulator $\mathcal{S}$

$\mathcal{S}_1(1^\lambda, 1^d)$:

- 1 $r \leftarrow \mathbb{F} \setminus [M]$
- 2 $\text{st} \leftarrow r$
- 3 Return st

$\mathcal{S}_2(1^\lambda, 1^d, \text{st})$:

- 4 If $\text{st} \notin \mathbb{F} \setminus [M]$:
- 5 Return $\perp$
- 6 $a_1, \dots, a_{\ell-2} \leftarrow \mathbb{F}$
- 7 $a \leftarrow G(a_1, \dots, a_{\ell-2}) \in \mathbb{F}$
- 8 Return $(a_1, \dots, a_{\ell-2}, a, 0) \in \mathbb{F}^{\ell+1}$

Figure 16: Simulator for reusable strong HVZK for BBG⁺.

using an argument similar to the one made in the proof of zero-knowledge in BBG⁺. First, we recall that an honest verifier samples r in order to check that the polynomials p and $G(f_1, \dots, f_{\ell-2})$ are constructed correctly by evaluating them at r using query vectors $\mathbf{q}_1, \dots, \mathbf{q}_\ell$. Among $(a_1, \dots, a_{\ell-2}, a_{\ell-1}, a_\ell)$, $a_\ell = 0$ and $a_{\ell-1} = G(a_1, \dots, a_{\ell-2})$ in the output of both games when $\mathbf{x} \in \mathcal{R}$. Since these values are identical, it suffices to show that the distribution of $(a_1, \dots, a_{\ell-2})$ is identical in the two games.

Here, $(a_1, \dots, a_{\ell-2})$ in the game $G_{\Pi, \mathcal{A}}^{\text{rshvzk-0}}$ simulates the values $(f_1(r), \dots, f_{\ell-2}(r))$ in the game $G_{\Pi, \mathcal{A}}^{\text{rshvzk-1}}$. For every $i \in [\ell-2]$, $f_i(r)$ can be written in terms of the Lagrange interpolating polynomials:

$$f_i(r) = \lambda_0(r) \cdot f_i(0) + \sum_{j=1}^M \lambda_j(r) \cdot f_i(j).$$

When $r \notin [M]$, the value of the zero-th interpolating polynomial $\lambda_0(r)$ is non-zero. We also know that each $f_i(0)$ is sampled uniformly from $\mathbb{F}$ in the ProofGen algorithm for each query $(\mathbf{x}, \mathbf{w})$. Thus, $\lambda_0(r) \cdot f_i(0)$ is uniform random and the distribution of $f_i(r)$ is uniform random as well.

On the other hand, the simulator in $\mathcal{S}_2$ samples $a_1, \dots, a_{\ell-2} \in \mathbb{F}$ uniformly at random. Therefore, the distribution of $(a_1, \dots, a_{\ell-2})$ is identical in the two games.

Hence, the responses from the Repeat oracle are identically distributed in both games and the advantage of any adversary is 0. $\square$

REMARK A.2. The FL-PCP described in Section 5 and Figure 9 satisfies reusable strong HVZK, the proof of which is identical to the above except that $\mathcal{S}_1$ also samples an additional vector $\psi \leftarrow \mathbb{F}^k$ and returns it as part of its output.

B PEAR Proof of Privacy

Here we prove Theorem 1 showing that PEAR is private against a malicious server. Let $\mathcal{A}$ be an adversary against the privacy of PEAR. Consider the game sequence in Fig. 17. The differences between consecutive games are highlighted in red. The first game $G_{\text{PEAR}, \mathcal{A}}^1$ is the indistinguishability game from Fig. 3. In the second game $G_{\text{PEAR}, \mathcal{A}}^2$, instead of encrypting $(\mathbf{x}^b \parallel \pi)$, we compute its inner product with each query vector, sample a vector $\mathbf{v}^*$ that has the same inner products with the query vectors, and encrypt that vector instead. The transition between these two games is justified using security of the IPFE scheme Δ and preimage sampleability of the inner product functionality. Next, in game $G_{\text{PEAR}, \mathcal{A}, \mathcal{S}}^3$, we use the FL-PCP simulator to generate the inner product values and sample

a vector that has these inner products with the query vectors. This game is independent of the bit b .

The following lemmas correspond to the games:

LEMMA 3. *For all $\lambda, d \in \mathbb{N}$, there exists an adversary $\mathcal{B}$ such that:*

$$\Pr[G_{\text{PEAR}, \mathcal{A}}^1(\lambda, d) \Rightarrow 1] - \Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda, d) \Rightarrow 1] = \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d).$$

LEMMA 4. *There exists a simulator S_{Π} such that the following equation holds:*

$$\Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda, d) \Rightarrow 1] = \Pr[G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda, d) \Rightarrow 1].$$

Finally, we can see that in game $\Pr[G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda)]$, the output of the game is independent of the bit b . Thus,

$$\Pr[G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda, d) \Rightarrow 1] = \frac{1}{2}.$$

Next, we combine this fact with the lemmas as follows:

$$\begin{aligned} \Pr[G_{\text{PEAR}, \mathcal{A}}^{\text{ind}}(\lambda) \Rightarrow 1] &= \Pr[G_{\text{PEAR}, \mathcal{A}}^1(\lambda) \Rightarrow 1] \\ &= \left(\Pr[G_{\text{PEAR}, \mathcal{A}}^1(\lambda) \Rightarrow 1] - \Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda) \Rightarrow 1] \right) \\ &\quad + \Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda) \Rightarrow 1] \\ &= \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d) + \Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda) \Rightarrow 1] \\ &= \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d) + \Pr[G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda) \Rightarrow 1] \\ &= \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d) + \frac{1}{2}. \end{aligned}$$

Finally,

$$\begin{aligned} \text{Adv}_{\text{PEAR}, \mathcal{A}}^{\text{ind}}(\lambda, d) &= 2 \cdot \Pr[G_{\text{PEAR}, \mathcal{A}}^{\text{ind}}(\lambda) \Rightarrow 1] - 1 \\ &= \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d). \end{aligned}$$

Now, it suffices to prove Lemmas 3 and 4.

PROOF OF LEMMA 3. First, consider adversary $\mathcal{B}$ provided in Fig. 18. The adversary simulates the encryption oracle for the adversary $\mathcal{A}$ by generating both messages $(\mathbf{x}^b || \boldsymbol{\pi})$ and $\mathbf{v}^*$ corresponding to the two games and making an IPFE encryption query with them. By inspection, we can see that when $\mathcal{B}$ has its challenge bit $\text{bit} = 0$, it perfectly simulates game $G_{\text{PEAR}, \mathcal{A}}^1$ in Fig. 17.

On the other hand, by definition of VecSamp, the inner products of vector $\mathbf{v}^*$ with every vector $\mathbf{q}_j$ is equal to the inner products of vector $(\mathbf{x}^b || \boldsymbol{\pi})$ with every $\mathbf{q}_j$ satisfying the validity requirement of the encryption oracle. Here we use preimage sampleability of the inner product function in the VecSamp to compute the vector $\mathbf{v}^*$. Thus, when $\mathcal{B}$ has its challenge bit $\text{bit} = 0$, it perfectly simulates $G_{\text{PEAR}, \mathcal{A}}^2$ for $\mathcal{A}$. Finally,

$$\Pr[G_{\text{PEAR}, \mathcal{A}}^1(\lambda, d) \Rightarrow 1] - \Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda, d) \Rightarrow 1] = \text{Adv}_{\text{IPFE}, \mathcal{B}}^{\text{ind}}(\lambda, d). \quad \square$$

PROOF OF LEMMA 4. The difference between games $G_{\text{PEAR}, \mathcal{A}}^2$ and $G_{\text{PEAR}, \mathcal{A}, S}^3$ is in the generation of the state st_{Π} and the inner product values $a_1, \dots, a_{\ell}$. The existence of a simulator $S_{\Pi, \mathcal{R}} = (S_1, S_2)$ whose output is identically distributed to the real state st_{Π} and inner product values $a_1, \dots, a_{\ell}$ is guaranteed by the reusable strong

HVZK property when the input instances are valid. Thus, if the distribution of these values is identical, then the output of the games is also identical for any adversary $\mathcal{A}$. Hence,

$$\Pr[G_{\text{PEAR}, \mathcal{A}}^2(\lambda, d) \Rightarrow 1] = \Pr[G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda, d) \Rightarrow 1]. \quad \square$$

C FL-PCP Optimization Proofs

Here, we show that this modified version of the FL-PCP instance satisfies Completeness, Soundness, and Strong honest-verifier zero-knowledge.

CLAIM C.1. *(Completeness.) Π' is complete. That is, for all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$, the verifier always accepts the proof. More precisely, for all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$, $\Pr[\text{Decision}(\text{st}, a_1, \dots, a_{L+1}) \Rightarrow 1] = 1$ where $\boldsymbol{\pi} \leftarrow \text{ProofGen}(1^{\lambda}, \mathbf{x}, \mathbf{w})$, $(\mathbf{q}_1, \dots, \mathbf{q}_{L+1}, \text{st}) \leftarrow \text{Query}(1^{\lambda}, 1^d)$, and $a_i \leftarrow \langle \mathbf{x}, \boldsymbol{\pi}, \mathbf{q}_i \rangle$ for all $i \in [L+1]$. The probability is taken over the coins of the ProofGen and Query algorithms.*

PROOF. First, we know that by construction of the output polynomial p , if the prover honestly generates the polynomials $f_1, \dots, f_L, p$, then the first L linear checks pass. Additionally, we can see that if $x_i \in \mathcal{R}$, then for all $i \in [k] : C_i(\mathbf{x}, \mathbf{w}) = 0$. Hence, $\psi_i \cdot C_i(\mathbf{x}, \mathbf{w}) = 0$ for any ψ_i , and thus the sum always equals 0 as required. $\square$

CLAIM C.2. *(Soundness.) Π' is sound. More precisely, for $(\mathbf{x}, \mathbf{w}) \in \mathbb{F}^d \times \mathbb{F}^h$, for any claimed proof $\boldsymbol{\pi}$ such that $|\boldsymbol{\pi}| = |\text{ProofGen}(1^{\lambda}, \mathbf{x}, \mathbf{w})|$, $(\mathbf{q}_1, \dots, \mathbf{q}_{L+1}, \text{st}) \leftarrow \text{Query}(1^{\lambda}, 1^d)$, and $a_i \leftarrow \langle \mathbf{x}, \boldsymbol{\pi}, \mathbf{q}_i \rangle$ for all $i \in [L+1]$ if $\text{Decision}(\text{st}, a_1, \dots, a_{L+1})$ outputs 1, then $(\mathbf{x}, \mathbf{w}) \notin \mathcal{R}$ with probability at most $\frac{\deg G}{|\mathbb{F}| - M} + \frac{1}{|\mathbb{F}|}$. The probability is taken over the coins of the ProofGen and Query algorithms.*

PROOF. Using the soundness argument of BBG⁺, by the Schwartz-Zippel Lemma, we know that the probability that

$$p(r) - G(f_1(r), \dots, f_L(r)) = 0 \text{ for } r \leftarrow \mathbb{F} \setminus [M] \text{ is equal to } \frac{\deg G}{|\mathbb{F}| - M}.$$

Next, consider the probability that:

$$\sum_{i=1}^k (\psi_i \cdot C_i(\mathbf{x}, \mathbf{w})) = 0$$

where at least one of the $C_i(\mathbf{x}, \mathbf{w}) \neq 0$. This probability can be computed as follows:

$$\Pr \left[\psi_k \cdot C_k(\mathbf{x}, \mathbf{w}) = - \sum_{i=1}^{k-1} \psi_i \cdot C_i(\mathbf{x}, \mathbf{w}) \mid \psi_i \leftarrow \mathbb{F} \ \forall i \in [k-1] \right] = \frac{1}{|F|}.$$

Thus, with probability $\frac{1}{|\mathbb{F}|}$, $\sum_{i=1}^k \psi_i \cdot C_i(\mathbf{x}, \mathbf{w}) = 0$, when $(\mathbf{x}, \mathbf{w}) \notin \mathcal{R}$. The verifier accepts when $(\mathbf{x}, \mathbf{w}) \notin \mathcal{R}$ with the total probability at most $\frac{\deg G}{|\mathbb{F}| - M} + \frac{1}{|\mathbb{F}|}$. $\square$

CLAIM C.3. *(Strong honest-verifier zero-knowledge.) Π' is SHVZK. More precisely, there exists a simulator S such that for all $(\mathbf{x}, \mathbf{w}) \in \mathcal{R}$ and $\lambda, d \in \mathbb{N}$, the distribution of $S(1^{\lambda}, 1^d)$ is identical to the following:*

$$(\text{st}, \langle \mathbf{x} || \boldsymbol{\pi}, \mathbf{q}_1 \rangle, \dots, \langle \mathbf{x} || \boldsymbol{\pi}, \mathbf{q}_{L+1} \rangle), (\mathbf{q}_1, \dots, \mathbf{q}_{L+1})$$

where $(\mathbf{q}_1, \dots, \mathbf{q}_{L+1}, \text{st}) \leftarrow \text{Query}(1^{\lambda}, 1^d)$ and $\boldsymbol{\pi} \leftarrow \text{ProofGen}(1^{\lambda}, \mathbf{x}, \mathbf{w})$.

Game $G_{\text{PEAR}, \mathcal{A}}^1(\lambda, d)$	Game $G_{\text{PEAR}, \mathcal{A}}^2(\lambda, d)$	Game $G_{\text{PEAR}, \mathcal{A}, S}^3(\lambda, d)$
INITIALIZE ($1^\lambda, 1^d$): 1 $b \leftarrow \{0, 1\}$ 2 $(\text{mpk}, \text{msk}) \leftarrow \Delta.\text{Setup}(1^\lambda, 1^d)$ 3 $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}_\Pi) \leftarrow \Pi_{\mathcal{R}}.\text{Query}(1^\lambda, 1^d)$ 4 For $j \in [\ell]$: 5 $\text{sk}_{q_j} \leftarrow \Delta.\text{KeyGen}(\text{msk}, \mathbf{q}_j)$ 6 $\text{vk} \leftarrow (\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_\Pi)$ 7 Return (mpk, vk) ENCO ($(\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1)$): 8 $\pi \leftarrow \Pi_{\mathcal{R}}.\text{ProofGen}(1^\lambda, \mathbf{x}^b, \mathbf{w}^b)$ 9 $\text{ct} \leftarrow \Delta.\text{Encrypt}(\text{mpk}, (\mathbf{x}^b \parallel \pi))$ 10 return ct FINALIZE (b'): 11 return $(b = b')$	INITIALIZE ($1^\lambda, 1^d$): 1 $b \leftarrow \{0, 1\}$ 2 $(\text{mpk}, \text{msk}) \leftarrow \Delta.\text{Setup}(1^\lambda, 1^d)$ 3 $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}_\Pi) \leftarrow \Pi_{\mathcal{R}}.\text{Query}(1^\lambda, 1^d)$ 4 For $j \in [\ell]$: 5 $\text{sk}_{q_j} \leftarrow \Delta.\text{KeyGen}(\text{msk}, \mathbf{q}_j)$ 6 $\text{vk} \leftarrow (\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_\Pi)$ 7 Return (mpk, vk) ENCO ($(\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1)$): 8 $\pi \leftarrow \Pi_{\mathcal{R}}.\text{ProofGen}(1^\lambda, \mathbf{x}^b, \mathbf{w}^b)$ 9 For $i \in [\ell]$: $a_i \leftarrow \langle (\mathbf{x}^b \parallel \pi), \mathbf{q}_i \rangle$ 10 $\mathbf{v}^* \leftarrow \text{VecSamp}((\mathbf{q}_1, a_1), \dots, (\mathbf{q}_\ell, a_\ell))$ 11 $\text{ct} \leftarrow \Delta.\text{Encrypt}(\text{mpk}, \mathbf{v}^*)$ 12 return ct FINALIZE (b'): 13 return $(b = b')$ VecSamp ($(\mathbf{y}_1, a_1), \dots, (\mathbf{y}_\ell, a_\ell)$): 14 Let $\mathbf{v}^* \in \mathbb{F}^{d+m}$ 15 Solve the system of linear equations for $\mathbf{v}^*$: 16 $(\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_\ell)^\top \cdot \mathbf{v}^* = (a_1 \parallel \dots \parallel a_\ell)$ 17 return $\mathbf{v}^*$	INITIALIZE ($1^\lambda, 1^d$): 1 $b \leftarrow \{0, 1\}$ 2 $(\text{mpk}, \text{msk}) \leftarrow \Delta.\text{Setup}(1^\lambda, 1^d)$ 3 $\text{st}_\Pi \leftarrow S_1(1^\lambda, 1^d)$ 4 $(\mathbf{q}_1, \dots, \mathbf{q}_\ell) \leftarrow \Pi_{\mathcal{R}}.\text{Query}(1^\lambda, 1^d; \text{st}_\Pi)$ 5 For $j \in [\ell]$: 6 $\text{sk}_{q_j} \leftarrow \Delta.\text{KeyGen}(\text{msk}, \mathbf{q}_j)$ 7 $\text{vk} \leftarrow (\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_\Pi)$ 8 Return (mpk, vk) ENCO ($(\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1)$): 9 $(a_1, \dots, a_\ell) \leftarrow S_2(1^\lambda, 1^d, \text{st}_\Pi)$ 10 $\mathbf{v}^* \leftarrow \text{VecSamp}((\mathbf{q}_1, a_1), \dots, (\mathbf{q}_\ell, a_\ell))$ 11 $\text{ct} \leftarrow \Delta.\text{Encrypt}(\text{mpk}, \mathbf{v}^*)$ 12 return ct FINALIZE (b'): 13 return $(b = b')$ VecSamp ($(\mathbf{y}_1, a_1), \dots, (\mathbf{y}_\ell, a_\ell)$): 14 Let $\mathbf{v}^* \in \mathbb{F}^{d+m}$ 15 Solve the system of linear equations for $\mathbf{v}^*$: 16 $(\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_\ell)^\top \cdot \mathbf{v}^* = (a_1 \parallel \dots \parallel a_\ell)$ 17 return $\mathbf{v}^*$

Figure 17: Games for proof of privacy for Theorem 1.
Changes between consecutive games are depicted in red.

Adversary $\mathcal{B}^{\text{KEYGENO}'(\cdot), \text{ENCO}'(\cdot, \cdot)}(1^\lambda, 1^d, \text{mpk})$
INITIALIZE : 1 $b \leftarrow \{0, 1\}$ 2 $(\mathbf{q}_1, \dots, \mathbf{q}_\ell, \text{st}_\Pi) \leftarrow \Pi_{\mathcal{R}}.\text{Query}(1^\lambda, 1^d)$ 3 For $j \in [\ell]$: 4 $\text{sk}_{q_j} \leftarrow \text{KEYGENO}'(\mathbf{q}_j)$ 5 $\text{vk} \leftarrow (\{\text{sk}_{q_j}\}_{j \in [\ell]}, \text{st}_\Pi)$ 6 Return (mpk, vk) ENCO ($(\mathbf{x}^0, \mathbf{w}^0), (\mathbf{x}^1, \mathbf{w}^1)$): 7 $\pi \leftarrow \Pi_{\mathcal{R}}.\text{ProofGen}(1^\lambda, \mathbf{x}^b, \mathbf{w}^b)$ 8 For $i \in [\ell]$: $a_i \leftarrow \langle (\mathbf{x}^b \parallel \pi), \mathbf{q}_i \rangle$ 9 $\mathbf{v}^* \leftarrow \text{VecSamp}((\mathbf{q}_1, a_1), \dots, (\mathbf{q}_\ell, a_\ell))$ 10 $\text{ct} \leftarrow \text{ENCO}'((\mathbf{x}^b \parallel \pi), \mathbf{v}^*)$ 11 return ct FINALIZE (b'): 12 return $(b = b')$ VecSamp ($(\mathbf{y}_1, a_1), \dots, (\mathbf{y}_\ell, a_\ell)$): 13 Let $\mathbf{v}^* \in \mathbb{F}^{d+m}$ 14 Solve the system of linear equations for $\mathbf{v}^*$: 15 $(\mathbf{y}_1 \parallel \dots \parallel \mathbf{y}_\ell)^\top \cdot \mathbf{v}^* = (a_1 \parallel \dots \parallel a_\ell)$ 16 return $\mathbf{v}^*$

Figure 18: IND adversary for privacy proof for PEAR.

PROOF. Note that the queries of the verifier are determined by the state of the verifier and the circuits. So, it suffices to simulate $(\text{st}, f_1(r), \dots, f_L(r), \langle (\mathbf{x}, \pi), \mathbf{q}_{L+1} \rangle)$. The required simulator S is as follows:

- (1) $r \leftarrow \mathbb{F} \setminus [M]$.
- (2) $\psi \leftarrow \mathbb{F}^k$.
- (3) $a_1, \dots, a_L \leftarrow \mathbb{F}$.
- (4) $\text{st} \leftarrow (r, \psi)$
- (5) Return $(\text{st}, a_1, \dots, a_L, 0) \in \mathbb{F}^{L+k+1}$.

The distribution of $r, a_1, \dots, a_L$ sampled in the simulator is identical to the distribution of $r, f_1(r), \dots, f_L(r)$ in the real interaction. This argument closely resembles the proof of strong honest verifier zero knowledge for the FL-PCP given in BBG⁺. All the values are uniform random in their respective domains.

Next, we know that an honest verifier samples the vector ψ uniformly from $\mathbb{F}$, thus the query vector $\mathbf{q}_{L+1}$ is identically distributed to the real $(L+1)$ -st query.

Finally, when $\mathbf{x}$ is valid, the final linear query outputs 0, thus the final output of the simulator is also perfect.

Hence, the simulator's entire output is identically distributed to the real interaction and Π' satisfies the SHVZK property. $\square$

D Unrestricted IPFE

The IPFE scheme must be unrestricted because the FL-PCP protocol's Query algorithm needs to sample random query vectors to verify the input. This scheme is proposed in [17] based on class groups, which we describe below. CLT operates within a DDH-hard group containing a DL-easy subgroup, meaning that there exists a polynomial time algorithm capable of solving discrete logarithm problem in this subgroup. Let Gen be the group generator. We specify that on inputs two parameters λ and μ , Gen outputs a tuple $(p, \tilde{s}, g, f, g_p, G, F, G^p)$ where, the set $(G, \cdot)$ is a cyclic group of order ps , s is an integer, $\tilde{s}$ is an upper bound of s , p is a μ -bit prime, and $\gcd(p, s) = 1$. The set $G^p = \{x^p, x \in G\}$ is the subgroup of order s of G , and F is subgroup of order p of G , so that $G = F \times G^p$. Finally f, g_p and $g = f \cdot g_p$ are generators of F, G^p and G , respectively. Since the discrete logarithm problem is easy in F , we let Solve be a deterministic polynomial time algorithm that solves the discrete logarithm problem in F .

Let $\mathbf{x}, \mathbf{y} \in \mathbb{Z}_p^d$, we present the construction based on the HSM assumption in Figure 19, and demonstrate that it is additively homomorphic. In [17], CLT.KeyGen is designed as a stateful algorithm,

CLT.Setup($1^\lambda, 1^\mu, 1^d$): 1 $(p, \bar{s}, f, g_p, G, F, G^p) \leftarrow \text{Gen}(1^\lambda, 1^\mu)$ 2 $\mathbf{s} = (s_1, \dots, s_d) \leftarrow \mathcal{D}$ 3 For $1 \leq i \leq d$: 4 $h_i = g_p^{s_i}$ 5 $\text{mpk} \leftarrow (p, \bar{s}, f, g_p, \{h_i\}_{i \in [d]})$ 6 $\text{msk} \leftarrow \mathbf{s}$ 7 Return (mpk, msk) CLT.KeyGen($\text{msk}, \mathbf{y}$): 8 $\mathbf{s} \leftarrow \text{msk}$ 9 $\text{sk}_y = \langle \mathbf{s}, \mathbf{y} \rangle$ 10 Return sk_y CLT.Encrypt($\text{mpk}, \mathbf{x}$): 11 $r \leftarrow \mathcal{D}_p$ 12 $\text{ct}_0 \leftarrow g_p^r$ 13 For $1 \leq i \leq d$: 14 $\text{ct}_i = f^{x_i} \cdot h_i^r$ 15 Return $(\text{ct}_0, \text{ct}_1, \dots, \text{ct}_d)$	CLT.Eval($\text{mpk}, \text{ct}_x, \text{sk}_y$): 16 $(\text{ct}_0, \text{ct}_1, \dots, \text{ct}_d) \leftarrow \text{ct}_x$ 17 $\text{ct} = \left(\prod_{i=1}^d \text{ct}_i^{y_i} \right) \cdot \text{ct}_0^{-\text{sk}_y}$ 18 Return $\text{Solve}(\text{ct})$ CLT.Decrypt(msk, ct_x): 19 $(\text{ct}_0, \text{ct}_1, \dots, \text{ct}_d) \leftarrow \text{ct}_x$ 20 $\mathbf{s} \leftarrow \text{msk}$ 21 For $1 \leq i \leq d$: 22 $\mathbf{x}_i \leftarrow \text{Solve}(\text{ct}_i \cdot \text{ct}_0^{-s_i})$ 23 Return $(\mathbf{x}_1, \dots, \mathbf{x}_d)$ CLT.Add($\text{ct}_{x_1}, \text{ct}_{x_2}$): 24 $(\text{ct}_{1,0}, \text{ct}_{1,1}, \dots, \text{ct}_{1,d}) \leftarrow \text{ct}_{x_1}$ 25 $(\text{ct}_{2,0}, \text{ct}_{2,1}, \dots, \text{ct}_{2,d}) \leftarrow \text{ct}_{x_2}$ 26 For $0 \leq i \leq d$: 27 $\text{ct}_i = \text{ct}_{1,i} \cdot \text{ct}_{2,i}$ 28 Return $(\text{ct}_0, \text{ct}_1, \dots, \text{ct}_d)$
--	---

Figure 19: Unrestricted IPFE construction instantiated by HSM in [17]

ensuring that keys cannot be queried for vectors that are linearly dependent over $\mathbb{Z}_p^d$ but independent over $\mathbb{Z}^d$. In our usage scenario, the analyst generates all inputs to CLT.KeyGen at once and can independently verify them before invoking CLT.KeyGen. Hence the stateful KeyGen is not necessary. Furthermore, we demonstrate that the scheme possesses the desired additive homomorphism property, as described in Section 3.3, enabling the server to aggregate ciphertexts and output their sum.

ADDITIVE HOMOMORPHISM OF CLT. For all $\lambda, \mu, d \in \mathbb{N}$, $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{Z}_p^d$, and $(\text{mpk}, \text{msk}) \leftarrow \text{CLT.Setup}(1^\lambda, 1^\mu, 1^d)$, we let $\text{ct}_{x_1} \leftarrow \text{CLT.Encrypt}(\text{mpk}, \mathbf{x}_1)$, $\text{ct}_{x_2} \leftarrow \text{CLT.Encrypt}(\text{mpk}, \mathbf{x}_2)$, and $\text{ct} \leftarrow \text{CLT.Add}(\text{ct}_{x_1}, \text{ct}_{x_2})$. We then have that $\text{ct}_0 = g_p^{r_1+r_2}$, $\text{ct}_i = f^{x_{1,i}+x_{2,i}} h_i^{r_1+r_2}$ for $1 \leq i \leq d$. Since r_1 and r_2 are randomly sampled, $\text{CLT.Add}(\text{ct}_{x_1}, \text{ct}_{x_2})$ is computationally indistinguishable from $\text{CLT.Encrypt}(\text{mpk}, \mathbf{x}_1 + \mathbf{x}_2)$.

INSTANTIATION DETAILS. As shown in [17], in practice, the distributions $\mathcal{D}$ and $\mathcal{D}_p$ can be implemented from the output of Gen more efficiently. Let $\mathcal{D}_\sigma$ represents the discrete Gaussian distribution over $\mathbb{Z}$ with parameter σ centered at 0. We choose $\mathcal{D} = \mathcal{D}_{\bar{s} \cdot p \cdot \sqrt{\lambda}}$, and $\mathcal{D}_p = \mathcal{D}_{\bar{s} \cdot \sqrt{\lambda}}$. The CLT scheme is instantiated using class groups of imaginary quadratic fields. We refer to [16, 17] for a full description of the implementation. We set both λ the security parameter and μ the size of prime p to 128. This results in the hidden order group G_p having 924-bit elements and those in G having 2084-bit elements.