

Verifiable Distributed Aggregation Functions

Hannah Davis
University of California, San Diego

Christopher Patton
Cloudflare

Mike Rosulek
Oregon State University

Phillipp Schoppmann
Google

2023/02/03

Abstract

The modern Internet is built on systems that incentivize collection of information about users. In order to minimize privacy loss, it is desirable to prevent these systems from collecting more information than is required for the application. The promise of *multi-party computation* is that data can be aggregated without revealing individual measurements to the data collector. This work offers a provable security treatment for “Verifiable Distributed Aggregation Functions (VDAFs)”, a class of multi-party computation protocols being considered for standardization by the IETF.

We propose a formal framework for the analysis of VDAFs and apply it to two candidate protocols. The first is based on the Prio system of Corrigan-Gibbs and Boneh (NSDI 2017). Prio is fairly mature and has been deployed in real-world applications. We prove that, with only minor changes, the current draft of the standardized version achieves our security goals. The second candidate is the recently proposed Poplar system from Boneh et al. (IEEE S&P 2021). The deployability of Poplar is less certain. One difficulty is that the interactive step requires two rounds of broadcast messages, whereas Prio requires just one. This makes Poplar less suitable for many deployment scenarios. We show the round complexity can be improved, at the cost of higher bandwidth.

Change Log

- 2023/02/03: Initial ePrint submission.

1 Introduction

Operating a complex software system, such as an operating system, web browser, or web service, often requires measuring the behavior of the system’s users. When used for a specific purpose, such measurements are often only consumed in some aggregated form, e.g., $F(m_1, \dots, m_{ct})$ for some specific function F , rather than the individual measurements m_1, \dots, m_{ct} . But in conventional systems, the measurements are revealed to the operator as a matter of course, resulting in an increased capability to surveil users. Consider the following motivating examples:

1. *Identifying misbehaving or malicious origins.* To detect bugs or attack vectors, a browser vendor might want to know how often establishing a connection to a given origin or loading a given web page triggers a specific event [47]. But logging these events and aggregating them in the clear risks exposing browser history.
2. *Measuring ad conversion rates.* Today advertising is an important revenue source for many web service providers. In order to accurately assess the value of an ad campaign, the service provider and advertiser might want to measure how many people who clicked on a given ad made a purchase [2].
3. *Classifying malicious client behavior.* Many operators benefit from the ability to classify (or predict) user behavior automatically, and in real-time. For example, anomaly detection systems use machine

learning models, trained and validated on requests from real clients, to classify fraudulent or otherwise malicious behavior [45].

These applications require only aggregates; by collecting individual measurements, the operator learns more information than is ultimately used for the intended purpose. One way out of this predicament is *multi-party computation (MPC)*, which allows computing some function of private inputs distributed across multiple parties, without revealing these private inputs. In this paper, we consider a class of MPC protocols in which the bulk of the computation is outsourced to a small set of non-colluding servers.

Recent attention from the MPC community on problems like these has yielded solutions that are practical enough for real-world deployment [31, 23, 17, 18, 5, 10]. Notable deployments include Mozilla’s Origin Telemetry project [47] and the COVID-19 Exposure Notification Private Analytics system developed jointly by Apple and Google [7]. The success of these projects spurred the formation of a working group within the Internet Engineering Task Force (IETF) whose objective is to standardize MPC for “Privacy-Preserving Measurement (PPM)” [1], thereby improving interoperability and providing a deployment roadmap for new schemes.

The primary goal of this paper is to lay the groundwork for the provable security analysis that will be needed to support this effort. We formalize a syntax and set of security definitions for a particular class of MPC protocols from the literature [23, 17, 18, 5] that conform to the operational and security considerations identified by the working group. Our definitions unify previous ones into an explicit, game-based framework that accounts for practical matters not attended to in prior work.

We apply our definitional framework to two constructions. The first is a candidate for standardization based on the Prio protocol due to Corrigan-Gibbs and Boneh [23]. We show that this protocol meets our security goals with only minor changes. Another candidate for standardization is the more recent Poplar protocol due to Boneh et al. [18]; we introduce and analyze a variant of this protocol that has improved round complexity.

Overview. The PPM working group plans to develop multiple protocol standards, one of which is the focus of this work. The *Distributed Aggregation Protocol (DAP)* standard [29] centers around the execution of a particular class of MPC protocols, called *Verifiable Distributed Aggregation Functions (VDAFs)* [9]. A VDAF is used to securely compute some **aggregation function** F over a set of measurements generated by the **clients**. To protect their privacy, the measurements are secret-shared and the computation of the aggregate is distributed amongst multiple, non-colluding aggregation servers (called **aggregators** hereafter). Execution of a VDAF involves four basic steps (illustrated in Figure 1):

- **Shard:** Each client shards its measurement m_i into **input shares** and sends one share to each aggregator. In this work, we sometimes refer to this sequence of input shares the client’s **report**.
- **Prepare:** After receiving a report from a client, the aggregators gossip amongst themselves in order to prepare their shares for aggregation. This involves refining the shares into an aggregatable form and verifying that the shares are “well-formed”, e.g., they correspond to an integer in a given range, or correspond to a one-hot vector (a vector that is non-zero in at most one position). We call the outputs of this process the **refined shares**.
- **Aggregate:** Once an aggregator has recovered the desired number of refined shares, it combines them into its share of the aggregate result, called an **aggregate share**. It then sends this to the data consumer, known as the **collector**.
- **Unshard:** Finally, the collector combines each of the aggregate shares into $F(m_1, \dots, m_{ct})$.

The key feature that makes VDAFs well-suited for Internet scale is that they are “streamable”, in the sense that the interactive part of the computation is fully parallelizable across all reports being aggregated. This also means that VDAFs can only be used for aggregation functions F that can be decomposed into f, g for which $F(m_1, \dots, m_{ct}) = f(g(m_1), \dots, g(m_{ct}))$, where g is the online computation (Shard/Prepare) and f is the offline computation (Aggregate/Unshard). Indeed, the goal is not to encompass all possible MPC schemes, but a particular, useful class of them. VDAFs can be used for a variety of aggregation tasks, including: Simple summary statistics, such as sum, mean, or standard deviation, quantile estimates, linear regression [23]; a step of a gradient descent [36]; and heavy hitters (see below).




Figure 1: Illustration of (left) sharding and preparation of a single measurement and (right) aggregation and unsharding of a set of measurements. All parameters are defined in Section 3.

Security goals. The primary goal for VDAFs is that they are **private** in the sense that the attacker learns nothing about the measurements m_1, \dots, m_{ct} beyond what it can infer from the aggregate result $F(m_1, \dots, m_{ct})$. We consider an active attacker who corrupts the collector and a fraction of the aggregators (typically all but one) and controls transmission of all messages in the protocol—except, of course, the input shares delivered to honest aggregators. Its corruptions are “static” in that the set of corrupt parties does not change over the course of the attack.

Another security consideration for VDAFs is that they are **robust** in the sense that the attacker cannot force the collector to compute anything other than the aggregate of honestly generated reports. Here the main consideration is malicious clients attempting to corrupt the aggregate result by sending malformed reports. For robustness we assume all of the aggregators execute the protocol correctly. (Otherwise, a corrupt aggregator could trivially corrupt the result by sending the collector a malformed aggregate share; there may be practical ways to prevent this, but doing so is out-of-scope for this work.)

We formalize these security notions in the game-playing paradigm [14]. First, in Section 3.2 we define privacy via an indistinguishability game $\text{Exp}_H^{\text{priv}}(A)$ played by an attacker A against VDAF H . The attacker interacts with the honest parties (i.e., the clients and uncorrupted aggregators) via a set of oracles. These oracles allow A to mount a kind of “chosen batch attack” in which the honest parties process one of two batches of measurements, and A ’s goal is to determine which was processed.

We formalize robustness via a game $\text{Exp}_H^{\text{robust}}(A)$ (Section 3.2). Here the attacker A —playing the role of a coalition of malicious clients—is given a single oracle that models the execution of the preparation step of VDAF execution on (invalid) reports. The attacker wins if an aggregator ever accepts an invalid share.

Note on the UC paradigm. An alternative approach, and one that is more conventional for MPC, is to formulate security in the Universal Composability (UC) framework [21]. This methodology would begin by specifying the “ideal functionality” for computing an aggregation function such that, for any VDAF that securely realizes this functionality, any suitable notion of either privacy or robustness would follow from the UC composition theorem.

While this methodology is attractive, it creates the following difficulty in our setting. Many applications of VDAFs may be willing to tolerate a loose robustness bound (i.e., a non-negligible probability of accepting an invalid share) if doing so leads to better performance or communication. On the other hand, no application can accept a loose bound for privacy. In order to reason about this tradeoff, it is necessary to obtain explicit, concrete bounds for privacy and robustness *separately*. A theorem in the UC framework yields only a single bound, for the “UC-realizability” of the ideal functionality; applying this result directly would lead to parameter choices that might be more conservative than strictly necessary for the given application.

Another consideration is to make our results accessible to the target audience. Applying the UC framework, and interpreting its results, involves a number of subtleties that can be easily misinterpreted when translated to practice.¹ One goal of our definitions is to make as explicit as possible all of the requirements an application like DAP [29] needs to meet in order to use VDAFs securely.

¹For a recent example, consider the standards for PAKEs (“Password-Authenticated Key Exchange”) developed by the CFRG. Most of these standards are based on protocols with analysis in the UC framework. For one protocol [4], one question left open by that analysis was how to securely instantiate the “session identifier”, one of the artifacts of the ideal functionality. The current draft offers recommendations for choosing the session identifier, but allows applications to ignore this entirely; a game-playing argument was used to justify this (cf [3, Appendix B]).

Constructions. The starting point for our work is draft-irtf-cfrg-vdaf-03 [9], the current draft of the VDAF specification at the time of writing. This document describes two VDAFs.

The first scheme, called Prio3, is based on Prio [23], but incorporates performance improvements from Boneh et al. [17] (hereafter BBCG+19). Prio3 can be used to compute a wide variety of aggregation functions due to its use of a *Fully Linear Proof (FLP)*. Briefly, an FLP is a special class of zero-knowledge proof that allows the client’s input measurement to be validated by the aggregators (e.g., ensure that it is a number in some pre-determined range) who have only secret shares of the input and proof. The FLP designed by BBCG+19 (see [17, Theorem 4.3]) and adopted by the draft (with minor modifications; see [9, Section 7.3]) is expressed in terms of some arithmetic circuit C that takes in the prover’s input x and a random string jr computed jointly by the prover and verifier. Computing this joint randomness, verifying the proof, and evaluating $C(x, jr)$ requires just one round of communication among the aggregators.

In Section 4, we prove Prio3 is both robust (Theorem 1) and private (Theorem 2) under the assumption that the underlying FLP is, respectively, *sound* and *honest-verifier zero-knowledge* as defined by BBCG+19. Our analysis unveiled a few subtle design issues in draft-03 that result in loose security bounds, which we address here.

The second scheme in the VDAF specification is called Poplar1 and is based on the recent Poplar protocol from Boneh et al. [18] (BBCG+21). Poplar is designed to solve the private “heavy hitters” problem in which each client submits an arbitrary bitstring α and the collector wants to compute the set of unique strings that occurred at least T times. The key idea of BBCG+21 is an extension of *distributed point functions (DPFs)* [30], where two aggregators hold a share of a “DPF key” that concisely represents a *point function*. A point function evaluates to 0 on every input, except for the distinguished point α , where the function evaluates to some $\beta \neq 0$. By secret sharing the DPF keys generated by the clients, the aggregators can count *how many* clients submitted a particular candidate string without revealing *which* clients submitted it.

Poplar1 makes use of an enriched primitive that BBCG+21 call an *incremental DPF (IDPF)*. IDPF keys can be queried not only at a given point, but a given *prefix*. That is, an *incremental point function* is one that evaluates to 0 on every input except for the set of strings that are a prefix of α . This new primitive gives rise to an efficient solution to the heavy hitters problem that involves running Poplar1 multiple times over the same set of IDPF keys, where each run begins with a set of candidate prefixes computed from the previous run.

To achieve robustness, Poplar1 uses a two-round multi-party computation in which the aggregators verify that the IDPF outputs are well-formed. That means that, compared to Prio3, the Poplar1 VDAF costs one additional round of communication, per report, during the preparation phase. The additional per-report latency is significant from an operational perspective.

In Section 5 we introduce *Doplar*, our modification to Poplar which achieves a one-round preparation. To achieve this, we combine FLPs and methods from distributed point functions in a novel way. We adopt a point-function verification method from De Castro and Polychroniadou [22], filling some minor gaps in their analysis along the way. We also introduce a new flavor of *delayed-input* FLPs, which may be of independent interest.

Related Work. Several works have considered private aggregate statistics, relying either on secret-sharing between non-colluding servers [24, 26, 28, 39, 41, 43], or on anonymization networks [49, 34, 20]. However, these works either do not provide privacy against malicious clients or rely on expensive zero-knowledge proofs.

A protocol for Secure Aggregation (SecAgg) in the single-server setting was presented by Bonawitz et al. [16] and subsequently improved by Bell et al. [12, 11]. While SecAgg can provide security against malicious parties, it relies on multiple rounds of interaction between clients and server.

The VDAF abstraction was designed to encompass the “streamable” MPC architecture of Prio and Poplar in which the interactive portion of the computation is fully parallelizable. Another example of a VDAF from the literature is the protocol of Addanki et al. [5], which uses boolean (bit-wise) secret sharing instead of arithmetic circuit to improve communication cost from client to aggregator. However, this comes at a cost of weaker privacy, since their protocol does not protect against malicious servers.

There are also protocols that do not fit neatly into the VDAF framework as specified, but which might be adapted into VDAFs in the future. Masked LARK [36] is a proposal by Microsoft for training machine learning models on private data, using secret-sharing and MPC between a set of aggregators. AdScale [31]

presents an aggregation system focused on private ads measurement. While designed for a single aggregation server, their construction appears to be amenable to our multi-server setting.

Other protocols in the literature share the same security goals of VDAFs, but do not have the same streaming architecture. One example is the recent “Oblivious Shuffling” protocol due to Anderson et al. [6], which involves an MPC, assisted by a third-party, for unlinking each report from the client that sent it. The online processing for this procedure intrinsically involves all of the reports being shuffled; for VDAFs, all of the online processing is per-report. Similarly, Bell et al. [10] present a protocol for computing sparse histograms with two aggregators that is more efficient than DPFs for large domains, but reveals differentially private views to the aggregators. Again, the protocol crucially relies on shuffling contributions from multiple users. Vogue [38] is a protocol for computing private heavy hitters using three non-colluding servers. The protocol is secure against malicious servers and clients, but again relies on shuffling. Finally, the STAR protocol [25] uses an anonymizing proxy to ensure the collector only learns “popular” measurements, while any measurement that occurs less than a pre-determined threshold is not revealed to any party.

In recent concurrent work, Mouris et al. [46] present another three-party, honest-majority protocol for computing heavy hitters. Their full protocol relies on a secure comparison protocol that is run after the aggregation phase, and thus doesn’t immediately fit our setting. However, we believe their input validation protocol can be adapted to obtain a VDAF for heavy hitters that has similar characteristics as our protocol in Section 5. (Indeed their core primitive, which they also call “Verifiable IDPF”, bears a striking resemblance to our own VIDPF abstraction.) Likewise, one could get robustness against malicious aggregators in the honest-majority setting by applying their “duplicate aggregator” technique to our protocols. We leave an exploration of how to combine our results to future work.

2 Preliminaries

This section describes cryptographic primitives on which our constructions are based. We begin with a bit of non-standard notation.

Notation. Let $[i..j]$ denote the set of integers $\{i, \dots, j\}$ and write $[i]$ as shorthand for the set $[1..i]$. If \vec{v} is a vector, let $\vec{v}[i]$ denote the i -th element of \vec{v} . Let $(x,)$ denote the singleton vector with value x and $()$ the empty vector.

In our pseudocode, all variables that are undeclared implicitly have the value \perp . Let $y \leftarrow_s \mathcal{S}$ denote sampling y uniformly from a finite set \mathcal{S} ; let $y \leftarrow_s A(x)$ denote execution of randomized algorithm A ; and let $y \leftarrow A(x; r)$ denote execution of randomized algorithm A with coins r . If X is a random variable with support $\{0, 1\}$ we let $\Pr[X]$ denote the probability that $X = 1$.

A table T is a map from unique keys to values; we write $T[K_1, \dots]$ to denote the value corresponding to key K_1, \dots . We sometimes write a dot “.” in place of one of the elements of the key, e.g., “ $T[K_1, \cdot]$ ” instead of “ $T[K_1, K_2]$ ”. We use this notation to denote the vector of values in the table that match the key pattern. For example, we write $T[K_1, \cdot]$ for the vector $(T[K_1, K_2^1], \dots, T[K_1, K_2^n])$ where $(K_1, K_2^1), \dots, (K_1, K_2^n)$ are all of the keys in the table prefixed by K_1 , in lexicographic order.

We measure an adversary’s runtime by the time it takes to run its experiment to completion, including evaluating its queries.

Pseudorandom Generators. The VDAF spec [9, Section 6.2] calls for a particular type of object they call “pseudorandom generator (PRG)”. Unlike the conventional PRGs, these objects are stateful. A PRG is comprised of the following algorithms:

- $\text{PRG.Init}(seed \in \{0, 1\}^\kappa, cntxt \in \{0, 1\}^*) \rightarrow st \in \mathcal{Q}$ takes a seed and context string to the initial PRG state. We call κ the **seed length** and \mathcal{Q} the **state space**.
- $\text{PRG.Next}(st \in \mathcal{Q}, \ell \in \mathbb{N}) \rightarrow (st' \in \mathcal{Q}, out \in \{0, 1\}^\ell)$ takes in the current PRG state and outputs a string of the desired length.

We also make use of an algorithm $\text{Expand}[\text{PRG}]$ that uses the given PRG to map a seed, and context string to a vector of integers over the modular ring \mathbb{Z}_p for the desired modulus p . We defer to [9, Section 6.2] for the full definition of $\text{Expand}[\text{PRG}]$.

In our security proofs, we model PRGs as random oracles [13]. In some cases, such as the distributed point functions (DPFs) in Section 5.1, constructions based on computational assumptions are known to be sufficient. We refer to Guo et al. [32, 33] for an overview of the state-of-the-art PRGs for DPFs and similar constructions.

Fully Linear Proof Systems. We recall the definition of FLP systems from BBCG+19 [17]. (Our formulation differs slightly, as we discuss below.) FLPs allow a prover to prove to a verifier, in zero-knowledge, that a secret-shared value has some property required by the application, e.g., the input is a number in the desired range, is a one-hot vector, etc. (The main construction of BBCG+19 allows the validity condition to be expressed in terms of an arithmetic circuit evaluated over the input, similar to more conventional zero-knowledge proof systems.) They are “fully linear” in the sense that verifying the proof involves computing a strictly linear function over both the input and proof. Hence, verification can be performed on secret-shared data, leveraging its additive homomorphism property. (This is contrast to prior work on “linear PCPs” [8, 15, 37] in which the verifier has linear access to the proof, but arbitrary access to the input.)

An FLP with finite field \mathbb{F} , proof length m , verifier length v , prover randomness length pl , joint randomness length jl , and query randomness length ql is a triple of algorithms FLP defined as follows:

- $\text{FLP.Prove}(x \in \mathbb{F}^n, jr \in \mathbb{F}^{jl}) \rightarrow \pi \in \mathbb{F}^m$ is the randomized **proof-generation** algorithm that takes in an input x and joint randomness jr and outputs a proof string $\pi \in \mathbb{F}^m$. We shall assume this algorithm generates random coins by sampling uniformly from \mathbb{F}^{pl} .
- $\text{FLP.Query}(x \in \mathbb{F}^n, \pi \in \mathbb{F}^m, jr \in \mathbb{F}^{jl}) \rightarrow \sigma \in \mathbb{F}^v$ is the randomized **query-generation** algorithm that takes in an input x , proof string π , and joint randomness jr and outputs a verifier string σ . We shall assume the random coins are sampled uniformly from \mathbb{F}^{ql} .
- $\text{FLP.Decide}(\sigma \in \mathbb{F}^v) \rightarrow acc \in \{0, 1\}$ is the deterministic **decision predicate** that takes in a verifier string σ and outputs a bit acc indicating whether the input is valid.

We require the field \mathbb{F} to have prime order; we occasionally denote its order by $\mathbb{F}.p$.

We say that FLP is *fully linear* if the query-generation algorithm computes a linear function of the input and proof. That is, there exists a function Q whose output is a matrix in $\mathbb{F}^{v \times (n+m)}$ and, for all inputs x , proofs π , joint randomnesses jr , and query randomnesses qr , it holds that $\text{Query}(x, \pi, jr; qr) = Q(jr; qr) \cdot (x \parallel \pi) \in \mathbb{F}^v$.

Associated with FLP is a language $\mathcal{L} \subseteq \mathbb{F}^n$. We say that FLP is **complete for \mathcal{L}** if the proof system outputs 1 whenever the input is in \mathcal{L} . That is, for all $x \in \mathcal{L}$ it holds that

$$\Pr \left[\text{Decide}(\sigma) : jr \leftarrow \mathbb{F}^{jl}; \pi \leftarrow \text{Prove}(x, jr); \sigma \leftarrow \text{Query}(x, \pi, jr) \right] = 1.$$

We define soundness of FLP in terms of experiment $\text{Err}_{\text{FLP}}(P^*)$ shown in Figure 2 associated with a malicious prover P^* . In this experiment, the prover commits to an invalid input $x \in \mathbb{F}^n \setminus \mathcal{L}$. Next, joint randomness jr is generated and given to P^* , who then generates a proof π . Finally, the verifier is run on x, π, jr ; the malicious prover “wins” if the verifier deems the input valid. We say FLP is **ϵ -sound for \mathcal{L}** if for all P^* it holds that $\Pr \left[\text{Err}_{\text{FLP}}(P^*) \right] \leq \epsilon$.

Let $\text{View}_{\text{FLP}}(x)$ denote the procedure defined in Figure 2. We say FLP is **δ -statistical, strong, honest-verifier zero-knowledge**—or, simply, **δ -private**—if the verifier’s view can be simulated without knowledge of the input. That is, there exists a randomized algorithm S such that for all $x \in \mathcal{L}$ it holds that

$$\sum_{\omega} \left| \Pr \left[\text{View}_{\text{FLP}}(x) = \omega \right] - \Pr \left[S() = \omega \right] \right| \leq \delta.$$

Comparison to Boneh et al. [17]. Our syntax diverges slightly from BBCG+19 in two main respects. First, we have tailored the syntax to 1.5-round, public-coin IOP systems (cf. [17, Section 3.2]), as this is the only type of system considered in the VDAF specification [9]. Following the spec, we refer to the “random challenge” as the “joint randomness”, as this allows us to more easily distinguish the challenge from the randomness consumed locally by the prover and verifier. Second, following the VDAF specification [9], we

Algorithm $\text{View}_{\text{FLP}}(x)$:	Algorithm $\text{Err}_{\text{FLP}}(P^*)$:
1 $jr \leftarrow_{\$} \mathbb{F}^{jl}; qr \leftarrow_{\$} \mathbb{F}^{ql}$	5 $(st_{P^*}, x) \leftarrow_{\$} P^*(\cdot); jr \leftarrow_{\$} \mathbb{F}^{jl}$
2 $\pi \leftarrow_{\$} \text{Prove}(x, jr)$	6 $\pi \leftarrow_{\$} P^*(st_{P^*}, jr)$
3 $\sigma \leftarrow \text{Query}(x, \pi, jr; qr)$	7 $\sigma \leftarrow_{\$} \text{Query}(x, \pi, jl)$
4 $\text{ret } jr \parallel qr \parallel \sigma$	8 $\text{ret } x \notin \mathcal{L} \wedge \text{Decide}(\sigma)$

Figure 2: Procedures for defining security of FLPs.

have adapted the syntax so that it describes explicitly the computations of the prover and verifier. Namely, our query-generation algorithm takes in the input and proof and outputs the verifier string consumed by the decision algorithm, whereas in BBCG+19, the query-generation algorithm outputs a description of the linear function used to compute the verifier string.

Our notion of FLP soundness differs slightly from BBCG+19 in that it explicitly requires the prover to “commit” to the invalid prior to the joint randomness being generated. This clarifies that the joint randomness needs to be independent of the input in order for soundness to be achievable.

Incremental Distributed Point Functions. A point function is a function that is zero everywhere except on a special input α ; an incremental point function is a function that is zero everywhere except on a special input α and all of its prefixes. Intuitively, one can imagine arranging the co-domain of this function into a complete, binary tree in which the nodes are labeled with prefixes; and for each node labeled p , its children are labeled with $p \parallel 0$ and $p \parallel 1$. Each node on the path to the leaf node α is assigned the value 1, and all other nodes are assigned the value 0. (See [18, Figure 4] for an illustration.)

An incremental point function that gives output $\vec{\beta}[\ell]$ on the length- ℓ prefix of α is defined formally as:

$$f_{\alpha, \vec{\beta}}(pfx \in \{0, 1\}^{\leq \eta}) = \begin{cases} \vec{\beta}[|pfx|] & \text{if } pfx \text{ is a prefix of } \alpha \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

An *Incremental Distributed Point Function (IDPF)* [18] is a concise secret sharing of an incremental point function. We recall the definition of an IDPF from Boneh et al. [18] and restrict it slightly to suit the constructions of [9]. An IDPF’s domain is the set of bitstrings of length at most η . For each input length ℓ , the IDPF generates outputs in the group \mathbb{G}_ℓ . We present definitions only for the case of 2 parties, since leading constructions are specialized for that case. Let η , and κ be positive integers, let \mathcal{M} be a set, and let \mathbb{G}_ℓ be a group for each $\ell \in [\eta]$. An IDPF is a pair of algorithms:

- $\text{IDPF.Gen}(\alpha \in \{0, 1\}^\eta, \vec{\beta} \in \mathbb{G}_1 \times \dots \times \mathbb{G}_\eta) \rightarrow (\{0, 1\}^\kappa)^2 \times \mathcal{M}$ is the **key generation** algorithm that takes a bitstring α and a vector $\vec{\beta}$ of point values, each of which is an element of the group \mathbb{G}_ℓ for the corresponding input length. It outputs a pair of key shares and a “public share” (an element of \mathcal{M}).
- $\text{IDPF.Eval}(id \in \{1, 2\}, key \in \{0, 1\}^\kappa, pub \in \mathcal{M}, pfx \in \{0, 1\}^\ell) \rightarrow \mathbb{G}_\ell$ is the **point-function evaluation** algorithm that takes in a shareholder index, an IDPF key share, a public share pub , and a prefix string of $\ell \leq \eta$ bits, then outputs a share of the IDPF output.

An IDPF is *correct* if for all $\alpha \in \{0, 1\}^\eta$, all $\vec{\beta} \in \mathbb{G}_1 \times \dots \times \mathbb{G}_\eta$, all $(key_1, key_2, pub) \in [\text{IDPF.Gen}(\alpha, \vec{\beta})]$, and all strings pfx of length $\ell \leq \eta$:

$$f_{\alpha, \vec{\beta}}(pfx) = \sum_{j \in \{1, 2\}} \text{IDPF.Eval}(j, key_j, pub, pfx).$$

We define *privacy* for an IDPF later in Section 5.1.

3 Security Model

3.1 Syntax

As discussed in Section 1, a VDAF can be thought of as a protocol for evaluating an aggregation function F that takes as input the vector of measurements generated by clients and outputs an aggregate result. In addition, the function may include an auxiliary “aggregation parameter” that allows the measurements to be “refined” to contain only the information of interest to the collector. Accordingly, prior to executing the VDAF, each aggregator’s state is initialized with this aggregation parameter.

Recall that execution of a VDAF proceeds in four distinct phases. (See Figure 1 for an illustration.) We formalize the computation of the parties in each phase as the component algorithms of a VDAF:

- $\text{Shard}(m \in \mathcal{I}, n \in \mathcal{N}) \rightarrow (msg_{\text{Init}} \in \mathcal{M}, \vec{x} \in \mathcal{X}^s)$ is the randomized **sharding** algorithm run by the client. It takes in the client’s input measurement m and a nonce n and returns an **initial message**² to be broadcasted to all aggregators and a sequence of **input shares**, one for each of the s aggregators.
- $\text{Prep}(\hat{j} \in [s], sk \in \mathcal{SK}, st \in \mathcal{Q}, n \in \mathcal{N}, m\vec{sg} \in \mathcal{M}^*, x \in \mathcal{X}) \rightarrow (sts \in \{\text{running}, \text{finished}, \text{failed}\}, out \in (\mathcal{Q} \times \mathcal{M}) \cup \mathcal{Y} \cup \{\perp\})$ is the deterministic, interactive **preparation** algorithm run by each aggregator during the online preparation process. Its inputs are the share index \hat{j} , the **verification key** shared by the aggregators sk , the current state st , the nonce ($n \in \mathcal{N}$), the most recent round of **broadcast messages** $m\vec{sg}$ (or $(msg_{\text{Init}},)$ if this is the first round), and the aggregator’s input share x . The preparation algorithm returns an indication sts of whether the process is **running**, **finished**, or **failed**. When the status is **running**, the output includes the aggregator’s next state and broadcast message $((st, msg) \in \mathcal{Q} \times \mathcal{M})$; and when the status is **finished**, the output includes an aggregator’s **refined share** ($y \in \mathcal{Y}$).
- $\text{Agg}(\vec{y} \in \mathcal{Y}^*) \rightarrow a \in \mathcal{A}$ is the deterministic *aggregation* algorithm run locally by each aggregator. It takes in a sequence of refined shares \vec{y} and outputs an **aggregate share** a .
- $\text{Unshard}(ct \in \mathbb{N}, \vec{a} \in \mathcal{A}^s) \rightarrow r \in \mathcal{O}$ is the deterministic *unsharding* algorithm used to compute the aggregate result r . Its inputs are the report count ct and aggregate shares \vec{a} .

The sets \mathcal{I} , \mathcal{N} , \mathcal{M} , \mathcal{X} , \mathcal{SK} , \mathcal{Q} , \mathcal{Y} , \mathcal{A} , and \mathcal{O} must also be defined by the VDAF. (We typically do so only implicitly.) In addition to these sets, the VDAF specifies a set $\mathcal{Q}_{\text{Init}} \subseteq \mathcal{Q}$ of possible **initial states**.

Our definitions of robustness and privacy for VDAFs (defined in the next section) require three additional syntactic properties. The first is a property we call **refinement consistency**. Intuitively, this property insists that, for a given initial state, the VDAF defines the set of refined measurements with respect to which the validity of the refined shares is to be verified. For example, in the VDAF of Section 5, the set of measurements are fixed-length bitstrings, while the refined measurements are one-hot vectors over a finite field. Formally, refinement consistency requires the existence of functions `refine` and `refineFromShares` such that for all m, n and $st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$,

$$\Pr[\text{refine}(st_{\text{Init}}, m) = \text{refineFromShares}(st_{\text{Init}}, msg, \vec{x}) : (msg, \vec{x}) \leftarrow_s \text{Shard}(m, n)] = 1.$$

Second, we require **aggregation consistency**, which means, roughly, that aggregating refined shares into aggregate shares, then unsharding, is equivalent to first unsharding the refined shares, then aggregating. To illustrate this idea, imagine arranging the refined shares into a matrix, where the rows correspond to aggregators and the columns to measurements: Aggregation consistency means that one can either add up the columns, then the rows, or add up the rows, then the columns. Formally, we require the existence of a function `finishResult` such that for all refined shares $y_1^1, \dots, y_{ct}^1, \dots, y_1^s, \dots, y_{ct}^s \in \mathcal{Y}$, it holds that

$$\begin{aligned} \text{Unshard}(ct, (\text{Agg}(y_1^1, \dots, y_{ct}^1), \dots, \text{Agg}(y_1^s, \dots, y_{ct}^s))) = \\ \text{finishResult}(ct, \text{Unshard}(1, (\text{Agg}(y_1^1), \dots, \text{Agg}(y_1^s)), \\ \dots, \text{Unshard}(1, (\text{Agg}(y_{ct}^1), \dots, \text{Agg}(y_{ct}^s))))). \end{aligned}$$

²This message is called the “public share” in the specification.

We will see that these notions of refinement and aggregation consistency, while fairly technical in nature, are trivial to show for natural constructions (including Prio3 and Doplar).

Lastly, our privacy definition allows the VDAF to be executed multiple times over the same batch of measurements, each time beginning with a new initial state. (This accounts for the iterative nature of IDPFs.) Depending on the VDAF, it may be necessary for aggregators to restrict the sequence of initial states to prevent trivial leakage. Accordingly, we require each VDAF to specify an **allowed-state** algorithm `validSt` that takes in the sequence of previous initial states and the next initial state and returns a bit indicating whether the next initial state is allowed.

3.2 Security

Three definitions are given for VDAFs. The first, completeness, is used to specify correct evaluation of an aggregation function. The others, robustness and privacy, roughly correspond to notions of the same names from [23, Appendix A].

Security considerations for DAP [29]. Recall from introduction that the DAP standard is designed to securely execute a VDAF in a real world network. Aspects of our security model can be thought of abstracting away the functionality provided by DAP. As such, many of our modeling decisions here amount to requirements that the DAP protocol must fulfill. We will highlight some of these considerations throughout this section.

Completeness. We require that, when executed honestly, the VDAF evaluates its aggregation function F correctly. We formalize non-adversarial execution of Π via procedure `Run Π` in Figure 3. Along with the VDAF Π , this procedure is parameterized by an initial state st_{Init} with which to configure the aggregators and a sequence of measurements and nonces to process into an aggregate result.

Algorithm `Run` processes the measurements as illustrated in Figure 1. First, each measurement is sharded into input shares by the submitting client (line 4), then refined into a set of refined shares by the aggregators (5–16). Next, the refined shares recovered by each aggregator are combined into an aggregate share (18). Finally, the aggregate shares are combined by the collector into the aggregate result (19).

Definition 1 (Completeness). *Let $F : \mathcal{Q}_{\text{Init}} \times \mathcal{I}^* \rightarrow \mathcal{O}$ be a function. We say that VDAF Π is complete for F if for all $\vec{m} \in \mathcal{I}^*$ and $\vec{n} \in \mathcal{N}^*$ for which $|\vec{m}| = |\vec{n}|$ and $st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$ it holds that*

$$\Pr [\text{Run}_{\Pi}(st_{\text{Init}}, \vec{m}, \vec{n}) = F(st_{\text{Init}}, \vec{m})] = 1,$$

where the probability is over the randomness of `Run` and its subroutines. We say that Π is **complete** if it is complete for some function F .

Robustness. We say that VDAF Π is robust if, when all of the aggregators execute the protocol correctly, all “valid” refined measurements will be aggregated, while any “invalid” measurement will be filtered out by the aggregators (with high probability). This property is captured via the game $\text{Exp}_{\Pi}^{\text{robust}}(A)$ defined in Figure 3. In this game the adversary, acting as a coalition of malicious clients, submits reports to the aggregators, eavesdrops on their communication, and observes the result of their computation. This functionality is modeled by the `Prep` oracle, which the adversary may query any number of times. It controls the nonce and initial state for each trial, but its oracle queries are subject to the restriction that, for each distinct nonce, the sequence of initial states must be valid (according to the allowed-state algorithm `validSt`).

Validity is defined in terms of the refinement-consistency algorithms (see Section 3.1). Let $\mathcal{V}_{st_{\text{Init}}} = \{\text{refine}_{st_{\text{Init}}}(m) : m \in \mathcal{I}\}$ be the set of refined measurements for initial state st_{Init} . The adversary wins the robustness game if, when run on initial state st_{Init} , initial message msg_{Init} , and input shares \vec{x} , either: (1) an aggregator accepts a share of an invalid refined measurement, i.e., one of the aggregators ends in state `finished`, but $y = \text{refineFromShares}(st_{\text{Init}}, msg_{\text{Init}}, \vec{x}) \notin \mathcal{V}_{st_{\text{Init}}}$ (see line 14 in Figure 3); or (2) the refined shares computed by the aggregators do not match the expected refined measurement, i.e., unsharding the refined shares does not result in y (line 16).

Definition 2 (Robustness). *Define the advantage of A in defeating the robustness of VDAF Π as*

$$\text{Adv}_{\Pi}^{\text{robust}}(A) = \Pr [\text{Exp}_{\Pi}^{\text{robust}}(A)].$$

*Informally, we say that Π is **robust** if for every efficient adversary A , the value of $\text{Adv}_{\Pi}^{\text{robust}}(A)$ is small.*

<p>Algorithm $\text{Run}_\Pi(st_{\text{Init}}, \vec{m}, \vec{n})$:</p> <pre> 1 $sk \leftarrow_s \mathcal{SK}$; $ct \leftarrow \vec{m}$ 2 // Shard/Prepare 3 for $\hat{k} \in [ct]$: 4 $(msg, \vec{x}) \leftarrow \Pi.\text{Shard}(\vec{m}[\hat{k}], \vec{n}[\hat{k}])$ 5 $\text{Msg}[0, 1] \leftarrow msg$ 6 for $\hat{j} \in [s]$: $\text{St}[\hat{j}] \leftarrow st_{\text{Init}}$ 7 for $\hat{\ell} \in [r + 1]$: 8 for $\hat{j} \in [s]$: 9 $(sts, out) \leftarrow \Pi.\text{Prep}(\hat{j}, sk, \text{St}[\hat{j}],$ 10 $\vec{n}[\hat{k}], \text{Msg}[\hat{\ell}-1, \cdot], \vec{x}[\hat{j}])$ 11 if $sts = \text{running}$: 12 $(\text{St}[\hat{j}], msg) \leftarrow out$ 13 $\text{Msg}[\hat{\ell}, \hat{j}] \leftarrow msg$ 14 else if $sts = \text{finished}$: 15 $\text{Out}[\hat{j}, \hat{k}] \leftarrow out$ 16 else if $sts = \text{failed}$: ret \perp 17 // Aggregate/Unshard 18 for $\hat{j} \in [s]$: $\vec{a}[\hat{j}] \leftarrow \Pi.\text{Agg}(\text{Out}[\hat{j}, \cdot])$ 19 ret $\Pi.\text{Unshard}(ct, \vec{a})$ </pre>	<p>Game $\text{Exp}_\Pi^{\text{robust}}(A)$:</p> <pre> 1 $sk \leftarrow_s \mathcal{SK}$; $w \leftarrow \text{false}$; $A^{\text{Prep}}(\cdot)$; ret w Prep($n \in \mathcal{N}, \vec{x} \in \mathcal{X}^s, msg_{\text{Init}} \in \mathcal{M}, st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$): 2 if not $\Pi.\text{validSt}(\text{Used}[n], st_{\text{Init}})$: ret \perp 3 $\text{Used}[n] \leftarrow \text{Used}[n] \parallel (st_{\text{Init}},)$ 4 $\text{Msg}[0, 1] \leftarrow msg_{\text{Init}}$ 5 $y \leftarrow \Pi.\text{refineFromShares}(st_{\text{Init}}, msg_{\text{Init}}, \vec{x})$ 6 for $\hat{j} \in [s]$: $\text{St}[\hat{j}] \leftarrow st_{\text{Init}}$ 7 for $\hat{\ell} \in [r + 1]$: 8 for $\hat{j} \in [s]$: 9 $(sts, out) \leftarrow \Pi.\text{Prep}(\hat{j}, sk, \text{St}[\hat{j}],$ 10 $n, \text{Msg}[\hat{\ell}-1, \cdot], \vec{x}[\hat{j}])$ 11 if $sts = \text{running}$: 12 $(\text{St}[\hat{j}], msg) \leftarrow out$ 13 $\text{Msg}[\hat{\ell}, \hat{j}] \leftarrow msg$ 14 else if $sts = \text{finished}$: $y_{\hat{j}} \leftarrow out$; $\tilde{w} \leftarrow [y \notin \mathcal{V}_{st_{\text{Init}}}]$ 15 else if $sts = \text{failed}$: pass 16 if not \tilde{w}: $\tilde{w} \leftarrow [y \neq \Pi.\text{Unshard}(1, \Pi.\text{Agg}(y_1), \dots, \Pi.\text{Agg}(y_s))]$ 17 $w \leftarrow w \vee \tilde{w}$; ret ($w, \text{Msg}$) </pre>
---	---

Figure 3: Left: Procedure for defining completeness of r -round, s -party VDAF Π . Right: Game for defining robustness of Π . Let $\mathcal{Q}_{\text{Init}} \subseteq \mathcal{Q}$ denote the set of valid initial states and, for each $st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$, let $\mathcal{V}_{st_{\text{Init}}} = \{\text{refine}_{st_{\text{Init}}}(m) : m \in \mathcal{I}\}$.

Remark 1. *If a VDAF is both robust and complete for some aggregation function F , then as long as the aggregators run the protocol correctly, the collector is guaranteed to correctly compute F over the measurements uploaded by honest clients. This is due to aggregation consistency requirement for VDAFs defined in Section 3.1.*

Privacy. We formalize privacy via the indistinguishability game $\text{Exp}_{\Pi,t}^{\text{priv}}(A)$ in the right panel of Figure 4. The game is associated with VDAF Π , adversary A , and **corruption threshold** t . We consider an attacker that controls the collector and statically corrupts at most t aggregators (lines 1–2). Using its **Prep** oracle (lines 16–28), the adversary controls transmission of all messages in the protocol, *except* for the honestly generated input shares sent to honest (uncorrupted) aggregators. We assume that the adversary also controls setup (see the **Setup** oracle on lines 11–15), meaning that it can pick the verification keys for honest aggregators (1) and the initial state of each run of the preparation phase (14). This captures the real-world setting of the DAP protocol [29], where one of the aggregators (the “leader”) effectively picks these values on behalf of the others (the “helpers”). Note that our game requires the secret key to be committed to prior to generating measurements: This is a deliberate restriction that was necessary to prove security of our constructions. (It is necessary for DAP to enforce this restriction.)

The initial state for each run is subject to the restriction imposed by the allowed-state algorithm defined by the VDAF (lines 11–13). (Accordingly, it is necessary for honest aggregators to enforce this restriction in the DAP protocol.)

The game asks A to distinguish execution of the protocol on two sets of measurements of its choosing. To capture this, the attacker is given an oracle **Shard** (lines 6–10) that models execution of the honest clients. This oracle takes in two measurements m_0, m_1 and shards m_b , where b is the challenge bit chosen at the start of the game, and returns the initial message and the input shares of the corrupted aggregators. The oracle chooses a nonce n from the nonce space \mathcal{N} at random. (Accordingly, the DAP protocol must arrange for clients to choose their nonces at random.)

To model an attacker that controls the collector, the game allows the adversary to learn the aggregate shares computed by honest aggregators. This is captured by the **Agg** oracle (lines 29–35). Queries to this oracle are subject to the restriction that the aggregate share does not trivially leak challenge bit:

<p>Game $\text{Exp}_{II,t}^{\text{priv}}(A)$:</p> <ol style="list-style-type: none"> 1 $(st_A, \mathcal{V}, (sk_j)_{j \in \mathcal{V}}) \leftarrow_s A()$ 2 if $\mathcal{V} + t \neq s$ return \perp 3 $b \leftarrow_s \{0, 1\}$ 4 $b' \leftarrow_s A^{\text{Shard, Setup, Prep, Agg}}(st_A)$ 5 ret $b = b'$ <p>Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):</p> <ol style="list-style-type: none"> 6 if $\text{Used}[\hat{k}] \neq \perp$: ret \perp 7 $n \leftarrow_s \mathcal{N}$ 8 $(\text{Pub}[\hat{k}], \text{In}[\hat{k}, \cdot]) \leftarrow_s II.\text{Shard}(m_b, n)$ 9 $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 10 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}})$ <p>Setup($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$):</p> <ol style="list-style-type: none"> 11 if $\text{Status}[\hat{i}, \hat{j}] \neq \perp$ 12 or not $II.\text{validSt}(\text{Setup}[\cdot, \hat{j}], st_{\text{Init}})$: 13 ret \perp 14 $\text{Setup}[\hat{i}, \hat{j}] \leftarrow st_{\text{Init}}$ 15 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{running}$ 	<p>Prep($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, \hat{k} \in \mathbb{N}, m\bar{s}g \in \mathcal{M}^*$):</p> <ol style="list-style-type: none"> 16 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$ or $\text{In}[\hat{k}, \hat{j}] = \perp$: ret \perp 17 if $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \perp$: 18 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \text{Setup}[\hat{i}, \hat{j}]; m\bar{s}g \leftarrow (\text{Pub}[\hat{k}],)$ 19 $(n, m_0, m_1) \leftarrow \text{Used}[\hat{k}]$ 20 $(sts, out) \leftarrow$ 21 $II.\text{Prep}(\hat{j}, sk_{\hat{j}}, \text{St}[\hat{i}, \hat{j}, \hat{k}], n, m\bar{s}g, \text{In}[\hat{k}, \hat{j}])$ 22 if $sts = \text{running}$: 23 $(st, msg) \leftarrow out; \text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow st$ 24 else if $sts = \text{finished}$: 25 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp; \text{Out}[\hat{i}, \hat{j}, \hat{k}] \leftarrow out$ 26 $\text{Batch}_0[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_0; \text{Batch}_1[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_1$ 27 else if $sts = \text{failed}$: $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp$ 28 ret (sts, msg) <p>Agg($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}$):</p> <ol style="list-style-type: none"> 29 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$: ret \perp 30 $(st_1, \dots, st_s) \leftarrow \text{Setup}[\hat{i}, \cdot]$ 31 if $F(st_{\hat{j}}, \text{Batch}_0[\hat{i}, \hat{j}, \cdot]) \neq F(st_{\hat{j}}, \text{Batch}_1[\hat{i}, \hat{j}, \cdot])$ 32 and $(\forall j, j' \in \mathcal{V}) st_j = st_{j'} \wedge sk_j = sk_{j'}$: 33 ret \perp 34 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{finished}$ 35 ret $II.\text{Agg}(\text{Out}[\hat{i}, \hat{j}, \cdot])$
---	--

Figure 4: Game for defining privacy of a complete, s -party VDAF II for corruption threshold ≥ 0 . Let F denote the aggregation function for which II is complete and let $\mathcal{Q}_{\text{Init}}$ its set of initial states. Let $\mathcal{T} = [s] \setminus \mathcal{V}$.

Namely, the aggregate of both batch of measurements specified by the adversary must be equal (31). (Tables $\text{Batch}_0, \text{Batch}_1$ keep track of the pairs of measurements m_0, m_1 passed to the **Shard** for which a given aggregator has recovered a refined share for a given initial state.) This restriction is analogous to the “leakage function” provided to the simulator in previous simulation-style definitions. See [23, Appendix A] and [18, Appendix A]. We consider something slightly stronger: If the honest aggregators disagree either on the initial state or the verification key, then we do not impose the restriction (32). This amounts to demanding that the aggregate shares leak nothing in this case.

Definition 3 (Privacy). *Let II be an s -party VDAF and let $t < s$ be a positive integer. Define the t -advantage of A in attacking the privacy of II as*

$$\text{Adv}_{II,t}^{\text{priv}}(A) = 2 \cdot \Pr[\text{Exp}_{II,t}^{\text{priv}}(A)] - 1.$$

Informally, we say that II is t -private if for every efficient A the value of $\text{Adv}_{II,t}^{\text{priv}}(A)$ is small.

4 Prio3

In this section we present our security analysis for Prio3 [9]. The starting point for this construction is an FLP system (Section 2) that defines the set of valid measurements; drawing on techniques from Boneh et al. [17], Prio3 exploits the full-linearity property to allow the aggregators to validate the secret shared input. However, in order for the resulting VDAF to be suitable for a particular aggregation function $F : \mathcal{I} \rightarrow \mathcal{O}$, we need the proof system to define how measurements (\mathcal{I}) are encoded as inputs to the prover and how refined shares are processed into the aggregate results (\mathcal{O}).

Definition 4 (Affine, aggregatable encodings [23, Sec. 5.]). *Let $F : \mathcal{I} \rightarrow \mathcal{O}$ be a function. An FLP system FLP admits an affine, aggregatable encoding for F if it defines the following algorithms:*

- $\text{FLP.Encode}(m \in \mathcal{I}) \rightarrow \text{inp} \in \mathbb{F}^n$ is an injective map from the domain of F to the input space \mathbb{F}^n of FLP.

<p>Algorithm Shard(m, n):</p> <pre> 1 $inp \leftarrow \text{Encode}(m)$ 2 for $\hat{j} \in [2..s]$: 3 $blind_{\hat{j}}, xseed_{\hat{j}}, pseed_{\hat{j}} \leftarrow_{\\$} \{0, 1\}^{\kappa}$ 4 $\vec{x}[\hat{j}] \leftarrow \text{RG}_2(xseed_{\hat{j}}, \hat{j})$ 5 $rseed[\hat{j}] \leftarrow \text{RG}_7(blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}])$ 6 $\vec{x}[1] \leftarrow inp - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ 7 $blind_1 \leftarrow_{\\$} \{0, 1\}^{\kappa}$ 8 $rseed[1] \leftarrow \text{RG}_7(blind_1, 1 \parallel n \parallel \vec{x}[1])$ 9 $jseed \leftarrow \text{RG}_6(0^{\kappa}, rseed)$; $jr \leftarrow \text{RG}_1(jseed, \varepsilon)$ 10 $ps \leftarrow_{\\$} \{0, 1\}^{\kappa}$; $pr \leftarrow \text{RG}_4(ps, \varepsilon)$ 11 $\vec{\pi}[1] \leftarrow \text{Prove}(inp, jr; pr)$ 12 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{\hat{j}=2}^s \text{RG}_3(pseed_{\hat{j}}, \hat{j})$ 13 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 14 for $\hat{j} \in [2..s]$: 15 $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 16 ret $(rseed, \vec{x})$ </pre> <p>Algorithm Unpack(\hat{j}, x):</p> <pre> 17 if $\hat{j} = 1$: $(inp, \pi, blind) \leftarrow x$ 18 else: 19 $(xseed, pseed, blind) \leftarrow x$ 20 $inp \leftarrow \text{RG}_2(xseed, \hat{j})$ 21 $\pi \leftarrow \text{RG}_3(pseed, \hat{j})$ 22 ret $(inp, \pi, blind)$ </pre>	<p>Algorithm Prep($\hat{j}, sk, st, n, m\vec{s}g, x$):</p> <pre> 23 if $st = \varepsilon$: // Process initial message from client 24 $(inp, \pi, blind) \leftarrow \text{Unpack}(\hat{j}, x)$ 25 $(rseed, \cdot) \leftarrow m\vec{s}g$; $rseed[\hat{j}] \leftarrow \text{RG}_7(blind, \hat{j} \parallel n \parallel inp)$ 26 $jseed \leftarrow \text{RG}_6(0^{\kappa}, rseed)$; $jr \leftarrow \text{RG}_1(jseed, \varepsilon)$ 27 $qr \leftarrow \text{RG}_5(sk, n)$ 28 $msg \leftarrow (\text{Query}(inp, \pi, jr; qr), rseed[\hat{j}])$ 29 $st \leftarrow (jseed, \text{Truncate}(inp))$ 30 ret $(\text{running}, st, msg)$ 31 // Process broadcast messages from aggregators 32 $(jseed, y) \leftarrow st$; $(\vec{v}\vec{f}s[\hat{j}], rseed[\hat{j}])_{\hat{j} \in [s]} \leftarrow m\vec{s}g$ 33 $acc \leftarrow \text{Decide}(\sum_{\hat{j}=1}^s \vec{v}\vec{f}s[\hat{j}])$ 34 if acc and $jseed = \text{RG}_6(0^{\kappa}, rseed)$: ret $(\text{finished}, y)$ 35 else ret (failed, \perp) </pre> <p>Algorithm Agg(\vec{y}):</p> <pre> 36 ret $\sum_{i=1}^{ \vec{y} } \vec{y}[i]$ </pre> <p>Algorithm Unshard(\vec{a}):</p> <pre> 37 ret $\text{Decode}(\sum_{i=1}^{ \vec{a} } \vec{a}[i])$ </pre> <p>Algorithm RG_i($seed, cntxt$):</p> <pre> 38 $l \leftarrow (jl, n, m, pl, ql)$ 39 if $i \leq 5$: ret $\text{Expand}[\text{PRG}](seed, \text{label}_i \parallel cntxt, \mathbb{F}.p, l[i])$ 40 else: ret $\text{PRG.Next}(\text{PRG.Init}(seed, \text{label}_i \parallel cntxt), \kappa)$ </pre>
---	--

Figure 5: Definition of 1-round, s -party VDAF Prio3[FLP, PRG]. Let $\text{label}_1, \dots, \text{label}_7$ be arbitrary, distinct bitstrings.

- $\text{FLP.Truncate}(inp \in \mathbb{F}^n) \rightarrow out \in \mathbb{F}^{ol}$ refines an FLP input into a format more suitable for aggregation. We call ol the output length.
- $\text{FLP.Decode}(ct \in \mathbb{N}, out \in \mathbb{F}^{ol}) \rightarrow a \in \mathcal{O}$ converts a refined, aggregated output out to its final form a . This computation may depend on the number of measurements ct .

Correctness requires that for all $ct \geq 0$ and $\vec{m} \in \mathcal{I}^{ct}$ it holds that

$$F(\vec{m}) = \text{Decode}\left(ct, \sum_{i \in [ct]} \text{Truncate}(\text{Encode}(\vec{m}[i]))\right).$$

Let FLP be an FLP system that admits an affine, aggregatable encoding for F and let PRG be a PRG. We specify the core algorithms of Prio3[FLP, PRG] in Figure 5. (This version includes some minor changes to idraft-irtf-cfrg-vdaf-03 [9], as we discuss below.) The sharding algorithm begins by encoding the measurement as prescribed by the FLP. It then splits the encoded measurement inp into shares, generates a proof of inp 's validity, and splits the proof into shares as well. The joint randomness jr passed to the proof generation algorithm is derived from the input shares following the Fiat-Shamir-style transform described—but not formally analyzed—in [17, Section 6.2.3]. During preparation, the aggregators collectively re-compute jr from their input shares. Each aggregator broadcasts a share of the verifier by running the FLP query-generation algorithm on its share of the input and proof. (The query randomness qr is derived from the shared verification key sk and the nonce n provided by the environment.) The FLP decision algorithm is used run on the combined verifier shares to decide validity.

The aggregators must derive the joint randomness prior to computing their verifier shares. In order to allow them to perform both computations in parallel in a single round, the client sends in its initial

message the sequence \vec{rseed} of “joint randomness hints” consisting of the intermediate values computed by the aggregators. This allows jr to be computed immediately on receipt of the input shares. However, since a malicious client might try to break robustness by transmitting malformed hints, the aggregators also verify the joint randomness was computed properly in the same flow.

Allowed initial states. The set of initial states for Prio3 is simply $\mathcal{Q}_{\text{Init}} = \{\varepsilon\}$. In our security analysis, we assume honest aggregators to process a batch at most once. Accordingly, the allowed-state algorithm $\text{Prio3}[\text{FLP}, \text{PRG}].\text{validSt}$ accepts only if the batch was not aggregated previously.

Consistency. The set of refined measurements includes any output of the affine, aggregatable encoding for FLP. That is, on input of $st_{\text{Init}} \in \{\varepsilon\}$ and $m \in \mathcal{I}$, the refinement algorithm $\text{Prio3}[\text{FLP}, \text{PRG}].\text{refine}$ first encodes m , then truncates and decodes it as prescribed by FLP. The refine-from-shares algorithm, $\text{Prio3}[\text{FLP}, \text{PRG}].\text{refineFromShares}$, unpacks each input share (see `Unpack` in Figure 5), extracts the shares of the FLP input, truncates them, adds them together, and decodes the result.

For aggregation consistency, we require the encoding scheme for FLP to be aggregation-consistent in a similar sense. Specifically, there must exist a function `finishResult` such that for all truncated outputs $out_1, \dots, out_{ct} \in \mathbb{F}^{ol}$ it holds that $\text{Decode}(ct, \sum_{k \in [ct]} out_k) = \text{finishResult}(ct, \text{Decode}(1, out_1), \dots, \text{Decode}(1, out_{ct}))$.

Changes to the specification [9]. Figure 5 differs from draft-03 of the VDAF spec in three ways. The most important change is to incorporate the nonce provided by the environment into the joint randomness computation. This turns out to be crucial for a tight robustness bound; without this change, we must contend with cases in which joint randomness is reused across reports.

Second, we have revised the domain separation tags for the PRG invocations so that each RG_i in Figure 5 can be treated as an independent random oracle.

Lastly, we have moved the joint randomness hints from the input shares into the client’s initial broadcast message. This change allowed us to simplify our proofs somewhat, but we do not believe it is essential for security. It also has the added benefit of reducing overall communication overhead for $s > 2$.

Security. Fix $s > 2$ and let $\Pi = \text{Prio3}[\text{FLP}, \text{PRG}]$ be as specified above. Let \mathcal{N} denote the nonce space for Π and let κ denote the seed length of PRG.

Theorem 1. *Modeling each RG_i in Figure 5 as a random oracle, if FLP is ϵ -sound (Section 2), then for every adversary A against the robustness of Π it holds that*

$$\text{Adv}_{\Pi}^{\text{robust}}(A) \leq (q_{\text{RG}} + q_{\text{Prep}}) \cdot \epsilon + \frac{q_{\text{RG}} + q_{\text{Prep}}^2}{2^{\kappa-1}},$$

where A makes q_{Prep} queries to `Prep` and a total of q_{RG} queries to its random oracles.

For reasonable choices of the seed size for PRG, the loosest term in this bound is $(q_{\text{RG}} + q_{\text{Prep}}) \cdot \epsilon$. The multiplicative loss of $q_{\text{RG}} + q_{\text{Prep}}$ reflects the adversary’s ability to partially control the randomness of the FLP insofar as it is able to use rejection sampling to obtain query and joint randomness with any property. The ϵ -soundness of FLP bounds the probability of violating soundness in a single interaction, but in a VDAF the attacker may interact with the underlying FLP once in each of its q_{Prep} queries to `Prep`, and it can use its queries to RG_1 to bias these interactions’ joint randomness.

Proof sketch. We sketch the security reduction here and defer the detailed proof to Appendix C.1. Our goal is to construct from A a malicious prover P^* for the soundness of FLP. The overall idea is to run A in a simulation of the robustness game for Π in which P^* ’s instance of the soundness experiment (Figure 2) is embedded in a random `Prep` query so that P^* wins its game precisely when A sets $w \leftarrow \text{true}$ for the first time in that query. The main difficulty is that P^* must arrange to use the joint randomness it received as input in its own game. To provide a consistent simulation of RG_1 , we need to arrange to extract the input to commit to from A ’s queries. This results in a union bound over all queries to RG_1 , either by the simulation of `Prep` or by A directly. \square

Remark 2. *For FLPs that do not make use of joint randomness (i.e., those for which $jl = 0$), queries to RG_1 can be disregarded, as this oracle is not used by Π . In particular, a similar reduction can be shown that results in a multiplicative loss of just q_{Prep} .*

Remark 3. *Although we have not addressed this explicitly in our specification, the extraction step of our security reduction relies on the encoding of the context string passed to each RG_i being invertible. (Similarly for Theorem 3.)*

Theorem 2. *Modeling each RG_i in Figure 5 as a random oracle, if FLP is δ -private, then for all $0 < t < s$ and attackers A it holds that*

$$\text{Adv}_{\Pi,t}^{\text{priv}}(A) \leq 2q_{\text{Shard}} \left(\delta + \frac{q_{\text{RG}} + q_{\text{Shard}}}{|\mathcal{N}|} + \frac{s \cdot q_{\text{RG}}}{2^{\kappa-1}} \right),$$

where A makes q_{Shard} queries to Shard and a total of q_{RG} queries to the random oracles.

Proof sketch. The full proof is given in Section C.2. The main idea is to arrange for A 's queries to its oracles to be independent of the challenge bit. We do so via a game-playing argument in which we incrementally revise the game until the outcome of each oracle is independent of the current state of the game. The last step involves a hybrid argument, where in each hybrid world we replace one invocation of the proof- and query-generation algorithms of FLP (see Figure 2) with invocation of the simulator hypothesized by the δ -privacy of FLP. This accounts for the multiplicative loss of q_{Shard} in the bound. \square

Remark 4. *Instead of using separate seeds for the input share, proof share, and blind, it may be safe to reuse the same seed for all three purposes, similar to the seed in Doplar (Section 5). This may only result in a slightly looser bound: This change would the attacker to test guesses of the input share because the joint randomness part would be derived from the same seed.*

5 Doplar

Poplar1 [9] solves the “heavy hitters” problem (as described in Section 1) using an IDPF (Section 2) in the following way. Two aggregators hold shares of an IDPF key generated by the clients. Each evaluates its IDPF key at a number of equal-length candidate prefixes. They expect that the output is non-zero for at most one of these candidates; to verify this, they execute an MPC to determine if they hold shares of a one-hot vector, and that the non-zero value is in the desired range (i.e., equal to one or zero). If verification succeeds, then each adds its share of the vector together with the other verified shares. The aggregate result is a vector representing the number of measurements prefixed by each candidate.

The “secure sketch” MPC of Boneh et al. [18] requires two rounds of communication between the aggregators. (Computing and verifying this sketch occurs during the preparation phase of VDAF evaluation.) In this section we propose an alternative strategy that, leveraging techniques in Section 4, requires just one.

Our first step is to factor the validity check into two, parallelizable computations. The first computation is solely responsible for checking that the vector of IDPF outputs is one-hot. In Section 5.1 we extend IDPFs (Section 2) into *verifiable* IDPFs (VIDPFs), which preserve the same privacy properties as IDPFs, but additionally verify the one-hotness of the refined shares. In Appendix A we show how to instantiate this primitives using a simple technique from de Castro and Polychroniadou [22].

The second computation checks that the *sum* of the elements of the vector is in the desired range. Our first idea is to perform this range check using an FLP (Section 2). This does not work, however, since a standard FLP requires the prover to know the statement it is proving: In our case, it does not know the value of the sum computed by the aggregators, since it does not know the candidate prefixes. To overcome this, we show in Section 5.2 how to transform an FLP into one that is *delayed input* [42]. Such a proof system allows a proof to be generated for a *set* of potential inputs such that the honest verifier accepts the proof for any input in this set, but rejects otherwise (with high probability). We define delayed-input FLP in Section 5.2 and defer the construction to Appendix B.

The result is the 1-round, 2-party VDAF presented in Section 5.3. The cost of this round reduction is a modest increase in overall communication, at least for the current instantiations of the VIDPF and delayed-input FLP. We compare the communication cost of Doplar and Poplar1 at the end of this section.

5.1 Verifiable IDPF

A **verifiable** IDPF (VIDPF) allows the dealer to prove to the shareholders that their shares represent a one-hot vector. For our purposes, we define a **one-hot vector** as a vector that is nonzero in *at most* one component (i.e., the all-zeroes vector is also one-hot). Verifiable function secret sharings (of which VIDPF is a special case) were previously considered in [19, 22], and a construction specifically for VIDPF was given in [22].

A VIDPF has two algorithms in addition to the usual `Gen`, `Eval`:

- `VIDPF.VEval`($id \in \{1, 2\}, key \in \{0, 1\}^\kappa, pub \in \mathcal{M}, \vec{x} \in (\{0, 1\}^\ell)^u$) $\rightarrow \{0, 1\}^* \times (\mathbb{G}_\ell)^u$ takes as input an IDPF share (private and public parts), and a sequence of IDPF inputs. It outputs a verification value and a sequence of output shares.
- `VIDPF.Verify`(h_1, h_2) $\rightarrow \{0, 1\}$ takes as input two verification values and returns a boolean.

We also overload the syntax of the plaintext evaluation function to take a vector of inputs, i.e., we let

$$f_{\alpha, \vec{\beta}}(\vec{x}) = \left(f_{\alpha, \vec{\beta}}(\vec{x}[1]), f_{\alpha, \vec{\beta}}(\vec{x}[2]), \dots \right).$$

We say VIDPF is *correct* if, for all $\alpha \in \{0, 1\}^\eta$, all $\vec{\beta} \in \mathbb{G}_1 \times \dots \times \mathbb{G}_\eta$, all $\vec{x} \in (\{0, 1\}^\ell)^*$, all $(key_1, key_2, pub) \in [\text{Gen}(\alpha, \vec{\beta})]$, all $(h_1, \vec{y}_1) \in [\text{VEval}(1, key_1, pub, \vec{x})]$, and all $(h_2, \vec{y}_2) \in [\text{VEval}(2, key_2, pub, \vec{x})]$:

- $\vec{y}_1 + \vec{y}_2 = f_{\alpha, \vec{\beta}}(\vec{x})$
- If $(\vec{y}_1 + \vec{y}_2)$ is a one-hot vector then $\mathcal{V}.\text{Verify}(h_1, h_2) = 1$

Theorem 3 requires VIDPF to be *extractable*. Intuitively, there should be an algorithm that can extract $\alpha, \vec{\beta}$ from adversarially generated VIDPF key shares. Then `VEval` must produce shares consistent with the incremental point function $f_{\alpha, \vec{\beta}}$, whenever `Verify` succeeds. (A similar property is formalized for IDPFs by BBCG+21.) This property implies, among other things, that if `Verify` succeeds, then shareholders are guaranteed to hold shares of a one-hot vector. We formalize this property below.

Definition 5 (Extractable VIDPF (cf. [18, Definition 7])). *Suppose that VIDPF is defined in terms of a random oracle with co-domain \mathcal{Y} . Refer to the game in Figure 6 associated to VIDPF, **extractor** E , and adversary A . Define A 's advantage in **fooling** E as $\text{Adv}_{\text{VIDPF}, E}^{\text{extract}}(A) = 2 \cdot \Pr[\text{Exp}_{\text{VIDPF}, E}^{\text{extract}}(A)] - 1$.*

Finally, our privacy reduction for Doplar (Theorem 4) requires the underlying VIDPF to be *private*, in the sense that one shareholder's view—consisting of its share key_j , the public share pub , and the other shareholder's verification value h —leaks nothing about the secrets α and β . Prior definitions of verifiable FSS—e.g., the one in de Castro and Polychroniadou [22]—only define privacy with respect to a single vector of evaluation points and single verification predicate, both of which are assumed to be known at the time of share generation. In our setting, shares are generated and only later is there a choice of evaluation points and verification predicates. The same shares may be evaluated many times, on different input vectors and with different verification predicates. This leads to a more interactive, and stronger, definition for privacy than in prior works.³

Definition 6. *Let $\text{Exp}_{\text{VIDPF}, S}^{\text{priv}}(A)$ be the privacy game for VIDPF, **simulator** $S = (S_1, S_2)$, and adversary A defined in Figure 6. Define the advantage of A in distinguishing S 's simulation from its view of VIDPF's execution as $\text{Adv}_{\text{VIDPF}, S}^{\text{priv}}(A) = 2 \cdot \Pr[\text{Exp}_{\text{VIDPF}, S}^{\text{priv}}(A)] - 1$.*

If this privacy game withholds the `Sketch` oracle from the adversary (shaded in Figure 6) then we obtain the privacy game for plain IDPFs, with the adversary's advantage defined analogously.

In Section A we formally describe a VIDPF construction that satisfies all the necessary security properties. The construction is heavily based on the verifiable DPF technique from [22].

³The game does not need to provide an oracle for `VIDPF.Verify` since it is a deterministic algorithm whose inputs are known to the adversary.

<p>Game $\text{Exp}_{\text{VIDPF},E}^{\text{extract}}(A)$:</p> <ol style="list-style-type: none"> 1 $b \leftarrow \{0, 1\}$; $(key_1, key_2, pub, st_A) \leftarrow A^{\text{RO}}()$ 2 if $b = 0$: $(\alpha, \vec{\beta}) \leftarrow E(key_1, key_2, pub, \text{Rand})$ 3 $b' \leftarrow A^{\text{RO, Eval}}(st_A)$; ret $b = b'$ <p><u>Eval</u>(\vec{x}):</p> <ol style="list-style-type: none"> 4 $(h_1, \vec{y}_1) \leftarrow \text{VIDPF.VEval}^{\text{RO}}(1, key_1, pub, \vec{x})$ 5 $(h_2, \vec{y}_2) \leftarrow \text{VIDPF.VEval}^{\text{RO}}(2, key_2, pub, \vec{x})$ 6 if $b = 0$ and $\text{VIDPF.Verify}^{\text{RO}}(h_1, h_2) = 1$: ret $f_{\alpha, \vec{\beta}}(\vec{x})$ 7 else: ret $\vec{y}_1 + \vec{y}_2$ <p><u>RO</u>(inp):</p> <ol style="list-style-type: none"> 8 if $\text{Rand}[inp] = \perp$: $\text{Rand}[inp] \leftarrow \mathcal{Y}$ 9 ret $\text{Rand}[inp]$ <hr/> <p>Game $\text{Exp}_{\text{VIDPF},S}^{\text{priv}}$($A$):</p> <ol style="list-style-type: none"> 10 $b \leftarrow \{0, 1\}$; $(st_A, \alpha, \vec{\beta}, \hat{j}) \leftarrow A()$ 11 if $b = 0$: $(key_j, pub) \leftarrow S_1(\hat{j})$ 12 else: $(key_1, key_2, pub) \leftarrow \text{VIDPF.Gen}(\alpha, \vec{\beta})$ 13 $b' \leftarrow A^{\text{Sketch}}(st_A, key_j, pub)$; ret $b = b'$ <p><u>Sketch</u>(\vec{x}):</p> <ol style="list-style-type: none"> 14 if $b = 0$: $h \leftarrow S_2(\hat{j}, key_j, pub, \vec{x})$ 15 else: $(h, \cdot) \leftarrow \text{VIDPF.VEval}(3 - \hat{j}, key_{3-\hat{j}}, pub, \vec{x})$ 16 ret h 	<p>Game $\text{Exp}_{\text{DFLP},S}^{\text{priv}}$($A$):</p> <ol style="list-style-type: none"> 1 $b \leftarrow \{0, 1\}$; $(\mathcal{X}, st_A) \leftarrow A()$ 2 if $b = 0$: $(st_S, jr, qr) \leftarrow S_1(\mathcal{X})$ 3 else: 4 $jr \leftarrow \mathbb{F}^{j^l}$; $qr \leftarrow \mathbb{F}^{q^l}$; $\Delta \leftarrow \mathbb{F}^{el}$ 5 $\pi \leftarrow \text{DFLP.Prove}(\mathcal{X}, \Delta, jr)$ 6 $(x, st_A) \leftarrow A(st_A, jr, qr)$; assert $x \in \mathcal{X}$ 7 if $b = 0$: $\sigma \leftarrow S(st_S)$ 8 else: $\sigma \leftarrow \text{DFLP.Query}(\text{DFLP.Encode}(\Delta, x), \Delta, \pi, jr; qr)$ 9 $b' \leftarrow A(st_A, \sigma)$; ret $b = b'$
--	---

Figure 6: Games for defining extractability (top-right), and privacy (bottom-left) VIDPFs and privacy of delayed-input FLP (right).

5.2 Delayed-Input FLPs

We introduce a new variant of fully linear proofs (FLPs), in which the prover does not know in advance which instance (i.e., input) will be used during verification. Instead, the proof is generated only knowing a set of possible instances; later, the proof is verified using one of those instances. For technical reasons, the proof and verification steps operate not on the instance, but on a *randomized encoding* of the instance. This extra randomness is useful in our eventual construction (Appendix B).

We adopt the terminology of **delayed-input**, which is standard in the study of (interactive) zero-knowledge protocols. In an interactive protocol with delayed input, the instance and witness need not be known/chosen until some intermediate round (often the prover’s final round). In our setting, the actual choice of instance/witness is not chosen until after the prover finishes “speaking”. The protocol of Lapidot and Shamir [42] is often regarded as the first ZK protocol with delayed input, while Katz and Ostrovsky [40] were the first to explicitly rely on the delayed input property while using a ZK proof in an application.

Definition 7. A **delayed-input FLP** DFLP consists of the following algorithms:

- $\text{DFLP.Encode}(\Delta \in \mathbb{F}^{el}, x \in \mathbb{F}^n) \rightarrow e \in \mathbb{F}^{n'}$: Takes as input encoding randomness Δ , and an input instance x . Returns an encoding of x ; we let n' denote the length of the encoding. The function $\text{Encode}(\Delta, \cdot)$ must be a linear function and invertible. We denote the inverse by Decode .
- $\text{DFLP.Prove}(\mathcal{X} \subseteq \mathbb{F}^n, \Delta \in \mathbb{F}^{el}, jr \in \mathbb{F}^{j^l}) \rightarrow \pi \in \mathbb{F}^m$: Takes as input a **set** of possible instances, encoding randomness Δ , and joint randomness jr . Produces output proof π .
- $\text{DFLP.Query}(e \in \mathbb{F}^{n'}, \Delta \in \mathbb{F}^{el}, \pi \in \mathbb{F}^m, jr \in \mathbb{F}^{j^l}; qr \in \mathbb{F}^{q^l}) \rightarrow \sigma \in \mathbb{F}^v$: Takes as input an encoded instance e , encoding randomness Δ , proof π , joint randomness jr , and query randomness qr . Returns a verifier σ . The function $\text{Query}(\cdot, \cdot, \cdot, jr; qr)$ must be linear.

<p>Algorithm Shard(α, n):</p> <pre> 1 // Construct the VIDPF key shares. 2 $seed_1, seed_2 \leftarrow \mathbb{S} \{0, 1\}^\kappa$ 3 for $\ell \in [\eta]$: 4 $\vec{\Delta}[\ell] \leftarrow \text{RG}_2(seed_1, n \parallel \ell \parallel 1)$ 5 $+ \text{RG}_2(seed_2, n \parallel \ell \parallel 2)$ 6 $\vec{\beta}[\ell] \leftarrow \text{DFLP.Encode}(\vec{\Delta}[\ell], 1)$ 7 $(key_1, key_2, pub) \leftarrow \mathbb{S} \text{VIDPF.Gen}(\alpha, \vec{\beta})$ 8 // Prepare the joint randomness. 9 $rseed[1] \leftarrow \text{RG}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 10 $rseed[2] \leftarrow \text{RG}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 11 $jseed \leftarrow \text{RG}_6(0^\kappa, rseed)$ 12 // Generate the level proofs. 13 for $\ell \in [\eta]$: 14 $jr \leftarrow \text{RG}_1(jseed, n \parallel \ell)$ 15 $\pi \leftarrow \mathbb{S} \text{DFLP.Prove}(\{0, 1\}, \vec{\Delta}[\ell], jr)$ 16 $\vec{pf}[\ell] \leftarrow \pi - \text{RG}_3(seed_2, n \parallel \ell)$ 17 // Prepare the initial message and input shares. 18 $x_1 \leftarrow (key_1, seed_1, \vec{pf})$ 19 $x_2 \leftarrow (key_2, seed_2)$ 20 $msg \leftarrow (pub, rseed)$ 21 ret (msg, x_1, x_2) </pre> <p>Algorithm Unpack(\hat{j}, x, n, ℓ):</p> <pre> 22 if $\hat{j} = 1$: $(key, seed, \vec{pf}) \leftarrow x$; $\pi \leftarrow \vec{pf}[\ell]$ 23 else: $(key, seed) \leftarrow x$; $\pi \leftarrow \text{RG}_3(seed, n \parallel \ell)$ 24 ret $(key, seed, \pi)$ </pre>	<p>Algorithm Prep($\hat{j}, sk, st, n, msg, x$):</p> <pre> 25 if $st \in \mathcal{Q}_{\text{Init}}$: // Process initial message from client 26 $(\ell, \vec{pf}x) \leftarrow st$; $u \leftarrow \vec{pf}x$ 27 $(pub, rseed) \leftarrow msg$; $(key, seed, \pi) \leftarrow \text{Unpack}(\hat{j}, x, n, \ell)$ 28 $\Delta \leftarrow \text{RG}_2(seed, n \parallel \ell \parallel \hat{j})$ 29 $rseed[\hat{j}] \leftarrow \text{RG}_5(seed, n \parallel \ell \parallel \hat{j} \parallel pub \parallel key)$ 30 $jseed \leftarrow \text{RG}_6(0^\kappa, rseed)$ 31 $jr \leftarrow \text{RG}_1(jseed, n \parallel \ell)$; $qr \leftarrow \text{RG}_4(sk, n \parallel \ell)$ 32 $(h, \vec{y}) \leftarrow \text{VIDPF.VEval}(\hat{j}, pub, key, \vec{pf}x)$ 33 $inp \leftarrow \sum_{i \in [u]} \vec{y}[i]$ 34 $\sigma \leftarrow \text{DFLP.Query}(inp, \Delta, \pi, jr; qr)$ 35 $msg \leftarrow (\sigma, rseed[\hat{j}], h)$; $st \leftarrow (jseed, (\text{DFLP.Decode}(\vec{y}[i])_{i \in [u]}))$ 36 ret $(\text{running}, st, msg)$ 37 // Process broadcast messages from aggregators 38 $(jseed, \vec{y}) \leftarrow st$; $((\sigma_1, rseed_1, h_1), (\sigma_2, rseed_2, h_2)) \leftarrow msg$ 39 $acc \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 40 if acc and $jseed = \text{RG}_6(0^\kappa, (rseed_1, rseed_2))$ 41 and $\text{VIDPF.Verify}(h_1, h_2)$: ret $(\text{finished}, \vec{y})$ 42 else: ret (failed, \perp) </pre> <p>Algorithm Agg(\vec{y}): ret $\sum_{i=1}^{ \vec{y} } \vec{y}[i]$</p> <p>Algorithm Unshard(\vec{a}): ret $\sum_{i=1}^{ \vec{a} } \vec{a}[i]$</p> <p>Algorithm RG_i($seed, cntxt$):</p> <pre> 43 $l \leftarrow (jl, el, m, ql)$ 44 if $i \leq 4$: ret $\text{Expand}[\text{PRG}](seed, \text{label}_i \parallel cntxt, \mathbb{F}.p, l[i])$ 45 else: ret $\text{PRG.Next}(\text{PRG.Init}(seed, \text{label}_i \parallel cntxt), \kappa)$ </pre>
--	--

Figure 7: Definition of 1-round, 2-party VDAF Doplar[VIDPF, DFLP, PRG]. Let $\text{label}_1, \dots, \text{label}_6$ be arbitrary, distinct bitstrings.

- $\text{DFLP.Decide}(\sigma \in \mathbb{F}^v) \rightarrow acc \in \{0, 1\}$: Takes as input query responses σ and returns a boolean.

If Prove is restricted to sets \mathcal{X} with $|\mathcal{X}| = k$ then we call the construction a **delayed-k-input FLP**.

A delayed-input FLP should satisfy the following properties:

- **Completeness** (with respect to language \mathcal{L}): For all $\mathcal{X} \subseteq \mathcal{L}$, all $x \in \mathcal{X}$, and all Δ :

$$\Pr[\text{Decide}(\sigma) : jr \leftarrow \mathbb{S} \mathbb{F}^{jl}; \pi \leftarrow \mathbb{S} \text{Prove}(\mathcal{X}, \Delta, jr); \\ \sigma \leftarrow \mathbb{S} \text{Query}(\text{Encode}(\Delta, x), \Delta, \pi, jr)] = 1.$$

- **Soundness** (with respect to \mathcal{L}): The scheme should be sound in the usual sense of FLPs, with respect to the language $\mathcal{L}^* = \{(\text{Encode}(\Delta, x), \Delta) \mid x \in \mathcal{L}\}$. In other words, it is hard for a malicious prover to generate a proof that verifies with respect to $(e, \Delta) \notin \mathcal{L}^*$.
- **Privacy**: In Figure 6 we define a game for delayed-input FLPs, in which the proof is generated using some set \mathcal{X} of candidates, and later verified with respect to a particular $x \in \mathcal{X}$. A delayed-input FLP is δ -private if there exists a simulator S such that every A 's advantage is $\text{Adv}_{\text{DFLP}, S}^{\text{priv}}(A) \leq \delta$, where

$$\text{Adv}_{\text{DFLP}}^{\text{priv}}(A) = 2 \cdot \Pr[\text{Exp}_{\text{DFLP}, S}^{\text{priv}}(A)] - 1.$$

5.3 Construction

We specify our construction $\text{Doplar}[\text{VIDPF}, \text{DFLP}, \text{PRG}]$ in Figure 7. Its three components are: A verifiable IDPF VIDPF with input length η ; a delayed-2-input FLP DFLP with input set $\{0, 1\}$, proof length m , encoded input length n , encoding randomness length el , joint randomness length jl , and query randomness length ql ; and a pseudorandom generator PRG (Section 2) with seed length κ . To be suitable for our construction, we must choose VIDPF and DFLP so that $\text{VIDPF}.\mathbb{G}_\ell = \text{DFLP}.\mathbb{F}^n$ for each $\ell \in [\eta]$.

To shard its measurement $\alpha \in \{0, 1\}^\eta$, the client begins by running the VIDPF key generator on α . The initial state for Doplar encodes the “level” ℓ at which the VIDPF shares are to be evaluated; each candidate prefix must have length ℓ . (Recall from Section 2 that (V)IDPFs can be thought of as shares of values arranged in a binary tree with nodes labeled by prefixes.) Thus, for each level of the VIDPF tree, the client generates a delayed-input proof of the refined shares’ validity; just as for Prio3 (Section 4), the joint randomness used at each level is derived from the aggregator’s input shares. The VIDPF output is programmed so that the sum of the output shares corresponds to an encoded input for the delayed-input FLP.

To prepare a report for aggregation, the aggregators evaluate their VIDPF key shares at the desired candidate prefixes, then interact in order to check that (1) the joint randomness was computed correctly, (2) their refined shares are one-hot, and (3) the sum of their refined shares is either one or zero.

Allowed initial states. An initial state is valid if it consists of a sequence of candidate prefixes all having the same length. Moreover, each of the prefixes must be distinct. An initial state is allowed for $\text{Doplar}[\text{VIDPF}, \text{DFLP}, \text{PRG}]$ if the prefix length is distinct from all previous states for the same report. That is, the allowed-state algorithm validSt only permits a new state $st = (\ell, p\vec{x})$ if ℓ is distinct for all previous states and each of the prefixes $p\vec{x}$ is distinct.

Remark 5. *Although not addressed in Boneh et al. [18] explicitly, this restriction on the candidate prefixes is necessary for Poplar as well, as re-using the correlated randomness shared by the client would reveal information about the secret-shared vector.*

Consistency. The set of refined measurements for Doplar are one-hot vectors over the field \mathbb{F} for which the non-zero element is equal to 0 or 1. For a given initial state $(\ell, p\vec{x})$, this can be computed from the VIDPF public share and key shares by evaluating the shares on each of the prefixes $p\vec{x}$. Aggregation consistency for Doplar is similarly straight-forward: Since the refined share space and aggregate share space are the same.

Security. Let $\Pi = \text{Doplar}[\text{VIDPF}, \text{DFLP}, \text{PRG}]$ as specified above. Let \mathcal{N} be the nonce space and let κ be the seed length for PRG.

Theorem 3. *Modeling each RG_i in Figure 7 as a random oracle, if DFLP is ϵ -sound, then for all t_A -time adversaries A and t_E -time extractors E there exists a $O(t_A + q_{\text{Prep}}t_E)$ -time adversary B for which*

$$\begin{aligned} \text{Adv}_{\Pi}^{\text{robust}}(A) &\leq 2(q_{\text{RG}} + q_{\text{Prep}}) \cdot \epsilon + \frac{(q_{\text{RG}} + 3q_{\text{Prep}})^2}{2^\kappa} \\ &\quad + q_{\text{Prep}} \cdot \text{Adv}_{\text{VIDPF}, E}^{\text{extract}}(B), \end{aligned}$$

where A makes q_{Prep} queries to Prep and a total of q_{RG} queries to its random oracles.

Proof sketch. The proof has a similar structure to Theorem 1 in that the last step is a reduction to the soundness of DFLP. However in order to use this, we must first revise the game so that the challenge input issued by the malicious prover P^* was constructed from the sum of refined shares that are otherwise valid (i.e., one-hot). Using the extractability property of VIDPF, we can simplify the winning condition by extracting the the input measurement from the adversary’s random oracle queries and use it to compute the refined measurement whenever the one-hotness check succeeds. Refer to Appendix C.3 for the full proof. \square

Theorem 4. *For all t_A -time adversaries A and t' -time simulators S, T there exist $O(t_A + q_{\text{Shard}}t')$ -time adversaries B, C for which*

$$\begin{aligned} \text{Adv}_{\Pi, 1}^{\text{priv}}(A) &\leq 2q_{\text{Shard}} \left(\text{Adv}_{\text{VIDPF}, S}^{\text{priv}}(B) + \eta \cdot \text{Adv}_{\text{DFLP}, T}^{\text{priv}}(C) \right. \\ &\quad \left. + \frac{\eta q_{\text{RG}} + q_{\text{Shard}}}{|\mathcal{N}|} + \frac{3q_{\text{RG}}}{2^{\kappa-1}} \right), \end{aligned}$$

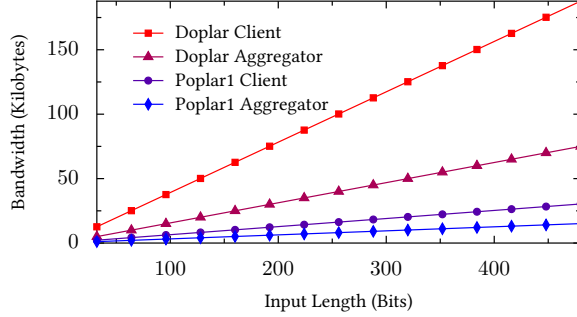


Figure 8: The total number of kilobytes sent by each party, per report, for various input lengths.

where each RG_i in Figure 7 is modeled as a random oracle, adversary A makes a total of q_{RG} queries to all of its random oracles and q_{Shard} queries to **Shard**.

Proof sketch. The reduction to DFLP privacy follows the same lines as Theorem 2 except there are $\eta \cdot q_{\text{Shard}}$ different hybrid worlds in the last step. Privacy of VIDPF is used to ensure that the simulation of the boundary world can be carried out without access to the input measurement. Refer to Appendix C.4 for the proof. \square

Comparison to Poplar1. To finish this section, we evaluate the bandwidth cost of our round-reduced variant of Poplar1. The parameters chosen for Poplar1 by the specification [9] match those in the performance evaluation conducted by Boneh et al. [17]. We therefore take these parameters as our basis for comparison. In the following, we have instantiated VIDPF and DFLP as described in Appendix A and Appendix B respectively.

Boneh et al. [17] claim a per-report robustness bound of roughly $2/|\mathbb{F}|$, where \mathbb{F} is the field chosen for the inner nodes.⁴ They choose a 62-bit field. In order to obtain the same robustness bound, while permitting the adversary at most 2^{64} queries to its random oracles, we need to use a 128-bit field for Doplar. For both constructions, we choose the PRG seed length to be $\kappa = 128$.

In Figure 8 we plot the communication cost of Doplar and Poplar1 for various choices of input lengths η . We plot the total number of kilobytes sent by each client and the total number of kilobytes sent by each aggregator, per report, over all η rounds of aggregation. As one would expect, the communication cost for Doplar scales linearly with the input length. However, the client’s bandwidth is about 6 times that of Poplar1; and the Aggregator’s bandwidth is about 5 times.

Level skipping. One way to improve bandwidth for both schemes is to “skip” IDPF evaluation at certain levels. For example, if we descend the IDPF tree in τ -bit increments instead of 1-bit increments, then (1) our VIDPF construction requires one-hot check material only in every τ -th level, and (2) the Doplar construction requires DFLPs only at every τ -th level.⁵ As a result, these major contributors to communication cost are reduced by a factor of τ . Additionally, the process of aggregating (traversing the tree of prefixes to find heavy hitters) requires fewer rounds by a factor of τ . The trade-off is that we consider more candidate prefixes at each level—i.e., at each step we consider the 2^τ descendants at depth τ from each candidate—but this cost is amortized over the batch.

Notably, the impact of this optimization is more significant for Doplar than for Poplar1. (For example, a “skip factor” of $\tau = 2$, i.e., skipping every other level, reduces the client’s overhead from 6 to 5 times that of Poplar1 with the same optimization.) This is primarily due to the reduction in the number of delayed-input proofs, which make up the bulk of the first input share. (The second input share compresses its shares of the proofs into a single PRG seed.)

⁴Poplar1 uses a smaller field for the inner nodes of the IDPF tree than the leaf nodes.

⁵The underlying (non-verifiable) IDPF is still organized as a binary tree, so its cost is not affected.

6 Conclusion and Future Work

The PPM working group’s ambition is to preserve user privacy even as software systems rely increasingly on gaining insights into user behavior. Our work aims to help ensure the standardization effort rests on firm formal foundations. However, it leaves open a number of directions for future work. We discuss two in the remainder.

Security analysis of DAP. The definitions in this paper apply to VDAFs, which are only a component of the DAP specification [29]. Thus, our work necessarily leaves open the security of the end-to-end protocol. There are two important questions. First, DAP is designed to inherit the security properties of VDAF. That is, we hope that whatever can be proven about the VDAF also holds when the VDAF is instantiated in the real-world environment in which DAP runs. One way to address this is to formulate the problem in terms of *indifferentiability* [48]: If DAP’s execution can be shown to be indistinguishable from the execution of the VDAF in the idealized environment described here, then any attack against DAP can be translated into an attack against the underlying VDAF.

The other important question is whether DAP meets its own security goals, which, depending on the application, might go beyond what can be achieved with a VDAF alone. Consider that whether MPC-style definitions like ours are enough for privacy depends intrinsically on the nature of the measurements being collected and how they are aggregated. It is one thing to ensure that we securely compute the aggregate; it is another to ensure that the aggregate itself does not leak “too much” information about the measurements. In particular, in many applications it will be useful to provide for differential privacy (DP) [27] in addition to secure computation. There are definitions of DP that extend to the multi-party setting [44, 51], and a number of works have considered MPC protocols for aggregate functionalities that also guarantee differential privacy of the outputs [50, 35, 10]. We hope to see future work extend this investigation to specific VDAFs.

Doplar improvements. For some applications, it would be useful for Doplar (or Poplar1) if the leaf output could be “weighted”, i.e., a number in range $\{a, \dots, b\}$ rather than $\{0, 1\}$. (Consider the ad-conversion use case from Section 1: It might be useful to know not only how many purchases were made per ad impression, but the total amount of money that was spent.) The delayed- k -input FLP paradigm may allow for this generalization, if schemes can be constructed for $k > 2$. (In this work, we only construct the delayed-2-input FLP needed for plain heavy hitters.)

There is also room for improvement in the communication cost of Doplar. Despite the round reduction, the higher bandwidth may be prohibitive for some applications. However, we are optimistic that the bandwidth can be improved. Future work should focus on the delayed-2-input FLP. The current instantiation (Appendix B), while simple, effectively doubles the proof size of the base FLP.

References

- [1] Privacy preserving measurement (2022), URL <https://datatracker.ietf.org/wg/ppm/about/>
- [2] Private Advertising Tecnology Community Group (2022), URL <https://www.w3.org/community/patcg/>
- [3] Abdalla, M., Haase, B., Hesse, J.: Security analysis of cpace. Cryptology ePrint Archive, Paper 2021/114 (2021), URL <https://eprint.iacr.org/2021/114>, <https://eprint.iacr.org/2021/114>
- [4] Abdalla, M., Haase, B., Hesse, J.: CPace, a balanced composable PAKE. Internet-Draft draft-irtf-cfrg-cpace-06, Internet Engineering Task Force (Jul 2022), URL <https://datatracker.ietf.org/doc/draft-irtf-cfrg-cpace/06/>, work in Progress
- [5] Addanki, S., Garbe, K., Jaffe, E., Ostrovsky, R., Polychroniadou, A.: Prio+: Privacy preserving aggregate statistics via boolean shares. Cryptology ePrint Archive, Report 2021/576 (2021), <https://ia.cr/2021/576>
- [6] Anderson, E., Chase, M., Durak, F.B., Ghosh, E., Laine, K., Weng, C.: Aggregate measurement via oblivious shuffling. Cryptology ePrint Archive, Report 2021/1490 (2021), <https://ia.cr/2021/1490>

- [7] Apple, Google: Exposure Notification Privacy-preserving Analytics (ENPA). White paper (2021), https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf
- [8] Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M.: Proof verification and the hardness of approximation problems. *J. ACM* **45**(3), 501–555 (may 1998), ISSN 0004-5411, doi:10.1145/278298.278306, URL <https://doi.org/10.1145/278298.278306>
- [9] Barnes, R., Patton, C., Schoppmann, P.: Verifiable Distributed Aggregation Functions. Internet-Draft draft-irtf-cfrg-vdaf-03, Internet Engineering Task Force (Aug 2022), URL <https://datatracker.ietf.org/doc/draft-irtf-cfrg-vdaf/03/>, work in Progress
- [10] Bell, J., Gascón, A., Ghazi, B., Kumar, R., Manurangsi, P., Raykova, M., Schoppmann, P.: Distributed, private, sparse histograms in the two-server model. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, pp. 307–321 (2022)
- [11] Bell, J., Gascón, A., Lepoint, T., Li, B., Meiklejohn, S., Raykova, M., Yun, C.: Acorn: Input validation for secure aggregation. Cryptology ePrint Archive (2022), URL <https://eprint.iacr.org/2022/1461>
- [12] Bell, J.H., Bonawitz, K.A., Gascón, A., Lepoint, T., Raykova, M.: Secure single-server aggregation with (poly) logarithmic overhead. In: Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security, pp. 1253–1269 (2020)
- [13] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security, pp. 62–73, CCS '93, ACM, New York, NY, USA (1993), ISBN 0-89791-629-8
- [14] Bellare, M., Rogaway, P.: The security of triple encryption and a framework for code-based game-playing proofs. In: Vaudenay, S. (ed.) Advances in Cryptology - EUROCRYPT 2006, pp. 409–426, Springer Berlin Heidelberg, Berlin, Heidelberg (2006), ISBN 978-3-540-34547-3
- [15] Bitansky, N., Chiesa, A., Ishai, Y., Paneth, O., Ostrovsky, R.: Succinct non-interactive arguments via linear interactive proofs. In: Sahai, A. (ed.) Theory of Cryptography, pp. 315–333, Springer Berlin Heidelberg, Berlin, Heidelberg (2013), ISBN 978-3-642-36594-2
- [16] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H.B., Patel, S., Ramage, D., Segal, A., Seth, K.: Practical secure aggregation for privacy-preserving machine learning. In: proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, pp. 1175–1191 (2017)
- [17] Boneh, D., Boyle, E., Corrigan-Gibbs, H., Gilboa, N., Ishai, Y.: Zero-knowledge proofs on secret-shared data via fully linear pcps. In: Boldyreva, A., Micciancio, D. (eds.) Advances in Cryptology – CRYPTO 2019, pp. 67–97, Springer International Publishing, Cham (2019), ISBN 978-3-030-26954-8
- [18] Boneh, D., Boyle, E., Corrigan-Gibbs, H., Gilboa, N., Ishai, Y.: Lightweight techniques for private heavy hitters. Cryptology ePrint Archive, Report 2021/017 (2021), <https://ia.cr/2021/017>
- [19] Boyle, E., Gilboa, N., Ishai, Y.: Function secret sharing: Improvements and extensions. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 1292–1303 (2016)
- [20] Brickell, J., Shmatikov, V.: Efficient anonymity-preserving data collection. In: Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 76–85 (2006)
- [21] Canetti, R.: Universally composable security. *J. ACM* **67**(5) (sep 2020), ISSN 0004-5411, doi:10.1145/3402457, URL <https://doi.org/10.1145/3402457>
- [22] de Castro, L., Polychroniadou, A.: Lightweight, maliciously secure verifiable function secret sharing. In: Dunkelman, O., Dziembowski, S. (eds.) Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part I, Lecture Notes in Computer Science, vol. 13275, pp. 150–179, Springer (2022), doi:10.1007/978-3-031-06944-4_6, URL https://doi.org/10.1007/978-3-031-06944-4_6

- [23] Corrigan-Gibbs, H., Boneh, D.: Prio: Private, robust, and scalable computation of aggregate statistics. In: 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), pp. 259–282, USENIX Association, Boston, MA (Mar 2017), ISBN 978-1-931971-37-9, URL <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/corrigan-gibbs>
- [24] Danezis, G., Fournet, C., Kohlweiss, M., Zanella-Béguelin, S.: Smart meter aggregation via secret-sharing. In: Proceedings of the first ACM workshop on Smart energy grid security, pp. 75–80 (2013)
- [25] Davidson, A., Snyder, P., Quirk, E., Genereux, J., Livshits, B., Haddadi, H.: Star: Secret sharing for private threshold aggregation reporting. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, pp. 697–710 (2022)
- [26] Duan, Y., Canny, J., Zhan, J.: {P4P}: Practical {Large-Scale}{Privacy-Preserving} distributed computation robust against malicious users. In: 19th USENIX Security Symposium (USENIX Security 10) (2010)
- [27] Dwork, C.: Differential privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) Automata, Languages and Programming, pp. 1–12, Springer Berlin Heidelberg, Berlin, Heidelberg (2006), ISBN 978-3-540-35908-1
- [28] Elahi, T., Danezis, G., Goldberg, I.: Privex: Private collection of traffic statistics for anonymous communication networks. In: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, pp. 1068–1079 (2014)
- [29] Geoghegan, T., Patton, C., Rescorla, E., Wood, C.A.: Distributed Aggregation Protocol for Privacy Preserving Measurement. Internet-Draft draft-ietf-ppm-dap-02, Internet Engineering Task Force (Sep 2022), URL <https://datatracker.ietf.org/doc/draft-ietf-ppm-dap/02/>, work in Progress
- [30] Gilboa, N., Ishai, Y.: Distributed point functions and their applications. In: Nguyen, P.Q., Oswald, E. (eds.) Advances in Cryptology – EUROCRYPT 2014, pp. 640–658, Springer Berlin Heidelberg, Berlin, Heidelberg (2014), ISBN 978-3-642-55220-5
- [31] Green, M., Ladd, W., Miers, I.: A protocol for privately reporting ad impressions at scale. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, p. 1591–1601, Association for Computing Machinery, New York, NY, USA (2016), ISBN 9781450341394, doi:10.1145/2976749.2978407, URL <https://doi.org/10.1145/2976749.2978407>
- [32] Guo, C., Katz, J., Wang, X., Yu, Y.: Efficient and secure multiparty computation from fixed-key block ciphers. In: 2020 IEEE Symposium on Security and Privacy (SP), pp. 825–841, IEEE (2020)
- [33] Guo, X., Yang, K., Wang, X., Zhang, W., Xie, X., Zhang, J., Liu, Z.: Half-tree: Halving the cost of tree expansion in cot and dpf. Cryptology ePrint Archive, Paper 2022/1431 (2022), URL <https://eprint.iacr.org/2022/1431>
- [34] Hohenberger, S., Myers, S., Pass, R., et al.: Anonize: A large-scale anonymous survey system. In: 2014 IEEE Symposium on Security and Privacy, pp. 375–389, IEEE (2014)
- [35] Humphries, T., Akhavan Mahdavi, R., Veitch, S., Kerschbaum, F.: Selective mpc: Distributed computation of differentially private key-value statistics. In: Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security, pp. 1459–1472 (2022)
- [36] III, J.J.P., Charles, D., Gilton, D., Jung, Y.H., Parsana, M., Anderson, E.: Masked lark: Masked learning, aggregation and reporting workflow (2021)
- [37] Ishai, Y., Kushilevitz, E., Ostrovsky, R.: Efficient arguments without short pcps. In: Twenty-Second Annual IEEE Conference on Computational Complexity (CCC’07), pp. 278–291 (2007), doi:10.1109/CCC.2007.10

- [38] Jangir, P., Koti, N., Kukkala, V.B., Patra, A., Gopal, B.R., Sangal, S.: Vogue: Faster computation of private heavy hitters. Cryptology ePrint Archive, Paper 2022/1561 (2022), URL <https://eprint.iacr.org/2022/1561>
- [39] Jawurek, M., Kerschbaum, F.: Fault-tolerant privacy-preserving statistics. In: International Symposium on Privacy Enhancing Technologies Symposium, pp. 221–238, Springer (2012)
- [40] Katz, J., Ostrovsky, R.: Round-optimal secure two-party computation. In: Annual International Cryptology Conference, pp. 335–354, Springer (2004)
- [41] Kursawe, K., Danezis, G., Kohlweiss, M.: Privacy-friendly aggregation for the smart-grid. In: International Symposium on Privacy Enhancing Technologies Symposium, pp. 175–191, Springer (2011)
- [42] Lapidot, D., Shamir, A.: Publicly verifiable non-interactive zero-knowledge proofs. In: Menezes, A., Vanstone, S.A. (eds.) Advances in Cryptology - CRYPTO '90, 10th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1990, Proceedings, Lecture Notes in Computer Science, vol. 537, pp. 353–365, Springer (1990), doi:10.1007/3-540-38424-3_26, URL https://doi.org/10.1007/3-540-38424-3_26
- [43] Melis, L., Danezis, G., De Cristofaro, E.: Efficient private statistics with succinct sketches. arXiv preprint arXiv:1508.06110 (2015)
- [44] Mironov, I., Pandey, O., Reingold, O., Vadhan, S.: Computational differential privacy. In: Annual International Cryptology Conference, pp. 126–142, Springer (2009)
- [45] Molteni, D.: Improving the WAF with machine learning. Cloudflare blog (2022), URL <https://blog.cloudflare.com/waf-ml/>
- [46] Mouris, D., Sarkar, P., Tsoutsos, N.G.: Plasma: Private, lightweight aggregated statistics against malicious adversaries with full security. Cryptology ePrint Archive, Paper 2023/080 (2023), URL <https://eprint.iacr.org/2023/080>, <https://eprint.iacr.org/2023/080>
- [47] Mozilla: Origin Telemetry (2022), URL <https://firefox-source-docs.mozilla.org/toolkit/components/telemetry/collection/origin.html>
- [48] Patton, C., Shrimpton, T.: Quantifying the security cost of migrating protocols to practice. In: Micciancio, D., Ristenpart, T. (eds.) Advances in Cryptology – CRYPTO 2020, pp. 94–124, Springer International Publishing, Cham (2020), ISBN 978-3-030-56784-2
- [49] Popa, R.A., Blumberg, A.J., Balakrishnan, H., Li, F.H.: Privacy and accountability for location-based aggregate statistics. In: Proceedings of the 18th ACM conference on Computer and communications security, pp. 653–666 (2011)
- [50] Roth, E., Noble, D., Falk, B.H., Haeberlen, A.: Honeycrisp: Large-scale differentially private aggregation without a trusted core. In: Proceedings of the 27th ACM Symposium on Operating Systems Principles, p. 196–210, SOSP '19, Association for Computing Machinery, New York, NY, USA (2019), ISBN 9781450368735, doi:10.1145/3341301.3359660, URL <https://doi.org/10.1145/3341301.3359660>
- [51] Schoppmann, P., Vogelsang, L., Gascón, A., Balle, B.: Secure and scalable document similarity on distributed databases: Differential privacy to the rescue. Proceedings on Privacy Enhancing Technologies **2**, 209–229 (2020)

A Instantiating VIDPF

In this section we present our proposed VIDPF construction.

De Castro-Polychroniadou technique. De Castro & Polychroniadou [22] (hereafter DP22) proposed the following simple and elegant technique to verify that a vector is one-hot. Consider a vector \vec{v} that is additively secret-shared $\vec{v} = \vec{v}_1 \oplus \vec{v}_2$. For simplicity, we describe the technique assuming that the sharing is with respect to XOR, since in that case the shares of zero are *identical strings*. The technique adapts readily to the more general case of additive shares over any group. Assume also that the parties have additive shares of a *binary* indicator vector $\vec{b} = \vec{b}_1 \oplus \vec{b}_2$, which is nonzero exactly in the same positions that \vec{v} is.

First, observe that the parties can easily verify whether they hold shares of an all-zeroes vector, since this happens if and only if their shares (as strings) are identical. They can simply exchange and compare hashes of their share-vectors (although see our remark below for a disclaimer about this idea). The technique of DP22 is to adjust a one-hot vector into an all-zeroes vector, with the help of the dealer.

Define

$$\text{adjust}(\vec{v}_i, \vec{b}_i, C) = \left(H(1, \vec{v}_i[1]) \oplus \vec{b}_i[1] \cdot C, \quad H(2, \vec{v}_i[2]) \oplus \vec{b}_i[2] \cdot C, \dots \right)$$

If \vec{v} and \vec{b} are nonzero in (only) position i^* , then set $C^* = H(i^*, \vec{v}_1[i^*]) \oplus H(i^*, \vec{v}_2[i^*])$. Now consider the result of both shareholders applying $\text{adjust}(\cdot, \cdot, C^*)$ to their shares:

- In positions $i \neq i^*$ where they share zero, we have $\vec{v}_1[i] = \vec{v}_2[i]$ and $\vec{b}_1[i] = \vec{b}_2[i]$. For these positions in the output of adjust , both parties will compute identical strings.
- In position i^* , the parties have $\vec{b}_1[i^*] \neq \vec{b}_2[i^*]$. By symmetry, suppose $\vec{b}_1[i^*] = 1$ and $\vec{b}_2[i^*] = 0$. Then the first party will compute

$$\begin{aligned} & H(i^*, \vec{v}_1[i^*]) \oplus C^* \\ &= H(i^*, \vec{v}_1[i^*]) \oplus (H(i^*, \vec{v}_1[i^*]) \oplus H(i^*, \vec{v}_2[i^*])) \\ &= H(i^*, \vec{v}_2[i^*]) \end{aligned}$$

and the second party will compute $H(i^*, \vec{v}_2[i^*])$ as well.

In all cases, both parties will compute the same output of adjust , which they can check for equality by exchanging and comparing hashes. Hence, the dealer will compute and the C^* value and include it in the parties' DPF keys. They can use C^* to perform their verification.

To see why the DP22 approach is sound, suppose the parties hold shares of a non-one-hot vector — i.e., it is nonzero at positions $i \neq i'$. Do both parties compute the same output of adjust ? This can only happen if C^* value somehow corrects both positions i and i' , and this happens only when

$$\begin{aligned} & H(i, \vec{v}_1[i]) \oplus H(i, \vec{v}_2[i]) = C^* = H(i', \vec{v}_1[i']) \oplus H(i', \vec{v}_2[i']) \\ \iff & H(i, \vec{v}_1[i]) \oplus H(i, \vec{v}_2[i]) \oplus H(i', \vec{v}_1[i']) \oplus H(i', \vec{v}_2[i']) = 0 \end{aligned}$$

The construction is therefore sound if it is hard to find “multi-collisions” of this form in H . In particular, if H is a random oracle with output length 4κ then an adversary making $q < 2^\kappa$ queries to H can find such a collision with probability bounded by $q^4/2^{4\kappa} \ll q/2^\kappa$.

Issues with DP22. We point out two issues with the technique of DP22, which we remedy in our subsequent presentation. First, the construction described in DP22 does not use the index i as a domain-separating first argument to H . If somehow $\vec{v}_1[i] = \vec{v}_1[i']$ and $\vec{v}_2[i] = \vec{v}_2[i']$, and the indices i, i' are not used as input to H , then a single C^* value $H(\vec{v}_1[i]) \oplus H(\vec{v}_2[i])$ would cause the parties to compute identical outputs of adjust in both positions i and i' . In this situation, it is possible that a non-one-hot vector can pass the verification, breaking soundness. Adding these domain separations, as we have in our presentation above, fixes the issue.

Second, party #1 holds its share \vec{v}_1 and the correction value $C^* = H(i^*, \vec{v}_1[i^*]) \oplus H(i^*, \vec{v}_2[i^*])$. Suppose this party has a guess for i^* and a guess for the nonzero value $v = \vec{v}_1[i^*] \oplus \vec{v}_2[i^*]$. Then she can verify this guess by checking whether $C^* = H(i^*, \vec{v}_1[i^*]) \oplus H(i^*, \vec{v}_1[i^*] \oplus v)$ — all values she knows. Hence, C^* exposes an offline dictionary attack on the secret values i^* and $\vec{v}[i^*]$. If $\vec{v}[i^*]$ is high entropy, then this is no vulnerability at all. But if $\vec{v}[i^*]$ is known to be a small value like 1 (as is the case in many applications), then this issue allows a corrupt shareholder to unilaterally learn i^* , violating privacy.⁶ We resolve this by simply ensuring

⁶The [22] always hashes a high-entropy PRG seed, so we do not claim a concrete attack against their construction. However, their privacy proof does not seem to use the fact that this argument to the hash is high entropy.

<pre> VIDPF.Gen($\alpha \in \{0, 1\}^\eta, \vec{\beta} \in \mathbb{G}_1 \times \dots \times \mathbb{G}_\eta$): 1 for $\ell \in [\eta]$: 2 $\vec{R}[\ell] \leftarrow_s \{0, 1\}^\kappa$ 3 $\vec{\beta}^*[\ell] \leftarrow (1, \vec{\beta}[\ell], \vec{R}[\ell])$ 4 $(key_1, key_2, pub) \leftarrow \text{IDPF.Gen}(\alpha, \vec{\beta}^*)$ 5 for $\ell \in [\eta]$: 6 $px \leftarrow \alpha[1 : \ell]$ 7 $(-, data_1, R_1) \leftarrow \text{IDPF.Eval}(1, key_1, pub, px)$ 8 $(-, data_2, R_2) \leftarrow \text{IDPF.Eval}(2, key_2, pub, px)$ 9 $\vec{C}[\ell] \leftarrow \text{RG}(px \parallel -data_1 \parallel -R_1)$ 10 $\oplus \text{RG}(px \parallel data_2 \parallel R_2)$ 11 $pub^* \leftarrow (pub, \vec{C})$ 12 ret (key_1, key_2, pub^*) </pre>	<pre> VIDPF.VEval(id, key, pub^*, \vec{x}): 12 $(pub, \vec{C}) \leftarrow pub^*$ 13 for $i \in [\vec{x}]$ 14 $\vec{y}[i] \leftarrow \text{IDPF.Eval}(id, key, pub, \vec{x}[i])$ 15 $(b, data[i], R) \leftarrow \vec{y}[i]$ 16 $h \leftarrow h \parallel \text{adjust}(id, key, pub^*, b, \vec{x}[i])$ 17 ret $(h, data)$ VIDPF.adjust(id, key, pub^*, b, x): // a helper procedure 18 $(pub, \vec{C}) \leftarrow pub^*$ 19 if $x = 0$: ret x // length of x as a bit string 20 $(-, d, R) \leftarrow \text{IDPF.Eval}(id, key, pub, b, x)$ 21 $prefix \leftarrow \text{adjust}(id, key, pub^*, b, x[1 : x - 1])$ 22 ret $prefix \parallel (\text{RG}(x \parallel (-1)^{id} d \parallel (-1)^{id} R) \oplus b \cdot \vec{C}[x])$ VIDPF.Verify(h_1, h_2): 23 ret $h_1 == h_2$ </pre>
---	---

Figure 9: VIDPF construction VIDPF[IDPF], based on any IDPF. If the VIDPF is to be instantiated with groups $\mathbb{G}_1, \dots, \mathbb{G}_\eta$ then the underlying IPDF is instantiated with groups $\tilde{\mathbb{G}}_1, \dots, \tilde{\mathbb{G}}_\eta$, where $\tilde{\mathbb{G}}_\ell = \{0, 1\} \times \mathbb{G}_\ell \times \{0, 1\}^\kappa$.

that the dealer encodes a random element at the one-hot position (in addition to a potentially low-entropy desired value).⁷

Extending to incremental DPF. The technique of DP22 is well-suited for DPFs. In an incremental DPF (IDPF), we can apply their technique to each prefix-length. However, this guarantees only that each prefix-length corresponds to some point function. It does not necessarily guarantee that the point functions of the different prefix-lengths satisfy the prefix condition that is needed in an IDPF.

In our construction, we extend the DP22 technique to IDPFs. For each evaluation of the IDPF — say, at point x — we compute the adjustment strings using the DP22 technique, for x and all of its prefixes. This alone is not enough to guarantee the prefix property. To “tie different prefix lengths together,” we ask the shareholders to compute the adjustment strings with respect to the *same sharings of the indicator bit*, for all the prefixes of x . We show that this forces the point functions at every prefix-length to be prefix-consistent.

Immediate Optimizations in an Implementation. Our construction evaluates the underlying IDPF on all prefixes of the given strings. Doing this naively would increase the computational costs by a factor of ℓ when evaluating on strings of length ℓ . However, these extra evaluations are essentially free in existing IDPFs — while evaluating at string x , these constructions already evaluate all prefixes of x along the way. A reasonable implementation of our VIDPF will take advantage of this fact.

The verification value h produced by VEval is a very long string, consisting of $\ell \cdot 4\kappa$ bits for each query point of length ℓ . If parties are to exchange these h values in an application of our VIDPF, it would account for a significant fraction of the total communication. However, the Verify algorithm that uses these h values merely checks them for equality. Therefore, it suffices for each party to send only a collision-resistant hash of their h value, which can have fixed length only 2κ . This optimization changes the concrete security bound for VIDPF soundness, by adding a term for the probability of finding a collision under the hash function.

Lemma 1. *Let IDPF be an IDPF and RG be a random oracle with outputs of length 4κ . Let A be an adversary making q queries to RG. There is a $O(t_A)$ -time adversary A' such that the construction VIDPF[IDPF] in*

⁷We have chosen to describe our VIDPF to use an underlying IDPF as a black-box. When this is the case, we must ensure that the IDPF outputs have sufficient entropy for the one-hotness check. If we were to instead to analyze our VIDPF (instantiated with a natural IDPF construction) as a *monolithic construction*, it is likely that the underlying IDPF would already have internal entropy available that could be used for the one-hotness check. I.e., we may be able to obtain smaller share sizes by exploiting internal properties of the underlying IDPF.

Figure 9 satisfies the following:

$$\begin{aligned}\text{Adv}_{\text{VIDPF}[\text{IDPF}],E}^{\text{extract}}(A) &\leq (q^4 + q^2)/2^{4\kappa} \\ \text{Adv}_{\text{VIDPF}[\text{IDPF}]}^{\text{priv}}(A) &\leq \text{Adv}_{\text{IDPF}}^{\text{priv}}(A') + q/2^\kappa\end{aligned}$$

Proof. Correctness of our construction follows from the discussion above, and is the same as in DP22.

Extractability: We begin with a few observations, which hold for all VIDPF keys, even adversarially generated ones:

Observation: If $\text{adjust}(1, \text{key}_1, \text{pub}^*, b_1, x) = \text{adjust}(2, \text{key}_2, \text{pub}^*, b_2, x)$, then $\text{adjust}(1, \text{key}_1, \text{pub}^*, b_1, x') = \text{adjust}(2, \text{key}_2, \text{pub}^*, b_2, x')$ as well, for every prefix x' of x . This follows trivially by inspection and the recursive nature of adjust . Note that the same b_1, b_2 are used for both x and x' .

Observation: Let $\text{pub}^* = (\text{pub}, \vec{C})$. Suppose $\text{adjust}(1, \text{key}_1, \text{pub}^*, b_1, x) = \text{adjust}(2, \text{key}_2, \text{pub}^*, b_2, x)$, and $\text{IDPF.Eval}(1, \text{key}_1, \text{pub}, x) = (-, y_1, R_1)$, and $\text{IDPF.Eval}(2, \text{key}_2, \text{pub}, x) = (-, y_2, R_2)$. Then:

1. If $b_1 = b_2$ then $\text{RG}(x \parallel -y_1 \parallel -R_1) = \text{RG}(x \parallel y_2 \parallel R_2)$. This includes the case where $(y_1, R_1) + (y_2, R_2) = (0, 0)$, making the two calls to RG identical. It also includes the case where these two calls to RG are a collision.
2. If $b_1 \neq b_2$ then $\vec{C}[|x|] = \text{RG}(x \parallel -y_1 \parallel -R_1) \oplus \text{RG}(x \parallel y_2 \parallel R_2)$.

This observation can be verified by inspection.

Let \mathcal{E}_1 denote the bad event that the adversary queries RG and observes a collision. If RG has outputs of length 4κ , and the adversary makes q oracle queries, then the probability of this bad event is bounded by $q^2/2^{4\kappa}$. When \mathcal{E}_1 does *not* happen, then in condition (1) above, only the case that $(y_1, R_1) + (y_2, R_2) = (0, 0)$ is possible.

Let \mathcal{E}_2 denote the bad event that the adversary makes any four queries to RG that satisfy:

$$\begin{aligned}\text{RG}(pfx \parallel -d_1 \parallel -R_1) \oplus \text{RG}(pfx \parallel d_2 \parallel R_2) = \\ \text{RG}(pfx' \parallel -d'_1 \parallel -R'_1) \oplus \text{RG}(pfx' \parallel d'_2 \parallel R'_2)\end{aligned}$$

for $pfx \neq pfx'$ and $d_1 + d_2 \neq 0$ and $d'_1 + d'_2 \neq 0$. (These conditions ensure that the four calls to RG must be on distinct inputs.) If RG has outputs of length 4κ , and the adversary makes q oracle queries, then the probability of this bad event is bounded by $q^4/2^{4\kappa}$. When \mathcal{E}_2 does *not* happen, then any value $C \in \{0, 1\}^{4\kappa}$ uniquely determines *at most one* pair of queries satisfying $C = \text{RG}(pfx \parallel -d_1 \parallel -R_1) \oplus \text{RG}(pfx \parallel d_2 \parallel R_2)$

We can apply the two observations inductively and obtain the following. If $\text{adjust}(1, \text{key}_1, \text{pub}^*, b_1, x) = \text{adjust}(2, \text{key}_2, \text{pub}^*, b_2, x)$ for $b_1 \neq b_2$, then every correction word $\vec{C}[\ell]$ must be of the form $\text{RG}(x[1 : \ell] \parallel \dots) \oplus \text{RG}(x[1 : \ell] \parallel \dots)$, for $\ell \leq |x|$. Then, provided that \mathcal{E}_2 does not happen, there is at most one x of length ℓ for which \vec{C} can be written in this way.

Combining all of these observations, we can define the extractor as shown in Figure 10.

Conditioned on the event that E doesn't abort (which happens only with probability $(q^4 + q^2)/2^{4\kappa}$), we claim that the adversary has no advantage in the extractability game.

Consider a query to $\text{Eval}(\vec{x})$ in the game, and assume the call to Verify succeeds. Then for every $\vec{x}[i]$, the corresponding calls to adjust produce identical output. If these calls to adjust have $b_1 = b_2$, and \mathcal{E}_1 has not happened, then the corresponding output $\vec{y}[i]$ must be 0. If these calls to adjust have $b_1 \neq b_2$, then $\vec{C}[\ell]$ must have the form $\text{RG}(\vec{x}[i] \parallel \dots) \oplus \text{RG}(\vec{x}[i] \parallel \dots)$. If \mathcal{E}_2 has not happened, then $\vec{x}[i]$ is in fact unique with this property, and therefore $\vec{x}[i]$ is a prefix of α computed by the extractor E . One can easily check that E extracts $\vec{\beta}[\ell]$ that is equal to the VIDPF output $\vec{y}[i]$. In other words, \vec{y} matches the output of $f_{\alpha, \vec{\beta}}$. Hence, the adversary's advantage is zero.

Privacy: Let S^{IDPF} be the simulator for privacy for the underlying IDPF. The simulator for our construction is given in Figure 11. We prove privacy in a series of hybrids, also illustrated in Figure 11. Game $\underline{\mathbf{G}}_0$ refers to the original experiment $\text{Exp}_{\text{VIDPF}}^{\text{priv}}$, where we have inlined the definition of S_2 for convenience. The $b = 0$ and $b = 1$ branches of the Sketch oracle differ only in whose shares are given as input to VIDPF.Eval . By the correctness of the scheme, the distinction doesn't matter, so the Sketch oracle is independent of b . Eliminating the conditional in the Sketch oracle, we obtain $\underline{\mathbf{G}}_1$, which is distributed identically to $\underline{\mathbf{G}}_0$.

```

 $E(key_1, key_2, pub^*, \text{Rand}):$ 
1  $(pub, \vec{C}) \leftarrow pub^*$ 
2 if  $\mathcal{E}_1$  or  $\mathcal{E}_2$ : // defined in the text, here with respect to oracle queries listed in Rand
3   abort
4  $\alpha \leftarrow$  empty string
5 for  $\ell \in [\eta]$ :
6   if  $\exists a \in \{0, 1\}, y_1, R_1, y_2, R_2$  such that
7      $\vec{C}[\ell] = \text{Rand}[(\alpha \| a) \| -y_1 \| -R_1] \oplus \text{Rand}[(\alpha \| a) \| y_1 \| R_2]$ 
8      $\alpha \leftarrow \alpha \| a$ 
9      $\vec{\beta}[\ell] \leftarrow y_1 + y_2$ 
10  else:  $\alpha \leftarrow \alpha \| 0$ ;  $\vec{\beta}[\ell] \leftarrow 0$ 
11 ret  $(\alpha, \vec{\beta})$ 

```

Figure 10: Extractor for the proof of Lemma 1.

$\underline{\mathcal{G}}_2$ is identical to $\underline{\mathcal{G}}_1$, but we have inlined the definition of VIDPF.Gen for convenience. By the correctness of the underlying IDPF, outputs of $\text{IDPF.Eval}(1, \cdot)$ and $\text{IDPF.Eval}(2, \cdot)$ are secret-shares of the appropriate plaintext values. So it has no effect on the adversary’s view to solve for the output of $\text{IDPF.Eval}(3 - \hat{j}, \cdot)$ using the plaintext values and the output of $\text{IDPF.Eval}(\hat{j}, \cdot)$, instead of using $key_{3-\hat{j}}$. In doing so, we obtain $\underline{\mathcal{G}}_3$ which is distributed identically to $\underline{\mathcal{G}}_2$.

Now notice that in $\underline{\mathcal{G}}_3$, the value $key_{3-\hat{j}}$ is never used. As such, we can replace line 26 (the call to IDPF.Gen) with a corresponding call to the simulator S^{IDPF} , which generates a simulated $key_{\hat{j}}$ and pub . Call the result $\underline{\mathcal{G}}_4$ (not pictured); this advantage in distinguishing $\underline{\mathcal{G}}_3$ from $\underline{\mathcal{G}}_4$ is at most ϵ_{priv} .

In $\underline{\mathcal{G}}_4$, the random values $\vec{R}[\ell]$ are used only to solve for $R_{3-\hat{j}}$, which is in turn used only as an argument to RG . Define a bad event that the adversary ever queries RG at an input of this form — i.e., of the form $\text{RG}(\cdot \| \cdot \| \vec{R}[\ell] - R_{\hat{j}})$. The probability of the bad event is bounded by $q/2^\kappa$ since $\vec{R}[\ell]$ is uniformly random. Conditioned on this bad event not happening, the results of these queries to RG are freshly random, and the value that is assigned to $\vec{C}[\ell]$ is uniform. In that case, the behavior of the game is independent of the challenge bit because S_1 also assigns uniform values to $\vec{C}[\ell]$. The advantage in guessing the challenge bit is therefore bounded by the probability of the bad event. \square

B Instantiating Delayed-Input FLP

Our main result is to construct a delayed-2-input FLP for use in Doplar.

Lemma 2. *The construction in Figure 12 (when suitably instantiated) is a delayed-2-input FLP with perfect completeness, soundness $4(n+2)/(|\mathbb{F}| - n - 2)$, and privacy $1/|\mathbb{F}|$, for \mathbb{F} -arithmetic circuits with n multiplication gates.*

The main idea of the construction is simple. The prover wishes to generate a proof that will work with either of two instances x_1 and x_2 . She simply generates a separate FLP proof for both instances x_1 and x_2 , and randomly permutes the two proofs. To verify the combined proof against some x , the verifier accepts iff either of the component proofs verifies against that x .

Completeness and soundness of this construction are relatively clear. However, the construction is not necessarily zero-knowledge. While verifying the combined proof, we expect to verify a component proof against a *proof that was generated for some other instance* — e.g., verify a proof generated for x_1 against x_2 . The standard zero-knowledge property of the underlying FLP does not apply to this situation. Indeed, since the Query function is linear, the result of querying a “mismatched” instance-proof pair will reveal “how far away” the instance is from the correct one.

We show that, when the underlying FLP is that of Boneh et al. [17], and extra randomness is introduced into the statement by means of the $\text{Encode}(\Delta, \cdot)$ function, even the queries to the “mismatched”

<pre> <u>$S_1(\hat{j})$:</u> 1 $(key, pub) \leftarrow S^{\text{IDPF}}()$ 2 for $\ell \in [\eta]$: $\vec{C}[\ell] \leftarrow_{\\$} \{0, 1\}^{4\kappa}$ 3 $pub^* \leftarrow (pub, \vec{C})$ 4 ret (key, pub^*) <u>$S_2(\hat{j}, key, pub, \vec{x})$:</u> 5 $(h, -) \leftarrow \text{VIDPF.VEval}(\hat{j}, key, pub, \vec{x})$ 6 ret h </pre> <hr/> <pre> Game $\underline{\mathbb{G}}_0$ $\underline{\mathbb{G}}_1$: 7 $b \leftarrow_{\\$} \{0, 1\}$ 8 $(st_A, \alpha, \vec{\beta}, \hat{j}) \leftarrow A()$ 9 if $b = 0$: $(key_j, pub^*) \leftarrow_{\\$} S_1(\hat{j})$ 10 else: $(key_1, key_2, pub^*) \leftarrow_{\\$} \text{VIDPF.Gen}(\alpha, \vec{\beta})$ 11 $b^* \leftarrow A^{\text{Sketch}}(st_A, key_j, pub^*)$ 12 ret $b = b^*$ <u>Sketch(\vec{x}):</u> 13 if $b = 0$: 14 // $h \leftarrow S_2(\hat{j}, key_j, pub^*, \vec{x})$ 15 $(h, -) \leftarrow \text{VIDPF.VEval}(\hat{j}, key_j, pub^*, \vec{x})$ 16 else: $(h, -) \leftarrow \text{VIDPF.VEval}(3 - \hat{j}, key_{3-\hat{j}}, pub^*, \vec{x})$ 17 ret h </pre>	<pre> Game $\underline{\mathbb{G}}_2$ $\underline{\mathbb{G}}_3$: 18 $b \leftarrow_{\\$} \{0, 1\}$ 19 $(st_A, \alpha, \vec{\beta}, \hat{j}) \leftarrow A()$ 20 if $b = 0$: $(key_j, pub^*) \leftarrow_{\\$} S_1(\hat{j})$ 21 else: 22 // $(key_1, key_2, pub^*) \leftarrow_{\\$} \text{VIDPF.Gen}(\alpha, \vec{\beta})$: 23 for $\ell \in [\eta]$: 24 $\vec{R}[\ell] \leftarrow_{\\$} \{0, 1\}^{\kappa}$ 25 $\vec{\beta}^*[\ell] \leftarrow (1, \vec{\beta}[\ell], \vec{R}[\ell])$ 26 $(key_1, key_2, pub) \leftarrow \text{IDPF.Gen}(\alpha, \vec{\beta}^*)$ 27 for $\ell \in [\eta]$: 28 $px \leftarrow \alpha[1 : \ell]$ 29 $(b_1, data_1, R_1) \leftarrow \text{IDPF.Eval}(1, key_1, pub, px)$ 30 $(b_2, data_2, R_2) \leftarrow \text{IDPF.Eval}(2, key_2, pub, px)$ 31 $(b_j, data_j, R_j) \leftarrow \text{IDPF.Eval}(\hat{j}, key_j, pub, px)$ 32 $(b_{3-\hat{j}}, data_{3-\hat{j}}, R_{3-\hat{j}})$ 33 $\leftarrow (1 \oplus b_j, \vec{\beta}[\ell] - data_j, \vec{R}[\ell] - R_j)$ 34 $\vec{C}[\ell] \leftarrow \text{RG}(px \parallel -data_1 \parallel R_1)$ 35 $\oplus \text{RG}(px \parallel data_2 \parallel R_2)$ 36 $pub^* \leftarrow (pub, \vec{C})$ 37 $b^* \leftarrow A^{\text{Sketch}}(st_A, key_j, pub^*)$ 38 ret $b = b^*$ <u>Sketch(\vec{x}):</u> 37 // $h \leftarrow S_2(\hat{j}, key_j, pub^*, \vec{x})$ 38 $(h, -) \leftarrow \text{VIDPF.VEval}(\hat{j}, key_j, pub^*, \vec{x})$ 39 ret h </pre>
--	--

Figure 11: Simulator and hybrids used in the proof of privacy for the VIDPF construction.

instance+proof can be simulated. Intuitively, the extra uncertainty of Δ blinds the results of the problematic queries.

Proof of Lemma 2. Figure 12 describes a delayed-input FLP that uses a basic FLP as a building block. Our claims in this proof rely on that FLP being instantiated using the construction of [17], also used in the VDAF draft specification (see [9, Section 7.3]). We recall the relevant aspects of that construction below, as needed.

In Doplar, we will use our DFLP construction for the language $\mathcal{L} = \{0, 1\}$ — i.e., we use it to prove that a value is zero or one. In this case, we instantiate the underlying FLP with the circuit:

$$C((s, t, \Delta), r) = (r \cdot s(s-1) + r^2 \cdot (s \cdot \Delta - t))^2, \quad (1)$$

where r denotes the joint randomness. This circuit recognizes the set of inputs $(s, t, \Delta) \in \{(0, 0, \Delta), (1, \Delta, \Delta)\}$. Note that FLP has input length $n = 3$ and joint-randomness length $jl = 1$; its circuit has 3 multiplication gates ($s(s-1)$, $s\Delta$, and the outer square). In the more general case, FLP will be instantiated for the language $\{(s, t, \Delta) \mid s \in \mathcal{L} \wedge s\Delta = t\}$. If the circuit for membership in \mathcal{L} has n multiplication gates, then FLP will be instantiated with a circuit with $n + 2$ multiplication gates.

Completeness follows immediately from the perfect completeness of the underlying FLP. The FLP of [17] has soundness $2n' / (|\mathbb{F}| - n')$ when its circuit has n' multiplication gates. We instantiate that FLP with $n' = n + 2$, and we also incur a factor 2 loss in soundness since our construction verifies two proofs in the underlying FLP. Hence, we obtain the soundness bound stated in the lemma.

<p><u>DFLP*.Prove($\{\vec{x}_1, \vec{x}_2\}, \Delta, jr$):</u></p> <pre> 1 $(jr_1, jr_2) \leftarrow jr$ 2 $\vec{e}_1 \leftarrow \text{FLP.Encode}(\Delta, \vec{x}_1)$ 3 $\vec{e}_2 \leftarrow \text{FLP.Encode}(\Delta, \vec{x}_2)$ 4 $b \leftarrow_s \{1, 2\}$ 5 $\pi_b \leftarrow_s \text{Prove}(\vec{e}_1, \Delta, jr_b)$ 6 $\pi_{3-b} \leftarrow_s \text{Prove}(\vec{e}_2, \Delta, jr_{3-b})$ 7 ret (π_1, π_2) </pre> <p><u>DFLP*.Query($\vec{e}, \Delta, jr; qr$):</u></p> <pre> 8 $(jr_1, jr_2) \leftarrow jr; (qr_1, qr_2) \leftarrow qr$ 9 $\sigma_1 \leftarrow_s \text{FLP.Query}(\vec{e}, \Delta, \pi_1, jr_1; qr_1)$ 10 $\sigma_2 \leftarrow_s \text{FLP.Query}(\vec{e}, \Delta, \pi_2, jr_2; qr_2)$ 11 ret (σ_1, σ_2) </pre>	<p><u>DFLP*.Decide(σ):</u></p> <pre> 12 $(\sigma_1, \sigma_2) \leftarrow \sigma$ 13 ret $\text{FLP.Decide}(\sigma_1)$ 14 $\vee \text{FLP.Decide}(\sigma_2)$ </pre> <p><u>DFLP*.Encode($\Delta \in \mathbb{F}, \vec{x} \in \mathbb{F}^n$):</u></p> <pre> 15 for $i \in [n]$: 16 $\vec{e}[i] \leftarrow \vec{x}[i]$ 17 $\vec{e}[i+n] \leftarrow \Delta \cdot \vec{x}[i]$ 18 ret \vec{e} </pre> <p><u>DFLP*.Decode($\vec{e} \in \mathbb{F}^{2n}$):</u></p> <pre> 19 ret $\vec{e}[1:n]$ </pre>
--	--

Figure 12: Delayed-2-input FLP construction $\text{DFLP}^*[\text{FLP}]$. The construction should be instantiated where FLP is the FLP for arithmetic circuits from [17].

The zero-knowledge simulator for our construction is given as S in Figure 13. To demonstrate privacy, we first consider the hybrid on the left of Figure 13. With the gray box included and white box excluded, the hybrid generates exactly the honest verifier’s view. In this game, both proofs are queried on x_c , the adversary’s choice. Note that proof $\pi_{b \oplus c}$ was generated with input x_c in mind, while $\pi_{b \oplus c \oplus 1}$ was not. Let $u = b \oplus c \oplus 1$, the index of the “mismatched” proof (i.e., π_u was generated with $x_{c \oplus 1}$ in mind, not x_c). By applying the linearity of $\text{Encode}(\Delta, \cdot)$ and $\text{Query}(\cdot, \cdot, \cdot)$, we can write:

$$\text{Query}(\text{Encode}(\Delta, x_c), \Delta, \pi_u, jr_u; qr_u) = \text{Query}(\text{Encode}(\Delta, x_{c \oplus 1}), \Delta, \pi_u, jr_u; qr_u) + \text{Query}(\text{Encode}(\Delta, x_c - x_{c \oplus 1}), 0, \vec{0}, jr_u; qr_u)$$

Making this change of notation in the game yields hybrid \underline{G}_1 (in Figure 13, gray boxes excluded and outlined box included). \underline{G}_1 is distributed identically to the original privacy game.

In each call to Query in \underline{G}_1 that involves a value π_i , we use the same input that was used to generate π_i . Hence, we can apply the zero-knowledge property of the underlying FLP to each such expression. In doing so, we obtain the hybrid \underline{G}_2 on the right of Figure 13. The underlying FLP of [17] has perfect zero-knowledge, so \underline{G}_2 is distributed identically to the original game.

To complete the proof, it suffices to show that σ_u is distributed pseudorandomly in $\mathbb{F}^3 \times \{0\}$, since the simulator samples σ_u uniformly from that set. In particular, when $\vec{\sigma}$ is distributed as in a simulated proof, Δ is random, and $d \neq 0$, what is the distribution on $\vec{\sigma} + \text{Query}(\text{Encode}(\Delta, d), \vec{0}, \vec{0}, \dots)$?

To answer this question, we must use specific properties of the FLP from [17]. We first briefly review the main idea behind their proof. The prover defines two polynomials L and R such that, for each multiplication gate i in the verification circuit, the value on its left wire is $L(i)$ and its right wire $R(i)$. Additionally, $L(0)$ and $R(0)$ are chosen uniformly. Define the “gadget” polynomial $G = L \times R$ — then $G(i)$ is the value of the output wire of the i th gate.

The proof vector π then consists of $L(0)$, $R(0)$, and the coefficients of the G polynomial. With that in mind, the Query algorithm makes 4 linear queries to the input + proof vector:

1. Obtain evaluations of the polynomial L as follows:
 - $L(0)$ is part of the proof vector.
 - For $i > 0$, if the left input to gate i is an input to the circuit, then $L(i)$ is given as part of the proof input/instance, to which Query has access.
 - Otherwise, the left input to gate i is the output of some other multiplication gate j . This value can be obtained as $G(j)$, since the coefficients of G are included in the proof vector.

Reconstruct L as the result of Lagrange interpolation over the points $\{(i, L(i))\}$. Evaluate this polynomial L at point qr (the query randomness).

2. Similarly, reconstruct R and evaluate it at point qr .
3. Evaluate the polynomial G at point qr .
4. Evaluate G at point m , where the output of verification circuit is the output wire of the m 'th multiplication gate.

Suppose the results of these queries are (r, s, t, u) ; the Decide algorithm checks that $t = rs$ and $u = 0$. The zero-knowledge property is that the result of the queries is distributed as $(r, s, rs, 0)$ for uniform $r, s \leftarrow_s \mathbb{F}$.

With Query as above, we now consider the distribution of

$$(r, s, rs, 0) + \text{Query}(\text{Encode}(\Delta, d), \vec{0}, \vec{0}, \dots),$$

where Δ, r, s are uniform in \mathbb{F} .

- The first component of this expression is uniform due to r .
- With overwhelming probability $1 - 1/|\mathbb{F}|$ we have $r \neq 0$. Conditioned on $r \neq 0$, the third component of the expression is uniform, since it is masked with rs , and s is uniform (even conditioned on the first component).
- Let q_4 be the 4th component of Query's output in the above expression. By definition of Query, q_4 is the result of evaluating G at point qr . But in this expression, the "proof vector" argument to Query is all zeroes, hence Query evaluates the all-zeroes polynomial and outputs $q_4 = 0$. Hence the 4th component of the overall expression is zero.
- Let q_2 be the second output of Query in the above expression. We see that q_2 is the result of evaluating polynomial R at point qr , after reconstructing R as described above. Fix a position i in which $\vec{d}[i] \neq 0$. Then the $(n+i)$ th position of $\vec{e} = \text{Encode}(\Delta, \vec{d})$ is $\vec{e}[i] = \vec{d}[i]\Delta$, and therefore is uniformly distributed when Δ is uniformly distributed.

The final multiplication gate in the verification circuit is the outermost square in (1). The input to this squaring operation is a linear combination that includes $\vec{e}[i]$. So as $\vec{e}[i]$ is uniformly distributed, the input to this multiplication gate is also uniformly distributed. Then the result of interpolating polynomial R (based on $\vec{e}[i]$ among other values) and evaluating R at qr is also uniformly distributed. In other words, q_2 is uniformly distributed over uniform choice of Δ , so the second component of the above expression is uniform.

Overall, we have shown that the distribution of σ_u in \mathcal{G}_2 of Figure 13 is statistical distance $1/|\mathbb{F}|$ from the simulator's distribution: uniform over $\mathbb{F}^3 \times \{0\}$. Hence in \mathcal{G}_2 the adversary has advantage bounded by $1/|\mathbb{F}|$ in \mathcal{G}_2 . \square

C Proofs of Theorems

C.1 Prio3 Robustness (Theorem 1)

We begin by instantiating the robustness game for Π in Figure 14. We expand the Prep algorithm and make a few simplifications to the game's internal notation and bookkeeping. First, the game $\text{Exp}^{\text{robust}}$ calls for a VDAF with an arbitrary number of rounds, but Prio3 constructions has just one round. Second, we know that the Prep algorithm will be called exactly twice for each aggregator, and that the initial broadcast message and state are empty. We Therefore unroll the loop of lines 5–14 of Figure 3 and evaluate those if-statements whose values are pre-determined. Third, we replace table St with a vector \vec{st} and remove table Msg altogether. (The transcript output by the oracle is now constructed at the end on line 21 on the left-hand panel of Figure 14.) Fourth, we evaluate Prep in parallel for all aggregators instead of in sequence; the order of these operations does not affect their results because aggregators do not share state. Fifth, we

<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $\text{Exp}_{\text{DFLP}, S}^{\text{priv}}(A): \quad \underline{\mathbf{G}}_1(A):$ </div> <pre style="font-family: monospace; font-size: 0.9em;"> 1 $b \leftarrow \{0, 1\}$ 2 $(\{x_0, x_1\}, st_A) \leftarrow A()$ 3 if $b = 0$: 4 $(st_S, (jr_0, jr_1), (qr_0, qr_1)) \leftarrow S_1()$ 5 else: 6 $jr_0, jr_1 \leftarrow \mathbb{F}^{j_l}; qr_0, qr_1 \leftarrow \mathbb{F}^{q_l}$ 7 $\Delta \leftarrow \mathbb{F}$ 8 // Prove($\{x_0, x_1\}, \Delta, jr$) 9 $b \leftarrow \{0, 1\}$ 10 $\pi_b \leftarrow \text{Prove}(\text{Encode}(\Delta, \vec{x}_0), \Delta, jr_b)$ 11 $\pi_{1 \oplus b} \leftarrow \text{Prove}(\text{Encode}(\Delta, \vec{x}_1), \Delta, jr_{1 \oplus b})$ 12 $(c, st_A) \leftarrow A(st_A, (jr_0, jr_1), (qr_0, qr_1))$ 13 if $b = 0$: $(\sigma_0, \sigma_1) \leftarrow S^*(st_S)$ 14 else: // Query($\text{Encode}(\Delta, x_c), \Delta, \pi, jr; qr$) 15 $\sigma_0 \leftarrow \text{Query}(\text{Encode}(\Delta, x_c), \Delta, \pi_0, jr_0; qr_0)$ 16 $\sigma_1 \leftarrow \text{Query}(\text{Encode}(\Delta, x_c), \Delta, \pi_1, jr_1; qr_1)$ 17 $\sigma_b \leftarrow \text{Query}(\text{Encode}(\Delta, x_0), \Delta, \pi_b, jr_b; qr_b)$ 18 $\sigma_{1 \oplus b} \leftarrow \text{Query}(\text{Encode}(\Delta, x_1), \Delta, \pi_{1 \oplus b}, jr_{1 \oplus b}; qr_{1 \oplus b})$ 19 $u \leftarrow b \oplus c \oplus 1$ // index of "mismatched" proof 20 $\sigma_u \leftarrow \sigma_u$ 21 + Query($\text{Encode}(\Delta, x_c - x_{1 \oplus c}), 0, \vec{0}, jr_u; qr_u$) 22 $b' \leftarrow A(st_A, (\sigma_0, \sigma_1))$ 23 ret $b == b'$ </pre>	<div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $\underline{\mathbf{G}}_2(A):$ </div> <pre style="font-family: monospace; font-size: 0.9em;"> 1 $b \leftarrow \{0, 1\}$ 2 $(\{x_0, x_1\}, st_A) \leftarrow A()$ 3 if $b = 0$: 4 $(st_S, (jr_0, jr_1), (qr_0, qr_1)) \leftarrow S_1()$ 5 else: 6 $\Delta \leftarrow \mathbb{F}^n$ 7 $b \leftarrow \{0, 1\}$ 8 $(jr_0, qr_0, \sigma_0) \leftarrow S_{\text{FLP}}()$ 9 $(jr_1, qr_1, \sigma_1) \leftarrow S_{\text{FLP}}()$ 10 $(c, st) \leftarrow A(st_A, (jr_0, jr_1), (qr_0, qr_1))$ 11 if $b = 0$: $(\sigma_0, \sigma_1) \leftarrow S^*(st_S)$ 12 else: 13 $u \leftarrow b \oplus c \oplus 1$ // index of "mismatched" proof 14 $\sigma_u \leftarrow \sigma_u$ 15 + Query($\text{Encode}(\Delta, x_c - x_{1 \oplus c}), 0, \vec{0}, jr_u; qr_u$) 16 $b' \leftarrow A(st_A, (\sigma_0, \sigma_1))$ 17 ret $b == b'$ </pre> <div style="border: 1px solid black; padding: 2px; margin-bottom: 5px;"> $S_1():$ <pre style="font-family: monospace; font-size: 0.8em;"> 1 $(jr_0, qr_0, \sigma_0) \leftarrow S()$ 2 $(jr_1, qr_1, \sigma_1) \leftarrow S()$ 3 ret $(st_S = (\sigma_0, \sigma_1), (jr_0, jr_1), (qr_0, qr_1))$ </pre> </div> <div style="border: 1px solid black; padding: 2px;"> $S_2(\sigma_0, \sigma_1):$ <pre style="font-family: monospace; font-size: 0.8em;"> 4 $b \leftarrow \{0, 1\}$ 5 $\sigma_b \leftarrow \mathbb{F}^3 \times \{0\}$ // overwrite σ_b 6 ret (σ_0, σ_1) </pre> </div>
---	--

Figure 13: Hybrids for zero-knowledge property of the delayed-2-input FLP construction.

perform the deterministic Decide operation only once since its result is the same for all aggregators. Finally, we replace each call to RG_i with a call to the corresponding random oracle RO_i . Let q_i denote the number of queries A makes to RO_i ; note that $q_{\text{RG}} = q_1 + \dots + q_7$.

We have also dropped winning condition on line 16 of Figure 3. By definition, $\Pi.\text{refineFromShares}(\varepsilon, \vec{x}) = \Pi.\text{Unshard}(1, \Pi.\text{Unshard}(1, (\Pi.\text{Agg}(\text{inp}[1]), \dots, \Pi.\text{Agg}(\text{inp}[s])))$, where $\text{inp}[\hat{j}]$ is the unpacked inner measurement share of input share $\vec{x}[\hat{j}]$ for each \hat{j} . Thus w can never be set by forcing the refined shares to mismatch the expected refined measurement.

Now we express the proof with a series of incrementally changed games, beginning with $\underline{\mathbf{G}}_1$ (c.f. Figure 14). The joint randomness for each aggregator \hat{j} is derived in $\text{Exp}_{\Pi}^{\text{robust}}$ from the seed $j\text{seed}_{\hat{j}}$ of that aggregator, which is also the state $\vec{st}[\hat{j}]$. In $\underline{\mathbf{G}}_1$, we instead derive joint randomness from $\vec{st}[1]$ for all aggregators, thus ensuring that the joint randomness is the same for everyone.

We build a wrapper adversary B for which

$$\text{Adv}_{\Pi}^{\text{robust}}(A) \leq \Pr[\underline{\mathbf{G}}_1(B)] + \frac{q_5}{2^\kappa}. \quad (2)$$

Adversary B only makes queries to Prep that set $\vec{st}[\hat{j}] = \vec{st}[1]$ for all \hat{j} . It accomplishes this by calculating $\vec{st}[\hat{j}]$ for every aggregator and Prep query made by A . If it finds that $\vec{st}[\hat{j}] = \vec{st}[1]$ for all aggregators, it forwards the query to its own Prep oracle. Otherwise, it runs Prep itself. B can perfectly simulate Prep except for line 9, because it does not know sk . Instead, B picks its own verification key sk' and uses sk' in line 9 where Prep would use sk . Adversary A can detect the substitution of sk' for sk in two cases: If A queries RO_5 on seed sk' ; or if the Prep oracle and B query RO_5 on the same context string. The latter event does not occur because each query to RO_5 contains a unique nonce. The former occurs with probability at most $\frac{q_5}{2^\kappa}$, because sk' is a uniformly random κ -bit string. The queries that B simulates would always set

<p>Game $\text{Exp}_H^{\text{robust}}(A)$ $\boxed{\mathbb{G}_1(A)}$:</p> <ol style="list-style-type: none"> 1 $w \leftarrow \text{false}; sk \leftarrow_s \{0, 1\}^\kappa$ 2 $A^{\text{RO.Prep}}(); \text{ret } w$ <p>$\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top$ 5 for $\hat{j} \in [s]$: 6 $(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}; \vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \vec{inp}[\hat{j}])$ 8 $\text{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed 10 $\boxed{jr \leftarrow \text{RO}_1(\vec{st}[\hat{j}], \varepsilon)}$ $\boxed{jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)}$ 11 $qr \leftarrow \text{RO}_5(sk, n)$ 12 $\vec{vfs}[\hat{j}] \leftarrow \text{Query}(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr)$ 13 $vf \leftarrow \sum_{\hat{j}=1}^s \vec{vfs}[\hat{j}]$ 14 $d \leftarrow \text{FLP.Decide}(vf)$ 15 for $\hat{j} \in [s]$: 16 $jseed_{\hat{j}} \leftarrow \vec{st}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, \vec{rseed}_{\hat{j}})$ 17 $acc_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 18 $w \leftarrow (w \vee [acc_{\hat{j}} \wedge \text{refineFromShares}(\varepsilon, \vec{x}) \notin \mathcal{L}])$ 19 ret $(w, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \vec{rseed}[\hat{j}]))_{\hat{j} \in [s]})$ 	<p>Adversary $B^{\text{RO.Prep}}()$:</p> <ol style="list-style-type: none"> 1 $sk' \leftarrow_s \{0, 1\}^\kappa$ 2 $A^{\text{RO.PrepSim}}()$ <p>$\text{PrepSim}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top; fwd \leftarrow \text{true}$ 5 for $\hat{j} \in [s]$: 6 $(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}; \vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \vec{inp}[\hat{j}])$ 8 $\text{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // Joint rand seed 10 if $\vec{st}[\hat{j}] \neq \vec{st}[1]$: $fwd \leftarrow \text{false}$ 11 $jr \leftarrow \text{RO}_1(\vec{st}[\hat{j}], \varepsilon)$ 12 $qr \leftarrow \text{RO}_5(sk', n)$ 13 $\vec{vfs}[\hat{j}] \leftarrow \text{Query}(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr)$ 14 if fwd: return $\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$ 15 ret $(\text{false}, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \text{rseed}[\hat{j}]))_{\hat{j} \in [s]})$
--	--

Figure 14: Left: Definition of game \mathbb{G}_1 for the proof of Theorem 1. Also shown is the robustness game for Π and adversary A with some simplifications applied. Right: Adversary B .

$acc_{\hat{j}} \leftarrow 0$ in line 17. Thus any query that would set $w \leftarrow \text{true}$ is forwarded to the Prep oracle by B , and B wins whenever A does. The claim follows.

Next, we use the full linearity of FLP to decompose FLP.Query into algorithm Q and a matrix multiplication operation, as shown in the left-hand panel of Figure 15. Q is a randomized algorithm, but it is executed deterministically with fixed input jr and coins qr . We may therefore call Q only once to eliminate redundancy. Finally, we sum the vectors $\vec{x}_{\hat{j}} \parallel \pi_{\hat{j}}$ before the multiplication instead of multiplying then summing the products. This preserves the output thanks to the associativity of matrix multiplication.

Full linearity is an information theoretic property that holds unconditionally for all inputs, proofs, and coins, so the new game computes the same verifier string vf . Thus

$$\Pr[\mathbb{G}_1(B)] = \Pr[\mathbb{G}_2(B)]. \quad (3)$$

We produce the next modified game, in the right-hand panel of Figure 15, to define variables $\text{inp} = \sum_{\hat{j}=1}^s \vec{inp}[\hat{j}]$ and $\pi = \sum_{\hat{j}=1}^s \vec{\pi}[\hat{j}]$ and invoke PRG.Query on inp, π directly. In addition, from Figure 5, we can see that $\text{inp} = \Pi.\text{refineFromShares}(\varepsilon, \vec{x})$, so we substitute inp into line 21. By the full linearity of FLP , we have that $Q(jr; qr) \cdot (\text{inp} \parallel \pi) = \text{FLP.Query}(\text{inp}, \pi, jr; qr)$. Again, these operations do not affect the adversary's view of Prep , and

$$\Pr[\mathbb{G}_2(B)] = \Pr[\mathbb{G}_3(B)]. \quad (4)$$

In the next game, we replace the pseudorandom query randomness qr with a fresh random string that is implicitly sampled by Query . We bound the difference in advantage between games \mathbb{G}_3 and \mathbb{G}_4 via a reduction B' to the pseudorandomness of RG_5 . The reduction honestly simulates \mathbb{G}_3 except in line 11, where it queries its challenge oracle on n and sets qr to the response. Because every nonce is unique, these queries are all distinct. When the challenge oracle is a random function, this is a perfect simulation of \mathbb{G}_4 ; otherwise it is a perfect simulation of \mathbb{G}_3 . Adversary B' makes q_{Prep} queries to its challenge oracle; when RG_5 is modeled as a random oracle, there is a maximum of q_5 random oracle queries.

The generic PRF advantage for a (q_5, q_{Prep}) -query attacker against a random oracle with domain $\{0, 1\}^\kappa$ is bounded by the probability $\frac{q_5}{2^\kappa}$ that the attacker makes a random oracle query containing sk . Thus

<p>Game $\underline{\mathbb{G}}_1(B)$ $\underline{\mathbb{G}}_2(B)$:</p> <ol style="list-style-type: none"> 1 $w \leftarrow \mathbf{false}$; $sk \leftarrow_s \{0, 1\}^\kappa$ 2 $B^{\text{RO.Prep}}(); \text{ret } w$ <p>$\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top$ 5 for $\hat{j} \in [s]$: 6 $(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}$; $\vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \vec{inp}[\hat{j}])$ 8 $\vec{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed 10 $jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)$; $qr \leftarrow \text{RO}_5(sk, n)$ 11 $\vec{vfs}[\hat{j}] \leftarrow \text{Query}(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr)$ 12 $vf \leftarrow \sum_{j=1}^s \vec{vfs}[\hat{j}]$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 13 $jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)$; $qr \leftarrow \text{RO}_5(sk, n)$ 14 $Z \leftarrow Q(jr; qr)$ 15 $vf \leftarrow Z \cdot \sum_{j=1}^s \vec{inp}[\hat{j}] \parallel \vec{\pi}[\hat{j}]$ </div> <ol style="list-style-type: none"> 16 $d \leftarrow \text{FLP.Decide}(vf)$ 17 for $\hat{j} \in [s]$: 18 $jseed_{\hat{j}} \leftarrow \vec{st}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, \vec{rseed})$ 19 $acc_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 20 $w \leftarrow (w \vee [acc_{\hat{j}} \wedge \text{refineFromShares}(\varepsilon, \vec{x}) \notin \mathcal{L}])$ 21 ret $(w, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \vec{rseed}[\hat{j}]))_{\hat{j} \in [s]})$ 	<p>Game $\underline{\mathbb{G}}_2(B)$ $\underline{\mathbb{G}}_3(B)$:</p> <ol style="list-style-type: none"> 1 $w \leftarrow \mathbf{false}$; $sk \leftarrow_s \{0, 1\}^\kappa$ 2 $B^{\text{RO.Prep}}(); \text{ret } w$ <p>$\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top$ 5 for $\hat{j} \in [s]$: 6 $(\vec{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}$; $\vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \vec{inp}[\hat{j}])$ 8 $\vec{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed 10 $jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)$ 11 $qr \leftarrow \text{RO}_5(sk, n)$ 12 $Z \leftarrow Q(jr; qr)$ 13 $vf \leftarrow Z \cdot \sum_{j=1}^s \vec{inp}[\hat{j}] \parallel \vec{\pi}[\hat{j}]$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 14 $\vec{inp} \leftarrow \sum_{j=1}^s \vec{inp}[\hat{j}]; \pi \leftarrow \sum_{j=1}^s \vec{\pi}[\hat{j}]$ 15 $vf \leftarrow \text{FLP.Query}(\vec{inp}, \pi, jr; qr)$ </div> <ol style="list-style-type: none"> 16 $d \leftarrow \text{FLP.Decide}(vf)$ 17 for $\hat{j} \in [s]$: 18 $jseed_{\hat{j}} \leftarrow \vec{st}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, \vec{rseed})$ 19 $acc_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 20 $w \leftarrow (w \vee [acc_{\hat{j}} \wedge \text{refineFromShares}(\varepsilon, \vec{x}) \notin \mathcal{L}])$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 21 $w \leftarrow (w \vee [acc_{\hat{j}} \wedge \vec{inp} \notin \mathcal{L}])$ </div> <ol style="list-style-type: none"> 22 ret $(w, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \vec{rseed}[\hat{j}]))_{\hat{j} \in [s]})$
--	--

Figure 15: Game $\underline{\mathbb{G}}_2$ (left) and game $\underline{\mathbb{G}}_3$ (right) for the proof of Theorem 1.

$$\Pr[\underline{\mathbb{G}}_3(B)] \leq \Pr[\underline{\mathbb{G}}_4(B)] + \frac{q_5}{2^\kappa}. \quad (5)$$

Our next game ($\underline{\mathbb{G}}_5$ defined in the right-hand panel of Figure 16) differs from $\underline{\mathbb{G}}_4$ as follows. We set a **bad** flag and force the adversary to lose if it makes two queries to Prep which derive their joint randomness from the same seed. Each query to Prep derives its joint randomness seed from a unique nonce, so duplicate seeds require a collision between two queries to RO_6 or between two vectors of hints. Both seeds and hints are randomly sampled by random oracles RO_6 and RO_7 respectively, so we limit the probability of both types of collision with a birthday bound over the q_{Prep} queries to Prep :

$$\frac{q_{\text{Prep}}^2}{2^{\kappa+1}} + \frac{q_{\text{Prep}}^2}{2^{\kappa \cdot s+1}} < \frac{q_{\text{Prep}}^2}{2^\kappa}.$$

Since the games are identical until **bad** gets set, we have

$$\Pr[\underline{\mathbb{G}}_4(B)] \leq \Pr[\underline{\mathbb{G}}_5(B)] + \frac{q_{\text{Prep}}^2}{2^\kappa}. \quad (6)$$

We are now ready to reduce to FLP soundness. To do so, we construct a malicious prover P^* in Figure 17 from B whose advantage in the FLP soundness experiment is related to B 's advantage in winning game $\underline{\mathbb{G}}_5$. Recall from Figure 2 that the prover is called twice, first to choose an input and a second time to generate a proof. The prover is given joint randomness jr in this second call, after committing to the input. Thus, in our reduction we must extract this input from B random oracle queries, then program the random oracle with jr before proceeding.

<p>Game $\underline{\mathbb{G}}_3(B)$ $\underline{\mathbb{G}}_4(B)$:</p> <ol style="list-style-type: none"> 1 $w \leftarrow \mathbf{false}$; $sk \leftarrow_{\\$} \{0, 1\}^\kappa$ 2 $B^{\text{RO.Prep}}();$ ret w <p>$\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top$ 5 for $\hat{j} \in [s]$: 6 $(\text{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}$; $\vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \text{inp}[\hat{j}])$ 8 $\text{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed 10 $jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)$ 11 $qr \leftarrow \text{RO}_5(sk, n)$ 12 $\text{inp} \leftarrow \sum_{\hat{j}=1}^s \text{inp}[\hat{j}]; \pi \leftarrow \sum_{\hat{j}=1}^s \vec{\pi}[\hat{j}]$ 13 $vf \leftarrow \text{FLP.Query}(\text{inp}, \pi, jr, qr)$ 14 $vf \leftarrow_{\\$} \text{FLP.Query}(\text{inp}, \pi, jr)$ 15 $d \leftarrow \text{FLP.Decide}(vf)$ 16 for $\hat{j} \in [s]$: 17 $jseed_{\hat{j}} \leftarrow \vec{st}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, \vec{rseed})$ 18 $\text{acc}_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 19 $w \leftarrow (w \vee [\text{acc}_{\hat{j}} \wedge \text{inp} \notin \mathcal{L}])$ 20 ret $(w, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \vec{rseed}[\hat{j}]))_{\hat{j} \in [s]})$ 	<p>Game $\underline{\mathbb{G}}_4(B)$ $\underline{\mathbb{G}}_5(B)$:</p> <ol style="list-style-type: none"> 1 $w \leftarrow \mathbf{false}$; $sk \leftarrow_{\\$} \{0, 1\}^\kappa$; $\mathcal{J} \leftarrow \emptyset$ 2 $B^{\text{RO.Prep}}();$ ret w <p>$\text{Prep}(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$:</p> <ol style="list-style-type: none"> 3 if $\text{Used}[n] \neq \perp$: ret \perp 4 $\text{Used}[n] \leftarrow \top$ 5 for $\hat{j} \in [s]$: 6 $(\text{inp}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 7 $(\vec{\rho}, \cdot) \leftarrow \text{msg}$; $\vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \text{inp}[\hat{j}])$ 8 $\text{rseed}[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 9 $\vec{st}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed <div style="border: 1px solid black; padding: 2px; margin: 5px 0;"> <ol style="list-style-type: none"> 10 if $\vec{st}[1] \in \mathcal{J}$: $\mathbf{bad} \leftarrow \mathbf{true}$ 11 $\mathcal{J} \leftarrow \mathcal{J} \cup \{\vec{st}[1]\}$ </div> <ol style="list-style-type: none"> 12 $jr \leftarrow \text{RO}_1(\vec{st}[1], \varepsilon)$ 13 $\text{inp} \leftarrow \sum_{\hat{j}=1}^s \text{inp}[\hat{j}]; \pi \leftarrow \sum_{\hat{j}=1}^s \vec{\pi}[\hat{j}]$ 14 $vf \leftarrow_{\\$} \text{FLP.Query}(\text{inp}, \pi, jr)$ 15 $d \leftarrow \text{FLP.Decide}(vf)$ 16 for $\hat{j} \in [s]$: 17 $jseed_{\hat{j}} \leftarrow \vec{st}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, \vec{rseed})$ 18 $\text{acc}_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 19 $w \leftarrow (w \vee [\neg \mathbf{bad} \wedge \text{acc}_{\hat{j}} \wedge \text{inp} \notin \mathcal{L}])$ 20 ret $(w, (\text{msg}_{\text{Init}}, (\vec{vfs}[\hat{j}], \vec{rseed}[\hat{j}]))_{\hat{j} \in [s]})$
---	--

Figure 16: Fourth and fifth intermediate games for the proof of Theorem 1.

The malicious prover P^* runs B in a simulation of $\underline{\mathbb{G}}_5$. Its random oracle queries are answered by lazy-evaluating a table Rand ; all oracle queries are handled the same way except for a distinguished query, which will be programmed using the jr string generated as part of the malicious prover's experiment. At the start of the simulation, the prover P^* samples $i^* \leftarrow_{\$} [q_1 + q_{\text{Prep}}]$. On the i^* unique invocation of RO_1 (see ROExt_1 in Figure 17), the prover checks the table Rand for a nonce n and input shares $\text{inp}_1, \dots, \text{inp}_s$ that give rise to the seed $jseed$ provided as input. If successful, the prover outputs $\text{inp}_1 + \dots + \text{inp}_s$ as its challenge input, awaits the response jr , and sets $\text{Rand}[1, jseed, \varepsilon] \leftarrow jr$ (line 11). It also records $n^* \leftarrow n$ for use later on.

The simulation of Prep queries is identical except after two events. First, the prover P^* halts and concedes if two Prep queries generate the same joint randomness seed $\vec{st}[1]$. (Adversary B loses in this case.) Second, if $n = n^*$, then P^* immediately halts and outputs the proof π computed on line 30. If the simulation has reached this point, then the probability that P^* wins its game is at least the probability that the game sets $w \leftarrow \mathbf{true}$ on line 36. Conditioning on the probability that P^* guesses the winning query to RO_1 , we have that

$$\Pr[\underline{\mathbb{G}}_5(B)] \leq (q_1 + q_{\text{Prep}}) \cdot \epsilon. \quad (7)$$

The claimed bound follows by gathering up all of the bounds across the games and simplifying. \square

C.2 Prio3 Privacy (Theorem 2)

We begin by instantiating the privacy game $\text{Exp}_{II,t}^{\text{priv}}$ for Prio3 VDAF II . Game $\underline{\mathbb{G}}_0$ in Figure 18 was constructed by inlining II 's constituent algorithms and cleaning up the control flow. In addition, calls to RG_i have been substituted with calls to a random oracle RO_i . Let q_i denote the number of queries A makes to RO_i ; note that $q_{\text{RG}} = q_1 + \dots + q_7$.

In our first game hop, we modify Shard oracle's behavior after setting flag \mathbf{bad}_1 on line 7. In the new game, $\underline{\mathbb{G}}_1$ (Figure 18), the nonce n is sampled without replacement, ensuring that each nonce is used is unique.

<p>Adversary $P^*[B](\cdot)$:</p> <pre> 1 $w \leftarrow \text{false}; sk \leftarrow_{\\$} \{0, 1\}^\kappa; \text{bad} \leftarrow \text{false}; \mathcal{J} \leftarrow \emptyset$ 2 $ctr \leftarrow 0; n^* \leftarrow \perp; i^* \leftarrow_{\\$} [q_1 + q_{\text{Prep}}]$ 3 $B^{\text{ROExt}_1, \text{RO}_2, \dots, \text{RO}_7 \text{PrepSim}(\cdot)}$; ret w ROExt₁(seed, cntxt): 4 if Rand[1, seed, cntxt] $\neq \perp$: ret $\text{RO}_1(\text{seed}, \text{cntxt})$ 5 $ctr \leftarrow ctr + 1$ 6 if $ctr = i^* \wedge (\exists n, (\text{blind}_{\hat{j}}, \text{inp}_{\hat{j}}, \rho_{\hat{j}})_{\hat{j} \in [s]})$ 7 $(\forall \hat{j}) \text{Rand}[7, \text{blind}_{\hat{j}}, \hat{j} \parallel n \parallel \text{inp}_{\hat{j}}] = \rho_{\hat{j}}$ 8 $\wedge \text{Rand}[6, 0^\kappa, (\rho_1, \dots, \rho_s)] = \text{seed}$: 9 output $\text{inp}_1 + \dots + \text{inp}_s$ and wait for jr. 10 $n^* \leftarrow n$; Rand[1, seed, cntxt] $\leftarrow jr$ 11 ret $\text{RO}_1(\text{seed}, \text{cntxt})$ RO_i(seed, cntxt): 12 $l \leftarrow (jl, n, m, pl, ql)$ 13 if Rand[i, seed, cntxt] $= \perp$: 14 if $i \leq 5$: Rand[i, seed, cntxt] $\leftarrow_{\\$} \mathbb{F}^{l[i]}$ 15 else: Rand[i, seed, cntxt] $\leftarrow_{\\$} \{0, 1\}^\kappa$ 16 ret Rand[i, seed, cntxt]</pre>	<p>PrepSim($n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}}$):</p> <pre> 17 if Used[$n$] $\neq \perp$: ret \perp 18 Used[n] $\leftarrow \top$ 19 for $\hat{j} \in [s]$: 20 $(\vec{\text{inp}}[\hat{j}], \vec{\pi}[\hat{j}], \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \vec{x}[\hat{j}])$ 21 $(\vec{\rho}, \cdot) \leftarrow \text{msg}; \vec{\rho}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \vec{\text{inp}}[\hat{j}])$ 22 $rseed[\hat{j}] \leftarrow \vec{\rho}[\hat{j}]$ 23 $\vec{\text{st}}[\hat{j}] \leftarrow \text{RO}_6(0^\kappa, \vec{\rho})$ // joint rand seed 24 if $\vec{\text{st}}[1] \in \mathcal{J}$: bad $\leftarrow \text{true}$; halt. 25 $\mathcal{J} \leftarrow \mathcal{J} \cup \{\vec{\text{st}}[1]\}$ 26 $jr \leftarrow \text{ROExt}_1(\vec{\text{st}}[1], \varepsilon)$ 27 $\text{inp} \leftarrow \sum_{\hat{j}=1}^s \vec{\text{inp}}[\hat{j}]; \pi \leftarrow \sum_{\hat{j}=1}^s \vec{\pi}[\hat{j}]$ 28 if $n = n^*$: output π and halt. 29 $vf \leftarrow_{\\$} \text{FLP.Query}(\text{inp}, \pi, jr)$ 30 $d \leftarrow \text{FLP.Decide}(vf)$ 31 for $\hat{j} \in [s]$: 32 $jseed_{\hat{j}} \leftarrow \vec{\text{st}}[\hat{j}]; jseed'_{\hat{j}} \leftarrow \text{RO}_6(0^\kappa, rseed)$ 33 $\text{acc}_{\hat{j}} \leftarrow d \wedge [jseed_{\hat{j}} = jseed'_{\hat{j}}]$ 34 $w \leftarrow (w \vee [\neg \text{bad} \wedge \text{acc}_{\hat{j}} \wedge \text{inp} \notin \mathcal{L}])$ 35 ret $(w, (\text{msg}_{\text{Init}}, (\vec{v}fs[\hat{j}], rseed[\hat{j}]))_{\hat{j} \in [s]})$</pre>
--	---

Figure 17: Malicious prover P^* for the proof of Theorem 1. The lookup in the random oracle table Rand on lines 6–8 can be performed efficiently by creating a reverse-lookup table; we omit the details for brevity.

Applying the Fundamental Lemma of Game Playing [14], and using a birthday bound for the probability of bad_1 getting set,

$$\Pr[\underline{\mathcal{G}}_0(A)] \leq \Pr[\underline{\mathcal{G}}_1(A)] + \frac{q_{\text{Shard}}^2}{|\mathcal{N}|}. \quad (8)$$

Next we replace the adversary A with one that controls all but one aggregator. We construct such an adversary B as a wrapper around A , and show that B wins with at least the probability of A . The adversary B , defined in Figure 19, presents four oracles ShardSim , SetupSim , PrepSim , and AggSim to adversary A , each emulating an oracle in game $\underline{\mathcal{G}}_1$. Algorithms SetupSim , PrepSim , AggSim are computed by B just as the respective oracles in game $\underline{\mathcal{G}}_1$ except that queries pertaining to aggregator z are forwarded to B 's own oracles. Algorithm ShardSim forwards A 's query to Shard in the natural way, but returns the shares of the aggregators deemed honest by A .

We claim that B perfectly simulates $\underline{\mathcal{G}}_1(A)$. This is obvious for ShardSim and for queries for which $\hat{j} = z$; in these cases, B simply forwards its queries to the appropriate oracles without changing their inputs. The only difference is that Shard returns more input shares than A requests; B stores these extra input shares for its own use and does not reveal them. By construction, this is a subset of the input shares returned by the query.

When A makes queries to SetupSim , PrepSim , or AggSim with $\hat{j} \in \mathcal{V} \setminus \{z\}$, our wrapper adversary performs the operations of Setup , Prep , or Agg respectively. Effectively, adversary B uses its stored input shares to fill in entries of tables In , Batch , Setup , Status , St , and Out exactly as the real privacy game would. Since each entry is disambiguated by its \hat{j} , there is no overlap with the tables maintained by the game; every table entry read by B must first have been written by B and thus all the information it needs to simulate the game perfectly is accessible. It follows that

$$\Pr[\underline{\mathcal{G}}_1(A)] = \Pr[\underline{\mathcal{G}}_1(B)]. \quad (9)$$

In the next game hop (Figure 20) we make some simplifying changes, including cleaning up the bad_1 flag and substituting $\{z\}$ for \mathcal{V} and simplifying accordingly. (We do not highlight this change in Figure 20, as it is fairly straightforward.) We also make the following breaking change: In game $\underline{\mathcal{G}}_2$, we program the

<p>Game $\underline{\mathcal{G}}_0(A)$ $\underline{\mathcal{G}}_1(A)$:</p> <ol style="list-style-type: none"> 1 $(st_A, \mathcal{V}, (sk_j)_{j \in \mathcal{V}}) \leftarrow_s A^{\text{RO}}(); \mathcal{T} \leftarrow [s] \setminus \mathcal{V}$ 2 if $\mathcal{V} + t \neq s$ return \perp 3 $b \leftarrow_s \{0, 1\}; b' \leftarrow_s A^{\text{RO, Shard, Setup, Prep, Agg}}(st_A)$ 4 ret $b = b'$ <p>Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):</p> <ol style="list-style-type: none"> 5 if $\text{Used}[\hat{k}] \neq \perp$: ret \perp 6 $n \leftarrow_s \mathcal{N}$ 7 if $n \in \mathcal{N}^*$: bad₁ \leftarrow true; $n \leftarrow_s \mathcal{N} \setminus \mathcal{N}^*$ 8 $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 9 $inp \leftarrow \text{Encode}(m_0)$ 10 for $\hat{j} \in [2..s]$: 11 $blind_{\hat{j}}, xseed_{\hat{j}}, pseed_{\hat{j}} \leftarrow_s \{0, 1\}^\kappa$ 12 $\vec{x}[\hat{j}] \leftarrow \text{RO}_2(xseed_{\hat{j}}, \hat{j})$ 13 $rseed[\hat{j}] \leftarrow \text{RO}_7(blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}])$ 14 $\vec{x}[1] \leftarrow inp - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ 15 $blind_1 \leftarrow_s \{0, 1\}^\kappa; ps \leftarrow_s \{0, 1\}^\kappa$ 16 $rseed[1] \leftarrow \text{RO}_7(blind_1, 1 \parallel n \parallel \vec{x}[1])$ 17 $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$; $jr \leftarrow \text{RO}_1(jseed, \varepsilon)$ 18 $pr \leftarrow \text{RO}_4(ps, \varepsilon)$ 19 $\vec{\pi}[1] \leftarrow \text{Prove}(inp, jr, pr)$ 20 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{\hat{j}=2}^s \text{RO}_3(pseed_{\hat{j}}, \hat{j})$ 21 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 22 for $\hat{j} \in [2..s]$: 23 $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 24 $\text{Pub}[\hat{k}] \leftarrow rseed$; $\text{In}[\hat{k}, \cdot] \leftarrow \vec{x}$ 25 $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 26 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}})$ <p>Setup($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, st_{\text{Init}} \in \{\varepsilon\}$):</p> <ol style="list-style-type: none"> 27 if $\text{Status}[\hat{i}, \hat{j}] \neq \perp$ or $\text{Setup}[\cdot, \hat{j}] > 0$: ret \perp 28 $\text{Setup}[\hat{i}, \hat{j}] \leftarrow st_{\text{Init}}$ 29 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{running}$ 	<p>Prep($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}, \hat{k} \in \mathbb{N}, m_{\text{msg}} \in \mathcal{M}^*$):</p> <ol style="list-style-type: none"> 30 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$ or $\text{In}[\hat{k}, \hat{j}] = \perp$: 31 ret \perp 32 if $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \perp$: 33 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \text{Setup}[\hat{i}, \hat{j}]$ 34 $m_{\text{msg}} \leftarrow (\text{Pub}[\hat{k}],)$ 35 $(n, m_0, m_1) \leftarrow \text{Used}[\hat{k}]$ 36 if $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \varepsilon$: // Process initial message from client 37 $(inp, \pi, blind) \leftarrow \text{Unpack}(\hat{j}, \text{In}[\hat{k}, \hat{j}])$ 38 $(rseed,)$ $\leftarrow m_{\text{msg}}$ 39 $rseed[\hat{j}] \leftarrow \text{RO}_7(blind, \hat{j} \parallel n \parallel inp)$ 40 $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$; $jr \leftarrow \text{RO}_1(jseed, \varepsilon)$ 41 $qr \leftarrow \text{RO}_5(sk_{\hat{j}}, n)$ 42 $msg \leftarrow (\text{Query}(inp, \pi, jr; qr), rseed[\hat{j}])$ 43 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow (jseed, \text{Truncate}(inp))$ 44 ret (running, msg) 45 // Process broadcast messages from aggregators 46 $(jseed, y) \leftarrow \text{St}[\hat{i}, \hat{j}, \hat{k}]$ 47 $(\vec{v}[\hat{j}], rseed[\hat{j}])_{\hat{j} \in [s]} \leftarrow m_{\text{msg}}$ 48 $acc \leftarrow \text{Decide}(\sum_{\hat{j}=1}^s \vec{v}[\hat{j}])$ 49 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp$ 50 if $acc = 0$ or $jseed \neq \text{RO}_6(0^\kappa, rseed)$: 51 ret (failed, \perp) 52 $\text{Out}[\hat{i}, \hat{j}, \hat{k}] \leftarrow y$ 53 $\text{Batch}_0[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_0$ 54 $\text{Batch}_1[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_1$ 55 ret (finished, \perp) <p>Agg($\hat{i} \in \mathbb{N}, \hat{j} \in \mathcal{V}$):</p> <ol style="list-style-type: none"> 56 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$: ret \perp 57 if $F(\text{Batch}_0[\hat{i}, \hat{j}, \cdot]) \neq F(\text{Batch}_1[\hat{i}, \hat{j}, \cdot])$ 58 and $(\forall j, j' \in \mathcal{V}) sk_j = sk_{j'}$: ret \perp 59 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{finished}$ 60 $\vec{y} \leftarrow \text{Out}[\hat{i}, \hat{j}, \cdot]$ 61 ret $\sum_{i=1}^{ \vec{y} } \vec{y}[i]$ <p>RO_i($seed, cntxt$):</p> <ol style="list-style-type: none"> 62 $l \leftarrow (jl, n, m, pl, ql)$ 63 if $\text{Rand}[i, seed, cntxt] = \perp$: 64 if $i \leq 5$: $\text{Rand}[i, seed, cntxt] \leftarrow_s \mathbb{F}^{l[i]}$ 65 else: $\text{Rand}[i, seed, cntxt] \leftarrow_s \{0, 1\}^\kappa$ 66 ret $\text{Rand}[i, seed, cntxt]$
---	---

Figure 18: Games $\underline{\mathcal{G}}_0$ and $\underline{\mathcal{G}}_1$ for the proof of Theorem 2. This game is identical to the privacy game for Π , except the **Shard**, **Prep**, and **Agg** algorithms have been inlined. Algorithm **Unpack** is as defined in Figure 5. The random oracles RO_i are lazy-evaluated in a table **Rand**.

Adversary $B^{\text{RO}}[A](\cdot)$:	Adversary $B^{\text{RO,Shard,Setup,Prep,Agg}}[A](st_B)$:
<ol style="list-style-type: none"> 1 $(st_A, \mathcal{V}, (sk_j)_{j \in \mathcal{V}}) \leftarrow_s A^{\text{RO}}(\cdot)$ 2 $z \leftarrow_s \mathcal{V}; \mathcal{V}' \leftarrow \{z\}; \mathcal{T} \leftarrow [s] \setminus \mathcal{V}$ 3 $st_B \leftarrow (st_A, z, \mathcal{T}, \mathcal{V}, (sk_j)_{j \in \mathcal{V}})$ 4 ret $(st_B, \mathcal{V}', (sk_z))$ 	<ol style="list-style-type: none"> 5 $(st_A, z, \mathcal{T}, \mathcal{V}, (sk_j)_{j \in \mathcal{V}}) \leftarrow st_B$ 6 $b' \leftarrow_s A^{\text{RO,ShardSim,SetupSim,PrepSim,AggSim}}(st_A)$ 7 ret b'

Figure 19: Wrapper adversary B for the proof of Theorem 2. Algorithms `SetupSim`, `PrepSim`, `AggSim` are evaluated by B just as the respective oracles in game $\underline{\mathsf{G}}_0$ except that queries pertaining to aggregator z are forwarded to B 's own oracles. Algorithm `ShardSim` forwards A 's query to `Shard` in the natural way, but returns the shares of the aggregators deemed honest by A .

table `Rand` with values chosen by the `Shard` oracle for the joint randomness, prover randomness, and query randomness. Accordingly, we pass these joint randomness and query randomness to the honest aggregator via its input share (`In`[k, z]; see line 29). This is to simplify bookkeeping in the next step.

Game $\underline{\mathsf{G}}_2$ is identical to game $\underline{\mathsf{G}}_1$ until programming `Rand` overwrites an already existing value on line 18, 19, 20, or 21.

- Line 18: Adversary B either has to guess $jseed$ or guess the input to `RO`₆ used to derive it. For the latter it must guess $rseed$ or all of the corresponding inputs to `RO`₇, which include the blinds generated by oracle `Shard`. Taking union bound over all the queries to `Shard`, the game overwrites `Rand` at this point with probability at most $q_1 q_{\text{Shard}}/2^\kappa + (q_6 + q_7)q_{\text{Shard}}/(2^{s \cdot \kappa})$.
- Line 19: B must guess the ps generated by oracle `Shard`, so the game overwrites the table with probability at most $q_4 q_{\text{Shard}}/2^\kappa$.
- Line 20: B must guess the nonce n generated by the oracle. The game overwrites the table here with probability at most $q_5 q_{\text{Shard}}/|\mathcal{N}|$.
- Line 21: B must guess $rseed$ or all of the corresponding inputs to `RO`₇, so the game overwrites the table with probability at most $(q_6 + q_7)q_{\text{Shard}}/(2^{s \cdot \kappa})$.

We bound the probability of B distinguishing between these games by the probability that any one of these events occurs, Gathering up the terms yields

$$\Pr[\underline{\mathsf{G}}_1(B)] \leq \Pr[\underline{\mathsf{G}}_2(B)] \tag{10}$$

$$+ \frac{(q_1 + q_4)q_{\text{Shard}}}{2^\kappa} + \frac{(q_6 + q_7)q_{\text{Shard}}}{2^{s \cdot \kappa - 1}} + \frac{q_5 q_{\text{Shard}}}{|\mathcal{N}|}. \tag{11}$$

In the next game hop (see $\underline{\mathsf{G}}_3$ in the left panel of Figure 21) we prepare to ensure that all of the input shares \vec{x} , proof shares $\vec{\pi}$, and the public share $rseed$ sampled by the `Shard` oracle are uniform random. We do so by sampling these values prior to processing m_b and programming the random oracle with the sample values (lines 3–12) so long as doing so does not overwrite existing values (see procedure `PO` on lines 37–40). The game sets a flag `bad`₄ if `Shard` would have overwritten an existing value. This does not change the adversary's view of the experiment, so

$$\Pr[\underline{\mathsf{G}}_2(B)] = \Pr[\underline{\mathsf{G}}_3(B)]. \tag{12}$$

Next, in game $\underline{\mathsf{G}}_4$ (top-right panel of Figure 21) we change oracle `Shard`'s behavior after `bad`₄ gets set. In particular, if ever `PO` is called on an input $(i, seed, cntxt, out)$ for which `Rand`[$i, seed, cntxt$], the value is overwritten. Game $\underline{\mathsf{G}}_4$ is identical to game $\underline{\mathsf{G}}_3$ until `bad`₄ gets set. Then we apply the Fundamental Lemma of Game Playing [14] to show that

$$\Pr[\underline{\mathsf{G}}_3(B)] \leq \Pr[\underline{\mathsf{G}}_4(B)] + \Pr[\underline{\mathsf{G}}_4(B) \text{ sets } \text{bad}_4] \tag{13}$$

$$\leq \Pr[\underline{\mathsf{G}}_4(B)] + \frac{((s-1)(q_2 + q_3) + s(q_7))q_{\text{Shard}}}{2^\kappa}. \tag{14}$$

<p>Game $\underline{\mathbb{G}}_1(B)$ $\underline{\mathbb{G}}_2(B)$:</p> <ol style="list-style-type: none"> 1 $(st_B, \{z\}, (sk_z,)) \leftarrow s B^{\text{RO}}(); \mathcal{T} \leftarrow [s] \setminus \{z\}$ 2 $b \leftarrow s \{0, 1\}; b' \leftarrow s B^{\text{RO, Shard, Setup, Prep, Agg}}(st_B)$ 3 $\text{ret } b = b'$ <p>Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):</p> <ol style="list-style-type: none"> 4 if $\text{Used}[\hat{k}] \neq \perp$: ret \perp 5 $n \leftarrow s \mathcal{N} \setminus \mathcal{N}^*; \mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 6 $\text{inp} \leftarrow \text{Encode}(m_b)$ 7 for $\hat{j} \in [2..s]$: 8 $\text{blind}_{\hat{j}}, \text{xseed}_{\hat{j}}, \text{pseed}_{\hat{j}} \leftarrow s \{0, 1\}^\kappa$ 9 $\vec{x}[\hat{j}] \leftarrow \text{RO}_2(\text{xseed}_{\hat{j}}, \hat{j})$ 10 $\text{rseed}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}])$ 11 $\vec{x}[1] \leftarrow \text{inp} - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ 12 $\text{blind}_1 \leftarrow s \{0, 1\}^\kappa; \text{ps} \leftarrow s \{0, 1\}^\kappa$ 13 $\vec{\text{rseed}}[1] \leftarrow \text{RO}_7(\text{blind}_1, 1 \parallel n \parallel \vec{x}[1])$ 14 $\text{jseed} \leftarrow \text{RO}_6(0^\kappa, \vec{\text{rseed}}); \text{jr} \leftarrow \text{RO}_1(\text{jseed}, \varepsilon)$ 15 $\text{pr} \leftarrow \text{RO}_4(\text{ps}, \varepsilon)$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 16 $\text{jseed} \leftarrow s \{0, 1\}^\kappa$ 17 $\text{jr} \leftarrow s \mathbb{F}^{jl}; \text{pr} \leftarrow s \mathbb{F}^{pl}; \text{qr} \leftarrow s \mathbb{F}^{ql}$ 18 $\text{Rand}[1, \text{jseed}, \varepsilon] \leftarrow \text{jr}$ 19 $\text{Rand}[4, \text{ps}, \varepsilon] \leftarrow \text{pr}$ 20 $\text{Rand}[5, sk_z, n] \leftarrow \text{qr}$ 21 $\text{Rand}[6, 0^\kappa, \vec{\text{rseed}}] \leftarrow \text{jseed}$ </div> <ol style="list-style-type: none"> 22 $\vec{\pi}[1] \leftarrow \text{Prove}(\text{inp}, \text{jr}; \text{pr})$ 23 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{\hat{j}=2}^s \text{RO}_3(\text{pseed}_{\hat{j}}, \hat{j})$ 24 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], \text{blind}_1)$ 25 for $\hat{j} \in [2..s]$: 26 $\vec{x}[\hat{j}] \leftarrow (\text{xseed}_{\hat{j}}, \text{pseed}_{\hat{j}}, \text{blind}_{\hat{j}})$ 27 $\text{Pub}[\hat{k}] \leftarrow \text{rseed}$ 28 $\text{In}[\hat{k}, \cdot] \leftarrow \vec{x}$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 29 $\text{In}[\hat{k}, z] \leftarrow (\vec{x}[z], \text{jseed}, \text{jr}, \text{qr})$ </div> <ol style="list-style-type: none"> 30 $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 31 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}})$ <p>Setup($\hat{i} \in \mathbb{N}, \hat{j} \in \{z\}, st_{\text{init}} \in \{\varepsilon\}$):</p> <ol style="list-style-type: none"> 32 if $\text{Status}[\hat{i}, \hat{j}] \neq \perp$ or $\text{Setup}[\cdot, \hat{j}] > 0$: ret \perp 33 $\text{Setup}[\hat{i}, \hat{j}] \leftarrow st_{\text{init}}$ 34 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{running}$ 	<p>Prep($\hat{i} \in \mathbb{N}, \hat{j} \in \{z\}, \hat{k} \in \mathbb{N}, \text{msg} \in \mathcal{M}^*$):</p> <ol style="list-style-type: none"> 35 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$ or $\text{In}[\hat{k}, \hat{j}] = \perp$: 36 ret \perp 37 if $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \perp$: 38 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \text{Setup}[\hat{i}, \hat{j}]$ 39 $\text{msg} \leftarrow (\text{Pub}[\hat{k}],)$ 40 $(n, m_0, m_1) \leftarrow \text{Used}[\hat{k}]$ 41 if $\text{St}[\hat{i}, \hat{j}, \hat{k}] = \varepsilon$: // Process initial message from client <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 42 $(\text{inp}, \pi, \text{blind}) \leftarrow \text{Unpack}(\hat{j}, \text{In}[\hat{k}, \hat{j}])$ </div> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 43 $(x, \text{jseed}, \text{jr}, \text{qr}) \leftarrow \text{In}[\hat{k}, \hat{j}]$ 44 $(\text{inp}, \pi, \text{blind}) \leftarrow \text{Unpack}(\hat{j}, x)$ </div> <ol style="list-style-type: none"> 45 $(\vec{\text{rseed}},) \leftarrow \text{msg}$ 46 $\vec{\text{rseed}}[\hat{j}] \leftarrow \text{RO}_7(\text{blind}, \hat{j} \parallel n \parallel \text{inp})$ <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <ol style="list-style-type: none"> 47 $\text{jseed} \leftarrow \text{RO}_6(0^\kappa, \vec{\text{rseed}}); \text{jr} \leftarrow \text{RO}_1(\text{jseed}, \varepsilon)$ 48 $\text{qr} \leftarrow \text{RO}_5(sk_z, n)$ </div> <ol style="list-style-type: none"> 49 $\text{msg} \leftarrow (\text{Query}(\text{inp}, \pi, \text{jr}; \text{qr}), \vec{\text{rseed}}[\hat{j}])$ 50 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow (\text{jseed}, \text{Truncate}(\text{inp}))$ 51 ret $(\text{running}, \text{msg})$ 52 // Process broadcast messages from aggregators 53 $(\text{jseed}, y) \leftarrow \text{St}[\hat{i}, \hat{j}, \hat{k}]$ 54 $(\vec{\text{vfs}}[\hat{j}], \vec{\text{rseed}}[\hat{j}])_{\hat{j} \in [s]} \leftarrow \text{msg}$ 55 $\text{acc} \leftarrow \text{Decide}(\sum_{\hat{j}=1}^s \vec{\text{vfs}}[\hat{j}])$ 56 $\text{St}[\hat{i}, \hat{j}, \hat{k}] \leftarrow \perp$ 57 if $\text{acc} = 0$ or $\text{jseed} \neq \text{RO}_6(0^\kappa, \vec{\text{rseed}})$: 58 ret (failed, \perp) 59 $\text{Out}[\hat{i}, \hat{j}, \hat{k}] \leftarrow y$ 60 $\text{Batch}_0[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_0$ 61 $\text{Batch}_1[\hat{i}, \hat{j}, \hat{k}] \leftarrow m_1$ 62 ret $(\text{finished}, \perp)$ <p>Agg($\hat{i} \in \mathbb{N}, \hat{j} \in \{z\}$):</p> <ol style="list-style-type: none"> 63 if $\text{Status}[\hat{i}, \hat{j}] \neq \text{running}$: ret \perp 64 if $F(\text{Batch}_0[\hat{i}, \hat{j}, \cdot]) \neq F(\text{Batch}_1[\hat{i}, \hat{j}, \cdot])$: ret \perp 65 $\text{Status}[\hat{i}, \hat{j}] \leftarrow \text{finished}$ 66 $\vec{y} \leftarrow \text{Out}[\hat{i}, \hat{j}, \cdot]$ 67 ret $\sum_{i=1}^{ \vec{y} } \vec{y}[i]$ <p>RO_i($\text{seed}, \text{cntxt}$):</p> <ol style="list-style-type: none"> 68 $l \leftarrow (jl, n, m, pl, ql)$ 69 if $\text{Rand}[i, \text{seed}, \text{cntxt}] = \perp$: 70 if $i \leq 5$: $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow s \mathbb{F}^{l[i]}$ 71 else: $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow s \{0, 1\}^\kappa$ 72 ret $\text{Rand}[i, \text{seed}, \text{cntxt}]$
--	--

Figure 20: Game $\underline{\mathbb{G}}_2$ for the proof of Theorem 2.

The probability that B sets the bad_4 flag in Game $\underline{\mathbf{G}}_4$ is the probability that B makes a random oracle query that gets overwritten on line 9, 10, 11, or 12. On each line, the random oracle is programmed with a uniform random string sampled by the oracle prior to being revealed to the adversary. Rolling out the for-loop on line 8 and taking a union bound over all Shard queries yields the claimed bound.

Next, in game $\underline{\mathbf{G}}_5$ (bottom-right panel of Figure 21) we simplify the Shard oracle by inlining calls to PO and replacing invocations of RO with corresponding value generated by the oracle. These changes do not change the view of the adversary, so

$$\Pr[\underline{\mathbf{G}}_4(B)] = \Pr[\underline{\mathbf{G}}_5(B)]. \quad (15)$$

Up to this point, we have constructed the “leader” input and proof shares differently than all of the other shares: we pick all other shares randomly, then set $\vec{x}[1] = \text{inp} - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ and $\vec{\pi}[1] = \pi - \sum_{\hat{j}=2}^s \vec{\pi}[\hat{j}]$. In our next game, we instead sample the “leader” shares randomly and compute the shares of the honest aggregator z in a distinguished manner: $\vec{x}[z] = \text{inp} - \sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}]$ and $\vec{\pi}[z] = \pi - \sum_{\hat{j} \in \mathcal{T}} \vec{\pi}[\hat{j}]$, where $\mathcal{T} = [s] \setminus \{z\}$. If $z = 1$, this changes nothing. Otherwise, consider that in $\underline{\mathbf{G}}_5$, we have

$$\vec{x}[1] = \text{inp} - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}] = \text{inp} - \left(\sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}] - \vec{x}[z] + \vec{x}[1] \right).$$

If we add $\vec{x}[z] - \vec{x}[1]$ to both sides of this equation, we can see that in $\underline{\mathbf{G}}_4$, it was already true that $\vec{x}[z] = \text{inp} - \sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}]$. The same holds true for $\vec{\pi}[z]$ by an analogous argument. Therefore the distributions of aggregators’ 1 and z ’s input and proof shares are unchanged between $\underline{\mathbf{G}}_4$ and $\underline{\mathbf{G}}_5$, and we have

$$\Pr[\underline{\mathbf{G}}_5(B)] = \Pr[\underline{\mathbf{G}}_6(B)]. \quad (16)$$

In our next game ($\underline{\mathbf{G}}_7$, defined in the right panel of Figure 22) we run the query algorithm for aggregator z in the Shard oracle and only send the result to Prep. The adversary cannot detect the timing of when this algorithm is run, so we have

$$\Pr[\underline{\mathbf{G}}_6(B)] = \Pr[\underline{\mathbf{G}}_7(B)]. \quad (17)$$

In the next game ($\underline{\mathbf{G}}_8$, defined in the left panel of Figure 23) we run View_{FLP} (as defined in Section 2) on input inp to get jr , qr , and a verifier σ and use these to compute Shard’s output. We have defined $\vec{x}[z] = \text{inp} - \sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}]$ and $\vec{\pi}[z] = \pi - \sum_{\hat{j} \in \mathcal{T}} \vec{\pi}[\hat{j}]$. Using the full linearity of FLP, we can the honest aggregator z ’s verifier share vfs in terms of jr , qr , σ and the corrupt aggregators’ shares, since:

$$\begin{aligned} \text{Query}(\text{inp}, \pi, jr; qr) &= \text{Query}(\vec{x}[z], \vec{\pi}[z], jr; qr) & (18) \\ &+ \sum_{\hat{j} \in \mathcal{T}} \text{Query}(\vec{x}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr) & (19) \end{aligned}$$

This revision to the game does not change the outcome of the experiment. However, since we do not have access to the prover randomness generated by $\text{View}_{\text{FLP}}(\text{inp})$, we can no longer consistently program the random oracle (see 19). Fortunately, to trigger this inconsistency, the adversary would have to guess the seed ps used to to program it prior to calling Shard. It follows that

$$\Pr[\underline{\mathbf{G}}_7(B)] \leq \Pr[\underline{\mathbf{G}}_8(B)] + \frac{q_4 q_{\text{Shard}}}{2^\kappa}. \quad (20)$$

Let S be the simulator hypothesized by δ -privacy of FLP. In the right panel of Figure 23 we define a series of hybrid games that replace View_{FLP} with a simulator S for the privacy of FLP. Recall from Section 2 that S outputs a string $jr \parallel qr \parallel \sigma$. In $\underline{\mathbf{G}}_9^i(B)$, the first $i - 1$ queries to Shard generate jr , qr , σ by calling $S()$; the remaining queries call View_{FLP} instead. This means that $\underline{\mathbf{G}}_9^1$ is identical to $\underline{\mathbf{G}}_8$, so

$$\Pr[\underline{\mathbf{G}}_8(B)] = \Pr[\underline{\mathbf{G}}_9^1(B)]. \quad (21)$$

For every $v \in \mathbb{F}^{jl \times ql \times v}$, we let $p_{i,v}$ denote the probability that B wins hybrid $\underline{\mathbf{G}}_9^i$, conditioned on the event that the i^{th} query to Shard sets $v = jr \parallel qr \parallel \sigma$. A union bound over all v shows that

$$\Pr[\underline{\mathbf{G}}_9^i(B)] = \sum_{v \in \mathbb{F}^{jl \times ql \times v}} p_{i,v}. \quad (22)$$

<p>Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$): Game $\underline{\mathbf{G}}_2$ $\underline{\mathbf{G}}_3$</p> <pre style="font-family: monospace; font-size: 0.9em;"> 1 if Used[\hat{k}] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 ($blind_1, \dots, blind_s$) $\leftarrow \mathbb{F}^{\kappa}$ 5 ($xseed_2, \dots, xseed_s$) $\leftarrow \mathbb{F}^{\kappa}$ 6 ($pseed_2, \dots, pseed_s$) $\leftarrow \mathbb{F}^{\kappa}$ 7 $rseed \leftarrow \mathbb{F}^{\kappa}$ 8 for $\hat{j} \in [s]$: 9 $\text{PO}_2(xseed_{\hat{j}}, \hat{j}, \vec{x}[\hat{j}])$ 10 $\text{PO}_3(pseed_{\hat{j}}, \hat{j}, \vec{\pi}[\hat{j}])$ 11 $\text{PO}_7(blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}], rseed[\hat{j}])$ 12 $\text{PO}_7(blind_1, 1 \parallel n \parallel \vec{x}[1], rseed[1])$ 13 $inp \leftarrow \text{Encode}(m_b)$ 14 for $\hat{j} \in [2..s]$: 15 $blind_{\hat{j}}, xseed_{\hat{j}}, pseed_{\hat{j}} \leftarrow \mathbb{F}^{\kappa}$ 16 $\vec{x}[\hat{j}] \leftarrow \text{RO}_2(xseed_{\hat{j}}, \hat{j})$ 17 $rseed[\hat{j}] \leftarrow \text{RO}_7(blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}])$ 18 $\vec{x}[1] \leftarrow inp - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ 19 $blind_1 \leftarrow \mathbb{F}^{\kappa}$; $ps \leftarrow \mathbb{F}^{\kappa}$ 20 $rseed[1] \leftarrow \text{RO}_7(blind_1, 1 \parallel n \parallel \vec{x}[1])$ 21 $jseed \leftarrow \mathbb{F}^{\kappa}$ 22 $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$; $qr \leftarrow \mathbb{F}^{ql}$ 23 $\text{Rand}[1, jseed, \varepsilon] \leftarrow jr$ 24 $\text{Rand}[4, ps, \varepsilon] \leftarrow pr$ 25 $\text{Rand}[5, sk_z, n] \leftarrow qr$ 26 $\text{Rand}[6, 0^{\kappa}, rseed] \leftarrow jseed$ 27 $\vec{\pi}[1] \leftarrow \text{Prove}(inp, jr; pr)$ 28 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{\hat{j}=2}^s \text{RO}_3(pseed_{\hat{j}}, \hat{j})$ 29 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 30 for $\hat{j} \in [2..s]$: 31 $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 32 $\text{Pub}[\hat{k}] \leftarrow rseed$ 33 $\text{In}[\hat{k}, \cdot] \leftarrow \vec{x}$ 34 $\text{In}[\hat{k}, z] \leftarrow (\vec{x}[z], jseed, jr, qr)$ 35 $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 36 ret ($n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}}$)</pre>	<p>Algorithm $\text{PO}_i(\text{seed}, \text{cntxt}, \text{out})$: Game $\underline{\mathbf{G}}_3$ $\underline{\mathbf{G}}_4$</p> <pre style="font-family: monospace; font-size: 0.9em;"> 1 if $\text{Rand}[i, \text{seed}, \text{cntxt}] = \perp$: 2 $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow \text{out}$ 3 else: 4 bad₄ $\leftarrow \text{true}$; $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow \text{out}$</pre> <hr/> <p>Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$): Game $\underline{\mathbf{G}}_5$</p> <pre style="font-family: monospace; font-size: 0.9em;"> 1 if Used[\hat{k}] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 ($blind_1, \dots, blind_s$) $\leftarrow \mathbb{F}^{\kappa}$ 5 ($xseed_2, \dots, xseed_s$) $\leftarrow \mathbb{F}^{\kappa}$ 6 ($pseed_2, \dots, pseed_s$) $\leftarrow \mathbb{F}^{\kappa}$ 7 $rseed \leftarrow \mathbb{F}^{\kappa}$ 8 for $\hat{j} \in [s]$: 9 $\text{Rand}[2, xseed_{\hat{j}}, \hat{j}] \leftarrow \vec{x}[\hat{j}]$ 10 $\text{Rand}[3, pseed_{\hat{j}}, \hat{j}] \leftarrow \vec{\pi}[\hat{j}]$ 11 $\text{Rand}[7, blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}]] \leftarrow rseed[\hat{j}]$ 12 $\text{Rand}[7, blind_1, 1 \parallel n \parallel \vec{x}[1]] \leftarrow rseed[1]$ 13 $inp \leftarrow \text{Encode}(m_b)$ 14 $\vec{x}[1] \leftarrow inp - \sum_{\hat{j}=2}^s \vec{x}[\hat{j}]$ 15 $ps \leftarrow \mathbb{F}^{\kappa}$ 16 $jseed \leftarrow \mathbb{F}^{\kappa}$ 17 $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$; $qr \leftarrow \mathbb{F}^{ql}$ 18 $\text{Rand}[1, jseed, \varepsilon] \leftarrow jr$ 19 $\text{Rand}[4, ps, \varepsilon] \leftarrow pr$ 20 $\text{Rand}[5, sk_z, n] \leftarrow qr$ 21 $\text{Rand}[6, 0^{\kappa}, rseed] \leftarrow jseed$ 22 $\vec{\pi}[1] \leftarrow \text{Prove}(inp, jr; pr)$ 23 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{\hat{j}=2}^s \text{RO}_3(pseed_{\hat{j}}, \hat{j})$ 24 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 25 for $\hat{j} \in [2..s]$: 26 $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 27 $\text{Pub}[\hat{k}] \leftarrow rseed$ 28 $\text{In}[\hat{k}, \cdot] \leftarrow \vec{x}$ 29 $\text{In}[\hat{k}, z] \leftarrow (\vec{x}[z], jseed, jr, qr)$ 30 $\text{Used}[\hat{k}] \leftarrow (n, m_0, m_1)$ 31 ret ($n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}}$)</pre>
---	---

Figure 21: Games $\underline{\mathbf{G}}_3$ (left), $\underline{\mathbf{G}}_4$ (top-right), and $\underline{\mathbf{G}}_5$ (bottom-right) for the proof of Theorem 2. Only the **Shard** is shown, as this is the only object that changes in each game hop.

Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):	Game \underline{G}_5 \underline{G}_6	Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):	Game \underline{G}_6 \underline{G}_7
<pre> 1 if Used[\hat{k}] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 $(blind_1, \dots, blind_s) \leftarrow \{0, 1\}^\kappa$ 5 $(xseed_2, \dots, xseed_s) \leftarrow \{0, 1\}^{\kappa s-1}$ 6 $(pseed_2, \dots, pseed_s) \leftarrow \{0, 1\}^{\kappa s-1}$ 7 $rseed \leftarrow \{0, 1\}^\kappa$ 8 $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$ 9 $inp \leftarrow \text{Encode}(m_b)$ 10 $\pi \leftarrow \text{Prove}(inp, jr; pr)$ 11 $\vec{x}[z] \leftarrow inp - \sum_{j \in \mathcal{T}} \vec{x}[\hat{j}]$ 12 $\vec{\pi}[z] \leftarrow \pi - \sum_{j \in \mathcal{T}} \vec{\pi}[\hat{j}]$ 13 for $\hat{j} \in [s]$: 14 Rand[2, $xseed_{\hat{j}}, \hat{j}$] $\leftarrow \vec{x}[\hat{j}]$ 15 Rand[3, $pseed_{\hat{j}}, \hat{j}$] $\leftarrow \vec{\pi}[\hat{j}]$ 16 Rand[7, $blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}]$] $\leftarrow \vec{rseed}[\hat{j}]$ 17 Rand[7, $blind_1, 1 \parallel n \parallel \vec{x}[1]$] $\leftarrow \vec{rseed}[1]$ 18 $inp \leftarrow \text{Encode}(m_b)$ 19 $\vec{x}[1] \leftarrow inp - \sum_{j=2}^s \vec{x}[\hat{j}]$ 20 $ps \leftarrow \{0, 1\}^\kappa$ 21 $jseed \leftarrow \{0, 1\}^\kappa$ 22 $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$; $qr \leftarrow \mathbb{F}^{ql}$ 23 Rand[1, $jseed, \varepsilon$] $\leftarrow jr$ 24 Rand[4, ps, ε] $\leftarrow pr$ 25 Rand[5, sk_z, n] $\leftarrow qr$ 26 Rand[6, $0^\kappa, rseed$] $\leftarrow jseed$ 27 $\vec{\pi}[1] \leftarrow \text{Prove}(inp, jr; pr)$ 28 $\vec{\pi}[1] \leftarrow \vec{\pi}[1] - \sum_{j=2}^s \vec{\pi}[\hat{j}]$ 29 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 30 for $\hat{j} \in [2..s]$: 31 $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 32 Pub[\hat{k}] $\leftarrow rseed$ 33 In[\hat{k}, \cdot] $\leftarrow \vec{x}$ 34 In[\hat{k}, z] $\leftarrow (\vec{x}[z], jseed, jr, qr)$ 35 Used[\hat{k}] $\leftarrow (n, m_0, m_1)$ 36 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{j \in \mathcal{T}})$ </pre>	<div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 0 auto;"> \underline{G}_5 </div> <div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 0 auto;"> \underline{G}_6 </div>	<pre> 1 if Used[\hat{k}] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 $(blind_1, \dots, blind_s) \leftarrow \{0, 1\}^\kappa$ 5 $(xseed_2, \dots, xseed_s) \leftarrow \{0, 1\}^{\kappa s-1}$ 6 $(pseed_2, \dots, pseed_s) \leftarrow \{0, 1\}^{\kappa s-1}$ 7 $\vec{rseed} \leftarrow \{0, 1\}^\kappa$; $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$ 8 $inp \leftarrow \text{Encode}(m_b)$ 9 $\pi \leftarrow \text{Prove}(inp, jr; pr)$ 10 $\vec{x}[z] \leftarrow inp - \sum_{j \in \mathcal{T}} \vec{x}[\hat{j}]$ 11 $\vec{\pi}[z] \leftarrow \pi - \sum_{j \in \mathcal{T}} \vec{\pi}[\hat{j}]$ 12 for $\hat{j} \in [s]$: 13 Rand[2, $xseed_{\hat{j}}, \hat{j}$] $\leftarrow \vec{x}[\hat{j}]$ 14 Rand[3, $pseed_{\hat{j}}, \hat{j}$] $\leftarrow \vec{\pi}[\hat{j}]$ 15 Rand[7, $blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}]$] $\leftarrow \vec{rseed}[\hat{j}]$ 16 Rand[7, $blind_1, 1 \parallel n \parallel \vec{x}[1]$] $\leftarrow \vec{rseed}[1]$ 17 $ps \leftarrow \{0, 1\}^\kappa$; $jseed \leftarrow \{0, 1\}^\kappa$; $qr \leftarrow \mathbb{F}^{ql}$ 18 Rand[1, $jseed, \varepsilon$] $\leftarrow jr$; Rand[4, ps, ε] $\leftarrow pr$ 19 Rand[5, sk_z, n] $\leftarrow qr$; Rand[6, $0^\kappa, rseed$] $\leftarrow jseed$ 20 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 21 for $\hat{j} \in [2..s]$: $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 22 Pub[\hat{k}] $\leftarrow \vec{rseed}$; In[\hat{k}, \cdot] $\leftarrow \vec{x}$ 23 In[\hat{k}, z] $\leftarrow (\vec{x}[z], jseed, jr, qr)$ 24 $vfs \leftarrow \text{Query}(\vec{x}[z], \vec{\pi}[z], jr; qr)$ 25 In[\hat{k}, \hat{j}] $\leftarrow (vfs, \vec{rseed}[z], \text{Truncate}(\vec{x}[z]), jseed)$ 26 Used[\hat{k}] $\leftarrow (n, m_0, m_1)$ 27 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{j \in \mathcal{T}})$ </pre>	<div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 0 auto;"> \underline{G}_6 </div> <div style="border: 1px solid black; padding: 2px; width: fit-content; margin: 0 auto;"> \underline{G}_7 </div>
<pre> 28 Prep($i \in \mathbb{N}, \hat{j} \in \{z\}, \hat{k} \in \mathbb{N}, m\vec{s}g \in \mathcal{M}^*$): 29 if Status[$i, \hat{j}$] $\neq \text{running}$ or In[\hat{k}, \hat{j}] = \perp: ret \perp 30 if St[i, \hat{j}, \hat{k}] = \perp: 31 St[i, \hat{j}, \hat{k}] $\leftarrow \text{Setup}[i, \hat{j}]; m\vec{s}g \leftarrow (\text{Pub}[\hat{k}],)$ 32 $(n, m_0, m_1) \leftarrow \text{Used}[\hat{k}]$ 33 if St[i, \hat{j}, \hat{k}] = ε: // Process initial message from client 34 $(x, jseed, jr, qr) \leftarrow \text{In}[\hat{k}, \hat{j}]$ 35 $(inp, \pi, blind) \leftarrow \text{Unpack}(\hat{j}, x); (\vec{rseed},) \leftarrow m\vec{s}g$ 36 $\vec{rseed}[\hat{j}] \leftarrow \text{RO}_7(blind, \hat{j} \parallel n \parallel inp)$ 37 $msg \leftarrow (\text{Query}(inp, \pi, jr; qr), \vec{rseed}[\hat{j}])$ 38 St[i, \hat{j}, \hat{k}] $\leftarrow (jseed, \text{Truncate}(inp))$ 39 $(vfs, rseed, y, jseed) \leftarrow \text{In}[\hat{k}, \hat{j}]$ 40 $msg \leftarrow (vfs, rseed); \text{St}[i, \hat{j}, \hat{k}] \leftarrow (jseed, y)$ 41 ret (running, msg) 42 // Process broadcast messages from aggregators 43 $(jseed, y) \leftarrow \text{St}[i, \hat{j}, \hat{k}]; (v\vec{f}s[\hat{j}], \vec{rseed}[\hat{j}])_{j \in [s]} \leftarrow m\vec{s}g$ 44 $acc \leftarrow \text{Decide}(\sum_{j=1}^s v\vec{f}s[\hat{j}]); \text{St}[i, \hat{j}, \hat{k}] \leftarrow \perp$ 45 if $acc = 0$ or $jseed \neq \text{RO}_6(0^\kappa, \vec{rseed})$: ret (failed, \perp) 46 Out[i, \hat{j}, \hat{k}] $\leftarrow y$ 47 Batch0[i, \hat{j}, \hat{k}] $\leftarrow m_0$; Batch1[i, \hat{j}, \hat{k}] $\leftarrow m_1$ 48 ret (finished, \perp) </pre>		<pre> 40 ret (running, msg) 41 // Process broadcast messages from aggregators 42 $(jseed, y) \leftarrow \text{St}[i, \hat{j}, \hat{k}]; (v\vec{f}s[\hat{j}], \vec{rseed}[\hat{j}])_{j \in [s]} \leftarrow m\vec{s}g$ 43 $acc \leftarrow \text{Decide}(\sum_{j=1}^s v\vec{f}s[\hat{j}]); \text{St}[i, \hat{j}, \hat{k}] \leftarrow \perp$ 44 if $acc = 0$ or $jseed \neq \text{RO}_6(0^\kappa, \vec{rseed})$: ret (failed, \perp) 45 Out[i, \hat{j}, \hat{k}] $\leftarrow y$ 46 Batch0[i, \hat{j}, \hat{k}] $\leftarrow m_0$; Batch1[i, \hat{j}, \hat{k}] $\leftarrow m_1$ 47 ret (finished, \perp) </pre>	

Figure 22: Game \underline{G}_6 (left) and game \underline{G}_7 (right) for the proof of Theorem 2.

Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):	Game $\underline{\mathbb{G}}_7$ $\underline{\mathbb{G}}_8$	Shard($\hat{k} \in \mathbb{N}, m_0, m_1 \in \mathcal{I}$):	Game $\underline{\mathbb{G}}_8$ $\underline{\mathbb{G}}_9^i$
<ol style="list-style-type: none"> 1 if Used$[\hat{k}] \neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 $(blind_1, \dots, blind_s) \leftarrow \mathbb{F}^s$ 5 $(xseed_2, \dots, xseed_s) \leftarrow \mathbb{F}^{s-1}$ 6 $(pseed_2, \dots, pseed_s) \leftarrow \mathbb{F}^{s-1}$ 7 $rseed \leftarrow \mathbb{F}^s$; $jr \leftarrow \mathbb{F}^{jl}$; $pr \leftarrow \mathbb{F}^{pl}$ 8 $inp \leftarrow \text{Encode}(m_b)$ 	<ol style="list-style-type: none"> 9 $\pi \leftarrow \text{Prove}(inp, jr; pr)$ 	<ol style="list-style-type: none"> 1 if Used$[\hat{k}] \neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 $\vec{x} \leftarrow \mathbb{F}^n$; $\vec{\pi} \leftarrow \mathbb{F}^m$ 4 $(blind_1, \dots, blind_s) \leftarrow \mathbb{F}^s$ 5 $(xseed_2, \dots, xseed_s) \leftarrow \mathbb{F}^{s-1}$ 6 $(pseed_2, \dots, pseed_s) \leftarrow \mathbb{F}^{s-1}$ 7 $rseed \leftarrow \mathbb{F}^s$ 8 $inp \leftarrow \text{Encode}(m_b)$ 	<ol style="list-style-type: none"> 9 $jr \parallel \parallel qr \parallel \sigma \leftarrow \text{View}_{\text{FLP}}(inp)$
<ol style="list-style-type: none"> 10 $jr \parallel \parallel qr \parallel \sigma \leftarrow \text{View}_{\text{FLP}}(inp)$ 		<ol style="list-style-type: none"> 10 $ctr \leftarrow ctr + 1$ 11 if $ctr < i$: $jr \parallel \parallel qr \parallel \sigma \leftarrow S()$ 12 else: $jr \parallel \parallel qr \parallel \sigma \leftarrow \text{View}_{\text{FLP}}(inp)$ 	
<ol style="list-style-type: none"> 11 $\vec{x}[z] \leftarrow inp - \sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}]$ 12 $\vec{\pi}[z] \leftarrow \pi - \sum_{\hat{j} \in \mathcal{T}} \vec{\pi}[\hat{j}]$ 		<ol style="list-style-type: none"> 13 $\vec{x}[z] \leftarrow inp - \sum_{\hat{j} \in \mathcal{T}} \vec{x}[\hat{j}]$ 	
<ol style="list-style-type: none"> 13 for $\hat{j} \in [s]$: 14 Rand[2, $xseed_{\hat{j}}, \hat{j}] \leftarrow \vec{x}[\hat{j}]$ 15 Rand[3, $pseed_{\hat{j}}, \hat{j}] \leftarrow \vec{\pi}[\hat{j}]$ 16 Rand[7, $blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}]] \leftarrow rseed[\hat{j}]$ 17 Rand[7, $blind_1, 1 \parallel n \parallel \vec{x}[1]] \leftarrow rseed[1]$ 18 $ps \leftarrow \mathbb{F}^{\kappa}$; $jseed \leftarrow \mathbb{F}^{\kappa}$; $qr \leftarrow \mathbb{F}^{ql}$ 19 Rand[1, $jseed, \varepsilon] \leftarrow jr$; Rand[4, $ps, \varepsilon] \leftarrow pr$ 20 Rand[5, $sk_z, n] \leftarrow qr$; Rand[6, $0^\kappa, rseed] \leftarrow jseed$ 21 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 22 for $\hat{j} \in [2..s]$: $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 23 Pub$[\hat{k}] \leftarrow rseed$; In$[\hat{k}, \cdot] \leftarrow \vec{x}$ 24 $vfs \leftarrow \text{Query}(\vec{x}[z], \vec{\pi}[z], jr; qr)$ 		<ol style="list-style-type: none"> 14 for $\hat{j} \in [s]$: 15 Rand[2, $xseed_{\hat{j}}, \hat{j}] \leftarrow \vec{x}[\hat{j}]$ 16 Rand[3, $pseed_{\hat{j}}, \hat{j}] \leftarrow \vec{\pi}[\hat{j}]$ 17 Rand[7, $blind_{\hat{j}}, \hat{j} \parallel n \parallel \vec{x}[\hat{j}]] \leftarrow rseed[\hat{j}]$ 18 Rand[7, $blind_1, 1 \parallel n \parallel \vec{x}[1]] \leftarrow rseed[1]$ 19 $jseed \leftarrow \mathbb{F}^{\kappa}$ 20 Rand[1, $jseed, \varepsilon] \leftarrow jr$ 21 Rand[5, $sk_z, n] \leftarrow qr$; Rand[6, $0^\kappa, rseed] \leftarrow jseed$ 22 $\vec{x}[1] \leftarrow (\vec{x}[1], \vec{\pi}[1], blind_1)$ 23 for $\hat{j} \in [2..s]$: $\vec{x}[\hat{j}] \leftarrow (xseed_{\hat{j}}, pseed_{\hat{j}}, blind_{\hat{j}})$ 24 Pub$[\hat{k}] \leftarrow rseed$; In$[\hat{k}, \cdot] \leftarrow \vec{x}$ 25 $vfs \leftarrow \sigma -$ 26 $\sum_{\hat{j} \in \mathcal{T}} \text{Query}(\vec{x}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr)$ 27 In$[\hat{k}, \hat{j}] \leftarrow (vfs, rseed[z], \text{Truncate}(\vec{x}[z]), jseed)$ 28 Used$[\hat{k}] \leftarrow (n, m_0, m_1)$ 29 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}})$ 	
<ol style="list-style-type: none"> 25 $vfs \leftarrow \sigma -$ 26 $\sum_{\hat{j} \in \mathcal{T}} \text{Query}(\vec{x}[\hat{j}], \vec{\pi}[\hat{j}], jr; qr)$ 			
<ol style="list-style-type: none"> 27 In$[\hat{k}, \hat{j}] \leftarrow (vfs, rseed[z], \text{Truncate}(\vec{x}[z]), jseed)$ 28 Used$[\hat{k}] \leftarrow (n, m_0, m_1)$ 29 ret $(n, \text{Pub}[\hat{k}], (\text{In}[\hat{k}, \hat{j}])_{\hat{j} \in \mathcal{T}})$ 			

Figure 23: Game $\underline{\mathbb{G}}_8$ (left) and game $\underline{\mathbb{G}}_9$ for the proof of Theorem 2.

We are now ready to bound the quantity $\Pr[\underline{\mathbb{G}}_9^{i+1}(B)] - \Pr[\underline{\mathbb{G}}_9^i(B)]$. The two games $\underline{\mathbb{G}}_9^{i+1}$ and $\underline{\mathbb{G}}_9^i$ differ only in the tuple v chosen by of the $(i+1)^{\text{th}}$ query to **Shard**: the former calls View_{FLP} and the latter calls S . We therefore decompose both probabilities over the possible choices of v , and substitute in the statement

$$\sum_{v \in \mathbb{F}^{jl \times ql \times v}} |\Pr[\text{View}_{\text{FLP}}(inp) = v] - \Pr[S() = v]| \leq \delta$$

that follows from the δ -privacy of FLP for all inp . Since $p_{i,v} \leq 1$ for all i and v and $\Pr[\text{View}_{\text{FLP}}(inp) =$

$$v] - \Pr[S() = v] \leq |\Pr[\text{View}_{\text{FLP}}(\text{inp}) = v] - \Pr[S() = v]|:$$

$$\begin{aligned} \Pr[\underline{\mathbf{G}}_9^{i+1}(B)] - \Pr[\underline{\mathbf{G}}_9^i(B)] &= \\ &= \sum_v p_{i,v} \cdot \Pr[\text{View}_{\text{FLP}}(\text{inp}) = v] - p_{i,v} \cdot \Pr[S() = v] \\ &= \sum_v p_{i,v} \cdot (\Pr[\text{View}_{\text{FLP}}(\text{inp}) = v] - \Pr[S() = v]) \\ &\leq \sum_v |\Pr[\text{View}_{\text{FLP}}(\text{inp}) = v] - \Pr[S() = v]| \\ &\leq \delta. \end{aligned}$$

A union bound over all $i \in [q_{\text{Shard}}]$ produces the final inequality:

$$\Pr[\underline{\mathbf{G}}_9^{q_{\text{Shard}}}(B)] - \Pr[\underline{\mathbf{G}}_9^1(B)] \leq \delta \cdot q_{\text{Shard}}. \quad (23)$$

Finally, we observe that game $\underline{\mathbf{G}}_9^{q_{\text{Shard}}}$ can now be rewritten so that the outcome is independent of the challenge bit b . Hence

$$\Pr[\underline{\mathbf{G}}_9^{q_{\text{Shard}}}(B)] = \frac{1}{2}. \quad (24)$$

Collecting bounds across all games and simplifying yields the theorem. \square

C.3 Doplar Robustness (Theorem 3)

The proof is by a game-playing argument. We begin with the game $\underline{\mathbf{G}}_0$ defined in Figure 24 played by the given adversary A . This game was constructed from $\text{Exp}_{\Pi}^{\text{robust}}(A)$ by applying the following revisions. First, we have replaced **Prep** with its implementation, rolled out the loops in the **Prep** oracle, and simplified some of the control flow. Second, we have removed the call to **refineFromShares** and set the purported refined measurement with the sum of the refined shares output by the calls to **VIDPF.VEval**. (This is equivalent by refinement consistency of Π .) Third, we use the fact that the allowed-state **validSt** algorithm for Π only permits **Prep** queries with unique (n, ℓ) pairs to make the contents of table **Used** more explicit. Finally, we lazy-evaluate each random oracle, denoted $\underline{\mathbf{RO}}_i$, with a table **Rand**. We use $\underline{\mathbf{RO}}'_i$ to denote the random oracle for **VIDPF**. By construction we have that

$$\text{Adv}_{\Pi}^{\text{robust}}(A) = \Pr[\underline{\mathbf{G}}_0(A)]. \quad (25)$$

In the remainder, we let q_i denote the number queries A makes to $\underline{\mathbf{RO}}_i$ and q'_i denote the number of queries A makes to $\underline{\mathbf{RO}}'_i$; note that $q_{\text{RG}} = q_1 + \dots + q_6 + q'$.

Similar to the proof of Theorem 1, note that we have dropped the winning condition on line 16 of the robustness game (Figure 3). The refined measurement computed from the input shares is equal to $\Pi.\text{Unshard}(1, (\Pi.\text{Agg}(\vec{y}_1), \Pi.\text{Agg}(\vec{y}_2)J)) = \vec{y}_1 + \vec{y}_2$, so this condition is never met by definition.

Next, in game $\underline{\mathbf{G}}_1$ (left panel of Figure 24) we revise the definition of the $\underline{\mathbf{RO}}$ oracle so that for each $i \in \{5, 6\}$, the values of $\text{Rand}[i, \text{seed}, \text{cntxt}]$ are sampled without replacement. The new game is identical to $\underline{\mathbf{G}}_2$ up to a collision in the output for either $\text{Rand}[5, \cdot, \cdot]$ or $\text{Rand}[6, \cdot, \cdot]$. Applying a birthday bound over all queries by A or by the **Prep** oracle yields

$$\Pr[\underline{\mathbf{G}}_0(A)] \leq \Pr[\underline{\mathbf{G}}_1(A)] + \frac{(q_5 + 2q_{\text{Prep}})^2}{2^{\kappa+1}} + \frac{(q_6 + 3q_{\text{Prep}})^2}{2^{\kappa+1}}. \quad (26)$$

Next, in game $\underline{\mathbf{G}}_2$ (right panel of Figure 25) we simplify the **Prep** oracle by substituting aggregator \hat{j} 's local computation of the joint randomness seed $j\text{seed}_{\hat{j}}$ with a direct computation of the seed $j\text{seed}$ from the parts ρ_1, ρ_2 computed on lines 9–10. Accordingly, We simplify the joint local randomness checks (lines 26–27) to just check if the purported hint $r\text{seed}[\hat{j}]$ matches the computed part $\rho_{\hat{j}}$ (28–29). This change is only detectable to the adversary if it can find a joint randomness seed and hints such that the check succeeds, but the aggregators compute distinct $j\text{seed}_1 \neq j\text{seed}_2$. This is impossible by construction (transition from $\underline{\mathbf{G}}_1$ to $\underline{\mathbf{G}}_2$), so

<p>Game $\underline{\mathbb{G}}_0(A)$ $\underline{\mathbb{G}}_1(A)$:</p> <p>1 $sk \leftarrow \mathcal{S}SK$; $w \leftarrow \text{false}$; $A^{\text{RO.Prep}(\cdot)}$; ret w</p> <p>$\text{Prep}(n, \vec{x}, msg_{\text{Init}}, st_{\text{Init}})$:</p> <p>2 $(\ell, p\vec{f}x) \leftarrow st$; $u \leftarrow p\vec{f}x$</p> <p>3 if $\text{Used}[n, \ell] \neq \perp$: ret \perp</p> <p>4 $\text{Used}[n, \ell] \leftarrow \top$</p> <p>5 $(pub, rseed) \leftarrow msg_{\text{Init}}$</p> <p>6 $(key_1, seed_1, \pi_1) \leftarrow \text{Unpack}(1, \vec{x}[1], n, \ell)$</p> <p>7 $(key_2, seed_2, \pi_2) \leftarrow \text{Unpack}(2, \vec{x}[2], n, \ell)$</p> <p>8 $\Delta_1 \leftarrow \text{RO}_2(seed_1, n \parallel \ell \parallel 1)$</p> <p>9 $\Delta_2 \leftarrow \text{RO}_2(seed_2, n \parallel \ell \parallel 2)$</p> <p>10 $\rho_1 \leftarrow \text{RO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$</p> <p>11 $\rho_2 \leftarrow \text{RO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$</p> <p>12 $jseed_1 \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \vec{rseed}[2]))$</p> <p>13 $jseed_2 \leftarrow \text{RO}_6(0^\kappa, (\vec{rseed}[1], \rho_2))$</p> <p>14 $\vec{jr}_1 \leftarrow \text{RO}_1(jseed_1, n \parallel \ell)$</p> <p>15 $\vec{jr}_2 \leftarrow \text{RO}_1(jseed_2, n \parallel \ell)$</p> <p>16 $qr \leftarrow \text{RO}_4(sk, n \parallel \ell \parallel \vec{jr}_1)$</p> <p>17 $(h_1, \vec{y}_1) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(1, pub, key_1, p\vec{f}x)$</p> <p>18 $(h_2, \vec{y}_2) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(2, pub, key_2, p\vec{f}x)$</p> <p>19 $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$</p> <p>20 $inp_1 \leftarrow \sum_{i \in [u]} \vec{y}_1[i]$</p> <p>21 $inp_2 \leftarrow \sum_{i \in [u]} \vec{y}_2[i]$</p> <p>22 $\sigma_1 \leftarrow \text{DFLP.Query}(inp_1, \Delta_1, \pi_1, \vec{jr}_1; qr)$</p> <p>23 $\sigma_2 \leftarrow \text{DFLP.Query}(inp_2, \Delta_2, \pi_2, \vec{jr}_2; qr)$</p> <p>24 $\vec{jseed} \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$</p> <p>25 $b_1 \leftarrow jseed_1 = \vec{jseed}$</p> <p>26 $b_2 \leftarrow jseed_2 = \vec{jseed}$</p> <p>27 $v \leftarrow \text{VIDPF.Verify}^{\text{RO}'}(h_1, h_2)$</p> <p>28 $d \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$</p> <p>29 if $\vec{y} \notin \mathcal{V}_{st_{\text{Init}}}$</p> <p>30 and $(b_1 \wedge v \wedge d)$ or $(b_2 \wedge v \wedge d)$: $w \leftarrow \text{true}$</p> <p>31 ret $(w, (msg_{\text{Init}}, ((\sigma_1, \rho_1, h_1), (\sigma_2, \rho_2, h_2))))$</p> <p>$\text{RO}_i(seed, cntxt)$:</p> <p>32 $l \leftarrow (jl, el, m, ql)$</p> <p>33 if $\text{Rand}[i, seed, cntxt] = \perp$:</p> <p>34 if $i \leq 4$: $\text{Rand}[i, seed, cntxt] \leftarrow \mathbb{F}^{l[i]}$</p> <p>35 else: $\text{Rand}[i, seed, cntxt] \leftarrow \{0, 1\}^\kappa$</p> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <p>36 $out \leftarrow \{0, 1\}^\kappa \setminus \mathcal{Q}_i$; $\mathcal{Q}_i \leftarrow \mathcal{Q}_i \cup \{out\}$</p> <p>37 $\text{Rand}[i, seed, cntxt] \leftarrow out$</p> </div> <p>38 ret $\text{Rand}[i, seed, cntxt]$</p> <p>$\text{RO}'(inp)$:</p> <p>39 if $\text{Rand}'[inp] = \perp$: $\text{Rand}'[inp] \leftarrow \mathcal{Y}$</p> <p>40 ret $\text{Rand}'[inp]$</p>	<p>$\text{Prep}(n, \vec{x}, msg_{\text{Init}}, st_{\text{Init}})$: $\underline{\mathbb{G}}_1$ $\underline{\mathbb{G}}_2$</p> <p>1 $(\ell, p\vec{f}x) \leftarrow st$; $u \leftarrow p\vec{f}x$</p> <p>2 if $\text{Used}[n, \ell] \neq \perp$: ret \perp</p> <p>3 $\text{Used}[n, \ell] \leftarrow \top$</p> <p>4 $(pub, rseed) \leftarrow msg_{\text{Init}}$</p> <p>5 $(key_1, seed_1, \pi_1) \leftarrow \text{Unpack}(1, \vec{x}[1], n, \ell)$</p> <p>6 $(key_2, seed_2, \pi_2) \leftarrow \text{Unpack}(2, \vec{x}[2], n, \ell)$</p> <p>7 $\Delta_1 \leftarrow \text{RO}_2(seed_1, n \parallel \ell \parallel 1)$</p> <p>8 $\Delta_2 \leftarrow \text{RO}_2(seed_2, n \parallel \ell \parallel 2)$</p> <p>9 $\rho_1 \leftarrow \text{RO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$</p> <p>10 $\rho_2 \leftarrow \text{RO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$</p> <div style="background-color: #f0f0f0; padding: 2px; margin: 2px 0;"> <p>11 $jseed_1 \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \vec{rseed}[2]))$</p> <p>12 $jseed_2 \leftarrow \text{RO}_6(0^\kappa, (\vec{rseed}[1], \rho_2))$</p> <p>13 $\vec{jr}_1 \leftarrow \text{RO}_1(jseed_1, n \parallel \ell)$</p> <p>14 $\vec{jr}_2 \leftarrow \text{RO}_1(jseed_2, n \parallel \ell)$</p> </div> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <p>15 $\vec{jseed} \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$</p> <p>16 $\vec{jr} \leftarrow \text{RO}_1(\vec{jseed}, n \parallel \ell)$</p> </div> <p>17 $qr \leftarrow \text{RO}_4(sk, n \parallel \ell \parallel \vec{jr}_1)$</p> <p>18 $(h_1, \vec{y}_1) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(1, pub, key_1, p\vec{f}x)$</p> <p>19 $(h_2, \vec{y}_2) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(2, pub, key_2, p\vec{f}x)$</p> <p>20 $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$</p> <p>21 $inp_1 \leftarrow \sum_{i \in [u]} \vec{y}_1[i]$</p> <p>22 $inp_2 \leftarrow \sum_{i \in [u]} \vec{y}_2[i]$</p> <p>23 $\sigma_1 \leftarrow \text{DFLP.Query}(inp_1, \Delta_1, \pi_1, \vec{jr}_1 \parallel \vec{jr}_2; qr)$</p> <p>24 $\sigma_2 \leftarrow \text{DFLP.Query}(inp_2, \Delta_2, \pi_2, \vec{jr}_2 \parallel \vec{jr}_1; qr)$</p> <div style="background-color: #f0f0f0; padding: 2px; margin: 2px 0;"> <p>25 $\vec{jseed} \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$</p> <p>26 $b_1 \leftarrow jseed_1 = \vec{jseed}$</p> <p>27 $b_2 \leftarrow jseed_2 = \vec{jseed}$</p> </div> <div style="border: 1px solid black; padding: 2px; margin: 2px 0;"> <p>28 $b_1 \leftarrow \rho_1 \neq \vec{rseed}[1]$</p> <p>29 $b_2 \leftarrow \rho_2 \neq \vec{rseed}[2]$</p> </div> <p>30 $v \leftarrow \text{VIDPF.Verify}^{\text{RO}'}(h_1, h_2)$</p> <p>31 $d \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$</p> <p>32 if $\vec{y} \notin \mathcal{V}_{st_{\text{Init}}}$</p> <p>33 and $(b_1 \wedge v \wedge d)$ or $(b_2 \wedge v \wedge d)$: $w \leftarrow \text{true}$</p> <p>34 ret $(w, (msg_{\text{Init}}, ((\sigma_1, \rho_1, h_1), (\sigma_2, \rho_2, h_2))))$</p>
---	--

Figure 24: Games $\underline{\mathbb{G}}_0$, $\underline{\mathbb{G}}_1$, and $\underline{\mathbb{G}}_2$ for the proof of Theorem 3. Let \mathcal{Y} denote the co-domain of the random oracle used by VIDPF.

$$\Pr[\underline{\mathbf{G}}_1(A)] = \Pr[\underline{\mathbf{G}}_2(A)]. \quad (27)$$

Next, in game $\underline{\mathbf{G}}_3$ (Figure 25), we make the following changes. First, we modify oracle $\underline{\mathbf{RO}}_4$ so that, for any query that coincides with the secret verification key sk sampled at the beginning of the game, the oracle immediately returns \perp without programming the table \perp . Second, we modify $\underline{\mathbf{Prep}}$ by replacing the call to

$$qr \leftarrow \underline{\mathbf{RO}}_4(sk, n \parallel \ell \parallel)$$

with

$$qr \leftarrow \text{Rand}[4, sk, n \parallel \ell] \leftarrow_{\$} \mathbb{F}^{ql}.$$

That way each call to $\underline{\mathbf{Prep}}$ samples fresh query randomness. The second change does not overwrite any value in Rand due to the first change. Thus the new game is identical to $\underline{\mathbf{G}}_2$ until the adversary makes a query to $\underline{\mathbf{RO}}_4$ with the seed equal to sk . Taking a union bound over all of A 's queries, we have that

$$\Pr[\underline{\mathbf{G}}_2(A)] \leq \Pr[\underline{\mathbf{G}}_3(A)] + \frac{q_4 q_{\text{Prep}}}{2^\kappa}. \quad (28)$$

In the last game, $\underline{\mathbf{G}}_4$ (right-hand panel of Figure 25), we use the extractability of VIDPF to simplify the winning condition. First, we change how the IDPF output vector \vec{y} is computed by $\underline{\mathbf{Prep}}$: If the one-hot check succeeds, i.e., v is set to 1 on line 24, then we use the extractor E to extract $(\alpha, \vec{\beta})$ from the transcript of the random oracle (7) and set \vec{y} to $f_{\alpha, \vec{\beta}}(p\vec{f}x)$. Second, we revise the winning condition (28) by requiring only that the sum of the elements of \vec{y} is not in the delayed-input set $\mathcal{X} = \{0, 1\}$ for DFLP. In particular, we no longer require \vec{y} to be one-hot for the adversary to win. (Recall that $\mathcal{V}_{st_{\text{init}}}$ is the set of one-hot vectors where the non-zero element is in \mathcal{X} .) These conditions are equivalent in the revised game, since (1) A cannot set w if $v = 0$, and if $v = 1$, vector \vec{y} is one-hot by definition.

We claim that there exists an $O(t_A + q_{\text{Prep}} t_E)$ -time adversary B for which

$$\Pr[\underline{\mathbf{G}}_3(A)] \leq \Pr[\underline{\mathbf{G}}_4(A)] + q_{\text{Prep}} \cdot \text{Adv}_{\text{VIDPF}, E}^{\text{extract}}(B). \quad (29)$$

The proof is by a hybrid argument. For each $i \in [q_{\text{Prep}}]$ let $\underline{\mathbf{G}}'_i$ be the game $\underline{\mathbf{G}}_3$ except that only the first i queries to $\underline{\mathbf{Prep}}$ are answered in the usual way; the remaining queries are answered as they are in game $\underline{\mathbf{G}}_4$. Adversary B first samples $i \leftarrow_{\$} [q_{\text{Prep}}]$ then runs $\underline{\mathbf{G}}'_i(A)$ as usual, except that it simulates $\underline{\mathbf{Prep}}$ queries for one of the reports using its own game. Specifically, after unpacking IDPF public share pub and key shares key_1, key_2 on lines 4–6, it pauses the simulation, outputs (pub, key_1, key_2) , and waits to be invoked again. On its second invocation, it resumes the simulation of the $\underline{\mathbf{Prep}}$ query until it reaches the computation of \vec{y} on lines 23–26: At this point it queries its own $\underline{\mathbf{Eval}}$ oracle on the candidate prefixes $p\vec{f}x$ and sets \vec{y} to the return value. Thereafter, it simulates the remainder of the game faithfully. If A sets $w \leftarrow \text{true}$ in its game, then B guesses 1; otherwise it guesses 0.

Let δ_1^i (resp. δ_0^i) denote the probability that B samples i and guesses 1 in the VIDPF extractability experiment, conditioned on the outcome of the coin toss being 1 (resp. 0). Then for all i ,

$$\text{Adv}_{\text{VIDPF}, E}^{\text{extract}}(A) \geq \frac{1}{q_{\text{Prep}}} (\delta_1^i - \delta_0^i). \quad (30)$$

Moreover, by construction we have that

$$\delta_1^i - \delta_0^i = \Pr[\underline{\mathbf{G}}'_i(A)] - \Pr[\underline{\mathbf{G}}'_{i+1}(A)]. \quad (31)$$

for all i . The claim follows from the observation that $\Pr[\underline{\mathbf{G}}_3(A)] = \Pr[\underline{\mathbf{G}}'_0(A)]$ and $\Pr[\underline{\mathbf{G}}_4(A)] = \Pr[\underline{\mathbf{G}}'_{q_{\text{Prep}}}(A)]$.

Consider what A must do to set $w \leftarrow \text{true}$ in game $\underline{\mathbf{G}}_4$. For some $\underline{\mathbf{Prep}}$ query, the delayed-input proof check must succeed when in fact the sum $\sum_{i \in [u]} \vec{y}[i]$ is not a valid encoded input. We bound A 's advantage in game $\underline{\mathbf{G}}_4$ by a reduction to the soundness of DFLP. Recall from the definition of soundness in Section 5.2 that the malicious prover P^* first commits to an encoded input (e, Δ) , then gets a fresh joint randomness jr , then picks a proof forgery π . It wins if $\text{DFLP.Decode}(e) \notin \mathcal{L}$ but the verifier deems the input valid (i.e., $\text{DFLP.Decide}(\text{DFLP.Query}(e, \Delta, \pi, jr; qr)) = 1$, where qr is a fresh query randomness sampled by the game).

consider the malicious prover P^* in Figure 26. The basic idea is that P^* simulates $\underline{G}_4(A)$ and extracts its commitment from queries to the random oracle. Specifically, the prover samples $i^* \leftarrow_{\$} [q_1 + q_{\text{Prep}}]$ at the beginning of the game, and for the i^* -th query to \underline{RO}_1 , it attempts to compute (e, Δ) as follows (see lines 15–19).

The prover maintains a reverse look-up table for random oracle queries for computing the query randomness (i.e., \underline{RO}_4), the joint randomness seed parts (\underline{RO}_5), and the joint randomness seed (\underline{RO}_5). On the i^* -th query, it looks for values $n, \ell, pub, key_1, key_2, seed_1$, and $seed_2$ that would be used by a query to \underline{Prep} . If successful, it uses these to construct its encoded input ($\text{DFLP.Encode}(\Delta, inp^*), \Delta$) to output in its game (20). It computes Δ as the sum of the Δ_j 's corresponding to that query (16–17). So how does it compute inp^* ? Well, in \underline{G}_4 , the \underline{Prep} query corresponding to i^* evaluates IDPF keys shares at a set of candidate prefixes \vec{pfx} chosen by the adversary. But because \vec{pfx} is not known at this point, the best we can do is guess. We therefore choose inp^* by sampling uniform randomly from the set $\mathcal{X} = \{0, 1\}$ of delayed-input values.

If extraction of the commitment is successful, then the prover outputs it, awaits the response from its game, and programs the table with the response jr (21). Thereafter, prover P^* runs $\underline{G}_5(A)$ as usual until a \underline{Prep} query is made for the session (n^*, ℓ^*) that coincides with the distinguished \underline{RO}_1 query i^* . At this point, the prover cannot compute the decision bit d and the verifier shares σ_1, σ_2 consistently, as it does not have access to the query randomness sampled by its game. Instead, it simply halts and outputs $\pi_1 + \pi_2$ as its proof forgery (35–37).

Observe that P^* 's simulation of $\underline{G}_5(A)$ is perfect up until the point it it halts and outputs its forgery. This is due to the full linearity of DFLP, which allows us to substitute the computation of the query-generation algorithm secret-shared data in \underline{G}_5 with the computation of the query-generation algorithm on plaintext inputs in the prover's soundness game. It follows that P^* wins precisely when A sets $w \leftarrow \text{true}$ in the call to \underline{Prep} that coincides with the distinguished session. Conditioning on the probability that P^* guesses the correct call to \underline{RO}_1 , and that we guessed the value of inp^* correctly, we conclude that

$$\Pr[\underline{G}_4(A)] \leq 2(q_1 + q_{\text{Prep}}) \cdot \epsilon. \quad (32)$$

The bound follows from gathering up each of the equations in simplifying.

C.4 Doplar Privacy (Theorem 4)

We begin with a game \underline{G}_0 (Figure 27) in which we instantiate $\text{Exp}_H^{\text{priv}}(A)$ in the random oracle model, in-line the sub-routines of H , and simplify the code. Calls to RG have been replaced with a random oracle \underline{RO} ; as usual, \underline{RO} is implemented by lazy-evaluating a table Rand . In the remainder, we let q_i denote the number of queries A makes to \underline{RO}_i . Another simplifying change we have made is to hard-code the index of the corrupt aggregator, which we denote by \hat{z} . (We denote the honest aggregator by z .) Accordingly, we have removed the share index \hat{j} from the oracle parameters and tables, as there is only one valid choice for these. (This is without loss of generality.) None of these changes impact the outcome of the experiment, so

$$\Pr[\text{Exp}_H^{\text{priv}}(A)] = \Pr[\underline{G}_0(A)]. \quad (33)$$

In game \underline{G}_1 (Figure 27) we revise the \underline{Shard} oracle by sampling the nonce without replacement (line 5). This ensures each report has a unique nonce, which will be useful in subsequent steps. By a birthday bound, we have that

$$\Pr[\underline{G}_0(A)] \leq \Pr[\underline{G}_1(A)] + \frac{q_{\text{Shard}}^2}{2^\kappa}. \quad (34)$$

In our next step, \underline{G}_2 (Figure 28), we modify the \underline{Shard} oracle such that, instead of querying the random oracle \underline{RO} , it *programs* the random oracle using a new sub-routine, PO (31–34). This ensures that the output of \underline{Shard} is not correlated with the game's current state, allowing us to treat the sampled values as fresh. This has a cost, however, since if any of the values programmed by the oracle overwrite existing values in table Rand , then the adversary will end up with an inconsistent view. We can bound this by considering the probability of any one of the following events occurring:

- Seed $seed_1$ or $seed_2$ sampled on line 4 coincides with a query to $\underline{\text{RO}}_2$ made by A (see lines 6–7). We write this as Rand_2 for short in the remainder.
- Seed $seed_1$ or $seed_2$ coincides with an element of Rand_5 (11–12).
- Vector \vec{rseed} sampled on lines 11–12 coincides with an element of Rand_6 (13).
- Seed $jseed$ sampled on line 13 coincides with an element of Rand_1 (16).
- Seed $seed_2$ coincides with an element of Rand_3 (18).

Because the nonces sampled by $\underline{\text{Shard}}$ are unique, and because each of this oracle queries encodes the nonce, we can be certain that points programmed into the table by each $\underline{\text{Shard}}$ query do not collide with one another. Indeed, it is only possible for these values to coincide with random oracle queries made by A . Apply a union bound over all q_{Shard} queries, we conclude that

$$\Pr[\underline{\mathbf{G}}_1(A)] \leq \Pr[\underline{\mathbf{G}}_2(A)] + \frac{q_2 q_{\text{Shard}}}{2^{\kappa-1}} + \frac{q_5 q_{\text{Shard}}}{2^{\kappa-1}} + \frac{q_6 q_{\text{Shard}}}{2^{2\kappa}} + \frac{q_1 q_{\text{Shard}}}{2^\kappa}. \quad (35)$$

In the next step, $\underline{\mathbf{G}}_3$ (Figure 29), we substitute calls to VIDPF.Gen and VIDPF.VEval with calls to the simulator $S = (S_{\text{VIDPF}}^1, S_{\text{VIDPF}}^2)$. The first part, S_{VIDPF}^1 , is used to simulate the public share corrupt aggregator’s key share (10); the second part, S_{VIDPF}^2 , is used to simulate the honest aggregators one-hot check, based on the output of the first (41). After this second point, we no longer compute the honest aggregator’s refined share \vec{y} consistently. Instead, we compute the *corrupt aggregator’s refined share* \vec{y} and compute the challenge input inp by subtracting the sum from the true sum for the input α_b (43–44).

There exists an adversary B for which

$$\Pr[\underline{\mathbf{G}}_2(A)] \leq \Pr[\underline{\mathbf{G}}_3(A)] + q_{\text{Shard}} \cdot \text{Adv}_{\text{VIDPF},S}^{\text{priv}}(B). \quad (36)$$

The proof is by a standard argument. In each hybrid game, we answer one more $\underline{\text{Shard}}$ query (and the corresponding $\underline{\text{Prep}}$ query) using S . Adversary B simply runs A in one of these hybrid games, chosen at random, and outputs whatever A outputs.

In game $\underline{\mathbf{G}}_4$ (Figure 30), we prepare for the $\underline{\text{Shard}}$ oracle for the reduction to DFLP privacy. The primary change is that we have $\underline{\text{Shard}}$ sample the query randomness qr that will be used to query the proof at each level (see line 18 in the left panel). This ensures that the query randomness is “committed” even before the query is made. We use the unpredictability of the nonce to bound the probability that this change leads to an inconsistent view of the experiment. In particular,

$$\Pr[\underline{\mathbf{G}}_3(A)] \leq \Pr[\underline{\mathbf{G}}_4(A)] + \frac{\eta q_4 q_{\text{Shard}}}{|\mathcal{N}|}. \quad (37)$$

In this step, we also make a couple of non-breaking changes. First, we in-line programming of the random oracle with the joint randomness and encoding randomness (16–17,19). Second, we store each proof and encoding randomness in tables P and D respectively. These changes are made to clarify the next step.

In game $\underline{\mathbf{G}}_5$ (Figure 30) we prepare the $\underline{\text{Prep}}$ oracle by re-arranging the proof query. In particular, we run the query-generation algorithm on the plaintext encoded input and proof, and generate the verifier share that is output by subtracting from the verifier (denoted $V[k, \ell]$; see line 19 of the right panel) the verifier share generated from the corrupt aggregator’s share. The adversary’s view is consistent with the previous game by the full linearity of DFLP.

Lastly, in game $\underline{\mathbf{G}}_6$ (not pictured) we modify the $\underline{\text{Prep}}$ oracle by replacing computation of the verifier from α_b with the DFLP-privacy simulator T . There exists an adversary C for which

$$\Pr[\underline{\mathbf{G}}_5(A)] \leq \Pr[\underline{\mathbf{G}}_6(A)] + \eta q_{\text{Shard}} \cdot \text{Adv}_{\text{DFLP},T}^{\text{priv}}(C). \quad (38)$$

The proof is by a hybrid argument, where each hybrid game $\underline{\mathbf{G}}'_{u,v}$ is defined as follows. For the first u reports and for the first v levels of the VIDPF tree, the verifier $V[u, v]$ is generated as specified in game $\underline{\mathbf{G}}_5$ (line 19 in the right panel of Figure 30); all other verifiers are generated by T as specified in game $\underline{\mathbf{G}}_6$. By construction,

$$\Pr[\underline{\mathbf{G}}_5(A)] - \Pr[\underline{\mathbf{G}}_6(A)] = \Pr[\underline{\mathbf{G}}'_{0,0}(A)] - \Pr[\underline{\mathbf{G}}'_{q_{\text{Shard}},\eta}(A)]. \quad (39)$$

Define DFLP-privacy attacker C as follows. (Refer to Figure 6.) On its first invocation, it simply outputs $\mathcal{X} = \{0, 1\}$ as the input set, as this is what is required by the game. On its next invocation, it is given joint randomness jr^* and query randomness qr^* . It proceeds by simulating A in a random hybrid game. It first samples $u^* \leftarrow_s [q_{\text{Shard}}]$ and $v^* \leftarrow_s [\eta]$. It then runs $\underline{\mathcal{G}}'_{u^*, v^*}(A)$ except:

- On the u^* -th query to Shard, for the v^* -th level, it uses jr^* and qr^* to program the random oracles for the joint and query randomness respectively.
- When A makes a Prep query corresponding to report u^* and level v^* , it halts and outputs x_b and awaits a response from its game. Upon being invoked once more on input σ , it sets $V[u^*, v^*] \leftarrow \sigma$ and continues the simulation.

Finally, when A halts, C halts and returns whatever A output. Then C perfectly simulates $\underline{\mathcal{G}}'_{u^*, v^*}(A)$ when the value of its challenge bit is 1, and it perfectly simulates $\underline{\mathcal{G}}'_{u^*, v^*+1}(A)$ when its challenge bit is equal to 0. The claimed bound follows from a standard conditioning argument.

To complete the proof, we note that

$$\Pr[\underline{\mathcal{G}}_6(A)] = \frac{1}{2}. \quad (40)$$

Gathering up all of the terms and simplifying yields the desired bound.

<u>Prep</u> ($n, \vec{x}, msg_{\text{Init}}, st_{\text{Init}}$): <u>G₂</u> <u>G₃</u>	<u>Prep</u> ($n, \vec{x}, msg_{\text{Init}}, st_{\text{Init}}$): <u>G₃</u> <u>G₄</u>
<pre> 1 ($\ell, p\vec{f}x$) $\leftarrow st; u \leftarrow p\vec{f}x$ 2 if Used[n, ℓ] $\neq \perp$: ret \perp 3 Used[n, ℓ] $\leftarrow \top$ 4 ($pub, rseed$) $\leftarrow msg_{\text{Init}}$ 5 ($key_1, seed_1, \pi_1$) $\leftarrow \text{Unpack}(1, \vec{x}[1], n, \ell)$ 6 ($key_2, seed_2, \pi_2$) $\leftarrow \text{Unpack}(2, \vec{x}[2], n, \ell)$ 7 $\Delta_1 \leftarrow \text{RO}_2(seed_1, n \parallel \ell \parallel 1)$ 8 $\Delta_2 \leftarrow \text{RO}_2(seed_2, n \parallel \ell \parallel 2)$ 9 $\rho_1 \leftarrow \text{RO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 10 $\rho_2 \leftarrow \text{RO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 11 $jseed \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$ 12 $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell)$ 13 $qr \leftarrow \text{RO}_4(sk, n \parallel \ell \parallel)$ 14 $qr \leftarrow \text{Rand}[4, sk, n \parallel \ell] \leftarrow_s \mathbb{F}^{ql}$ 15 (h_1, \vec{y}_1) $\leftarrow \text{VIDPF.VEval}^{\text{RO}'}(1, pub, key_1, p\vec{f}x)$ 16 (h_2, \vec{y}_2) $\leftarrow \text{VIDPF.VEval}^{\text{RO}'}(2, pub, key_2, p\vec{f}x)$ 17 $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$ 18 $inp_1 \leftarrow \sum_{i \in [u]} \vec{y}_1[i]$ 19 $inp_2 \leftarrow \sum_{i \in [u]} \vec{y}_2[i]$ 20 $\sigma_1 \leftarrow \text{DFLP.Query}(inp_1, \Delta_1, \pi_1, jr; qr)$ 21 $\sigma_2 \leftarrow \text{DFLP.Query}(inp_2, \Delta_2, \pi_2, jr; qr)$ 22 $b_1 \leftarrow \rho_1 \neq rseed[1]$ 23 $b_2 \leftarrow \rho_2 \neq rseed[2]$ 24 $v \leftarrow \text{VIDPF.Verify}^{\text{RO}'}(h_1, h_2)$ 25 $d \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 26 if $\vec{y} \notin \mathcal{V}_{st_{\text{Init}}}$ 27 and $(b_1 \wedge v \wedge d)$ or $(b_2 \wedge v \wedge d)$: $w \leftarrow \text{true}$ 28 ret $(w, (msg_{\text{Init}}, ((\sigma_1, \rho_1, h_1), (\sigma_2, \rho_2, h_2))))$ </pre>	<pre> 1 ($\ell, p\vec{f}x$) $\leftarrow st; u \leftarrow p\vec{f}x$ 2 if Used[n, ℓ] $\neq \perp$: ret \perp 3 Used[n, ℓ] $\leftarrow \top$ 4 ($pub, rseed$) $\leftarrow msg_{\text{Init}}$ 5 ($key_1, seed_1, \pi_1$) $\leftarrow \text{Unpack}(1, \vec{x}[1], n, \ell)$ 6 ($key_2, seed_2, \pi_2$) $\leftarrow \text{Unpack}(2, \vec{x}[2], n, \ell)$ 7 if T[n] = \perp: T[n] $\leftarrow_s E(key_1, key_2, pub, \text{Rand}')$ 8 $\Delta_1 \leftarrow \text{RO}_2(seed_1, n \parallel \ell \parallel 1)$ 9 $\Delta_2 \leftarrow \text{RO}_2(seed_2, n \parallel \ell \parallel 2)$ 10 $\rho_1 \leftarrow \text{RO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 11 $\rho_2 \leftarrow \text{RO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 12 $jseed \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$ 13 $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell)$ 14 $qr \leftarrow \text{Rand}[4, sk, n \parallel \ell] \leftarrow_s \mathbb{F}^{ql}$ 15 (h_1, \vec{y}_1) $\leftarrow \text{VIDPF.VEval}^{\text{RO}'}(1, pub, key_1, p\vec{f}x)$ 16 (h_2, \vec{y}_2) $\leftarrow \text{VIDPF.VEval}^{\text{RO}'}(2, pub, key_2, p\vec{f}x)$ 17 $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$ 18 $inp_1 \leftarrow \sum_{i \in [u]} \vec{y}_1[i]$ 19 $inp_2 \leftarrow \sum_{i \in [u]} \vec{y}_2[i]$ 20 $\sigma_1 \leftarrow \text{DFLP.Query}(inp_1, \Delta_1, \pi_1, jr; qr)$ 21 $\sigma_2 \leftarrow \text{DFLP.Query}(inp_2, \Delta_2, \pi_2, jr; qr)$ 22 $b_1 \leftarrow \rho_1 \neq rseed[1]$ 23 $b_2 \leftarrow \rho_2 \neq rseed[2]$ 24 $v \leftarrow \text{VIDPF.Verify}^{\text{RO}'}(h_1, h_2)$ 25 if $v = 1$: $(\alpha, \vec{\beta}) \leftarrow_s T[n]; \vec{y} \leftarrow f_{\alpha, \vec{\beta}}(p\vec{f}x)$ 26 else $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$ 27 $d \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 28 if $\vec{y} \notin \mathcal{V}_{st_{\text{Init}}}$ $(\sum_{i \in [u]} \vec{y}[i]) \notin \mathcal{X}$ 29 and $(b_1 \wedge v \wedge d)$ or $(b_2 \wedge v \wedge d)$: $w \leftarrow \text{true}$ 30 ret $(w, (msg_{\text{Init}}, ((\sigma_1, \rho_1, h_1), (\sigma_2, \rho_2, h_2))))$ </pre>
<p><u>RO_i</u>($seed, cntxt$):</p> <pre> 29 if $i = 4 \wedge seed = sk$: ret \perp 30 $l \leftarrow (jl, el, m, ql)$ 31 if Rand[$i, seed, cntxt$] = \perp: 32 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow_s \mathbb{F}^{l[i]}$ 33 else: 34 $out \leftarrow_s \{0, 1\}^\kappa \setminus \mathcal{Q}_i$; $\mathcal{Q}_i \leftarrow \mathcal{Q}_i \cup \{out\}$ 35 Rand[$i, seed, cntxt$] $\leftarrow out$ 36 ret Rand[$i, seed, cntxt$] </pre>	

Figure 25: Games G₃ and G₄ for the proof of Theorem 3. Let $\mathcal{X} = \{0, 1\}$ denote the delayed-input set for DFLP.

```

Adversary  $\mathcal{P}^*[A]()$ :
1  $i^* \leftarrow \mathcal{R}[q_1 + q_{\text{Prep}}]$ ;  $n^*, \ell^* \leftarrow \perp$ ;  $ctr \leftarrow 0$ 
2  $sk \leftarrow \mathcal{R}[\mathcal{SK}]$ ;  $w \leftarrow \text{false}$ ;  $A^{\text{ROExt, PrepSim}}()$ 

ROExt $_i(\text{seed}, \text{cntxt})$ :
3 if  $i = 4 \wedge \text{seed} = sk$ : ret  $\perp$ 
4  $l \leftarrow (jl, el, m, ql)$ 
5 if  $\text{Rand}[i, \text{seed}, \text{cntxt}] = \perp$ :
6   if  $i = 1$ :
7      $ctr \leftarrow ctr + 1$ 
8     if  $i = i^* \wedge$ 
9       if  $(\exists n, \ell, \text{pub}, \text{key}_1, \text{key}_2, \rho_1, \rho_2, \text{seed}_1, \text{seed}_2)$ 
10         $\wedge \rho_1 = \text{Rand}[5, \text{seed}_1, n \parallel \ell \parallel 1 \parallel \text{pub} \parallel \text{key}_1]$ 
11         $\wedge \rho_2 = \text{Rand}[5, \text{seed}_2, n \parallel \ell \parallel 2 \parallel \text{pub} \parallel \text{key}_2]$ 
12         $\wedge \text{seed} = \text{Rand}[6, 0^\kappa, (\rho_1, \rho_2)]$ :
13           $(n^*, \ell^*) \leftarrow (n, \ell)$ 
14          // We don't know  $p\vec{f}x$ , so guess what the sum will be!
15           $\text{inp}^* \leftarrow \{0, 1\}$ 
16           $\Delta_1 \leftarrow \text{ROExt}_2(\text{seed}_1, n \parallel \ell \parallel 1)$ 
17           $\Delta_2 \leftarrow \text{ROExt}_2(\text{seed}_2, n \parallel \ell \parallel 2)$ 
18           $\Delta \leftarrow \Delta_1 + \Delta_2$ 
19           $e \leftarrow \text{DFLP.Encode}(\Delta, \text{inp}^*)$ 
20          output  $(e, \Delta)$  and wait for  $jr$ .
21           $\text{Rand}[1, \text{seed}, \text{cntxt}] \leftarrow jr$ 
22        else:  $\text{Rand}[1, \text{seed}, \text{cntxt}] \leftarrow \mathbb{F}^{jl}$ 
23      else if  $i \in \{2, 3, 4\}$ :  $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow \mathbb{F}^{l[i]}$ 
24      else:
25         $\text{out} \leftarrow \{0, 1\}^\kappa \setminus \mathcal{Q}_i$ ;  $\mathcal{Q}_i \leftarrow \mathcal{Q}_i \cup \{\text{out}\}$ 
26         $\text{Rand}[i, \text{seed}, \text{cntxt}] \leftarrow \text{out}$ 
27      ret  $\text{Rand}[i, \text{seed}, \text{cntxt}]$ 

ROExt' $(\text{inp})$ :
28 if  $\text{Rand}'[\text{inp}] = \perp$ :  $\text{Rand}'[\text{inp}] \leftarrow \mathcal{Y}$ 
29 ret  $\text{Rand}'[\text{inp}]$ 

PrepSim $(n, \vec{x}, \text{msg}_{\text{Init}}, \text{st}_{\text{Init}})$ :
30  $(\ell, p\vec{f}x) \leftarrow \text{st}$ ;  $u \leftarrow |p\vec{f}x|$ 
31 if  $\text{Used}[n, \ell] \neq \perp$ : ret  $\perp$ 
32  $\text{Used}[n, \ell] \leftarrow \top$ 
33  $(\text{pub}, r\vec{seed}) \leftarrow \text{msg}_{\text{Init}}$ 
34  $(\text{key}_1, \text{seed}_1, \pi_1) \leftarrow \text{Unpack}(1, \vec{x}[1], n, \ell)$ 
35  $(\text{key}_2, \text{seed}_2, \pi_2) \leftarrow \text{Unpack}(2, \vec{x}[2], n, \ell)$ 
36 if  $(n^*, \ell^*) = (n, \ell)$ : output  $\pi_1 + \pi_2$  and halt.
37 if  $\text{T}[n] = \perp$ :  $\text{T}[n] \leftarrow \mathcal{R}[E(\text{key}_1, \text{key}_2, \text{pub}, \text{Rand}')]$ 
38  $\Delta_1 \leftarrow \text{RO}_2(\text{seed}_1, n \parallel \ell \parallel 1)$ 
39  $\Delta_2 \leftarrow \text{RO}_2(\text{seed}_2, n \parallel \ell \parallel 2)$ 
40  $\rho_1 \leftarrow \text{RO}_5(\text{seed}_1, n \parallel \ell \parallel 1 \parallel \text{pub} \parallel \text{key}_1)$ 
41  $\rho_2 \leftarrow \text{RO}_5(\text{seed}_2, n \parallel \ell \parallel 2 \parallel \text{pub} \parallel \text{key}_2)$ 
42  $j\text{seed} \leftarrow \text{RO}_6(0^\kappa, (\rho_1, \rho_2))$ 
43  $jr \leftarrow \text{RO}_1(j\text{seed}, n \parallel \ell)$ 
44  $qr \leftarrow \text{Rand}[4, sk, n \parallel \ell] \leftarrow \mathbb{F}^{ql}$ 
45  $(h_1, \vec{y}_1) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(1, \text{pub}, \text{key}_1, p\vec{f}x)$ 
46  $(h_2, \vec{y}_2) \leftarrow \text{VIDPF.VEval}^{\text{RO}'}(2, \text{pub}, \text{key}_2, p\vec{f}x)$ 
47  $\text{inp}_1 \leftarrow \sum_{i \in [u]} \vec{y}_1[i]$ 
48  $\text{inp}_2 \leftarrow \sum_{i \in [u]} \vec{y}_2[i]$ 
49  $\sigma_1 \leftarrow \text{DFLP.Query}(\text{inp}_1, \Delta_1, \pi_1, jr; qr)$ 
50  $\sigma_2 \leftarrow \text{DFLP.Query}(\text{inp}_2, \Delta_2, \pi_2, jr; qr)$ 
51  $b_1 \leftarrow \rho_1 \neq r\vec{seed}[1]$ 
52  $b_2 \leftarrow \rho_2 \neq r\vec{seed}[2]$ 
53  $v \leftarrow \text{VIDPF.Verify}^{\text{RO}'}(h_1, h_2)$ 
54 if  $v = 1$ :  $(\alpha, \vec{\beta}) \leftarrow \mathcal{R}[\text{T}[n]]$ ;  $\vec{y} \leftarrow f_{\alpha, \vec{\beta}}(p\vec{f}x)$ 
55 else  $\vec{y} \leftarrow \vec{y}_1 + \vec{y}_2$ 
56  $d \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 
57 if  $(\sum_{i \in [u]} \vec{y}[i]) \notin \mathcal{X}$ 
58   and  $(b_1 \wedge v \wedge d)$  or  $(b_2 \wedge v \wedge d)$ :  $w \leftarrow \text{true}$ 
59 ret  $(w, (\text{msg}_{\text{Init}}, ((\sigma_1, \rho_1, h_1), (\sigma_2, \rho_2, h_2))))$ 

```

Figure 26: Malicious prover P^* against the soundness of DFLP for the proof of Theorem 3.

```

Game  $\underline{\mathsf{G}}_0(A)$   $\underline{\mathsf{G}}_1(A)$ :
1   $(st_A, \{z\}, (sk, )) \leftarrow_s A^{\text{RO}}(); \tilde{z} \leftarrow 3 - z$ 
2   $b \leftarrow_s \{0, 1\}; b' \leftarrow_s A^{\text{RO, Shard, Setup, Prep, Agg}}(st_A)$ 
3  ret  $b = b'$ 

Shard( $\hat{k} \in \mathbb{N}, \alpha_0, \alpha_1 \in \mathcal{I}$ ):
4  if Used[ $\hat{k}$ ]  $\neq \perp$ : ret  $\perp$ 
5   $n \leftarrow_s \mathcal{N}$   $n \leftarrow_s \mathcal{N} \setminus \mathcal{N}^*; \mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 
6  // Construct the VIDPF key shares.
7   $seed_1, seed_2 \leftarrow_s \{0, 1\}^\kappa$ 
8  for  $\ell \in [\eta]$ :
9     $D[\hat{k}, \ell] \leftarrow \text{RO}_2(seed_1, n \parallel \ell \parallel 1)$ 
10    $+ \text{RO}_2(seed_2, n \parallel \ell \parallel 2)$ 
11    $\vec{\beta}[\ell] \leftarrow \text{Encode}(D[\hat{k}, \ell], 1)$ 
12    $(key_1, key_2, pub) \leftarrow_s \text{VIDPF.Gen}(\alpha_b, \vec{\beta})$ 
13   // Prepare the joint randomness.
14    $rseed[1] \leftarrow \text{RO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 
15    $rseed[2] \leftarrow \text{RO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 
16    $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$ 
17   // Generate the level proofs.
18   for  $\ell \in [\eta]$ :
19      $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell)$ 
20      $\pi \leftarrow \text{DFLP.Prove}(\{0, 1\}, D[\hat{k}, \ell], jr)$ 
21      $\vec{p}\hat{f}[\ell] \leftarrow \pi - \text{RO}_3(seed_2, n \parallel \ell)$ 
22   // Prepare the initial message and input shares.
23    $x_1 \leftarrow (key_1, seed_1, \vec{p}\hat{f})$ 
24    $x_2 \leftarrow (key_2, seed_2)$ 
25    $\text{In}[\hat{k}] \leftarrow x_z$ 
26    $\text{Pub}[\hat{k}] \leftarrow (pub, rseed)$ 
27    $\text{Used}[\hat{k}] \leftarrow (n, \alpha_0, \alpha_1)$ 
28   ret  $(n, \text{Pub}[\hat{k}], (x_2, ))$ 

Setup( $\hat{i} \in \mathbb{N}, st_{\text{Init}} \in \mathcal{Q}_{\text{Init}}$ ):
29  $(\ell, \vec{p}\hat{f}x) \leftarrow st_{\text{Init}}$ 
30 if Status[ $\hat{i}$ ]  $\neq \perp$  or  $\ell \in \mathcal{U}$  or  $\vec{p}\hat{f}x$  not distinct: ret  $\perp$ 
31  $\mathcal{U} \leftarrow \mathcal{U} \cup \{\ell\}$ 
32 Setup[ $\hat{i}$ ]  $\leftarrow st_{\text{Init}}$ ; Status[ $\hat{i}$ ]  $\leftarrow \text{running}$ 

Prep( $\hat{i} \in \mathbb{N}, \hat{k} \in \mathbb{N}, m\hat{s}g \in \mathcal{M}^*$ ):
33 if Status[ $\hat{i}$ ]  $\neq \text{running}$  or In[ $\hat{k}$ ] =  $\perp$ : ret  $\perp$ 
34 if St[ $\hat{i}, \hat{k}$ ] =  $\perp$ : St[ $\hat{i}, \hat{k}$ ]  $\leftarrow \text{Setup}[\hat{i}]$ 
35  $(n, \alpha_0, \alpha_1) \leftarrow \text{Used}[\hat{k}]$ 
36 if St[ $\hat{i}, \hat{k}$ ]  $\in \mathcal{Q}_{\text{Init}}$ : // Process initial message from client
37  $(\ell, \vec{p}\hat{f}x) \leftarrow \text{St}[\hat{i}, \hat{k}]; u \leftarrow |\vec{p}\hat{f}x|$ 
38  $(pub, rseed) \leftarrow \text{Pub}[\hat{k}]$ 
39  $(key, seed, \pi) \leftarrow \text{Unpack}(z, \text{In}[\hat{k}], n, \ell)$ 
40  $\Delta \leftarrow \text{RO}_2(seed, n \parallel \ell \parallel z)$ 
41  $rseed[z] \leftarrow \text{RO}_5(seed, n \parallel \ell \parallel z \parallel pub \parallel key)$ 
42  $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$ 
43  $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell); qr \leftarrow \text{RO}_4(sk, n \parallel \ell)$ 
44  $(h, \vec{y}) \leftarrow \text{VIDPF.VEval}(z, pub, key, \vec{p}\hat{f}x)$ 
45  $inp \leftarrow \sum_{i \in [u]} \vec{y}[i]$ 
46  $\sigma \leftarrow \text{DFLP.Query}(inp, \Delta, \pi, jr; qr)$ 
47  $msg \leftarrow (\sigma, rseed[z], h)$ 
48 St[ $\hat{i}, \hat{k}$ ]  $\leftarrow (jseed, (\text{DFLP.Decode}(\vec{y}[i])_{i \in [u]}))$ 
49 ret  $(\text{running}, msg)$ 
50 // Process broadcast messages from aggregators
51  $(jseed, \vec{y}) \leftarrow \text{St}[\hat{i}, \hat{k}]; \text{St}[\hat{i}, \hat{k}] \leftarrow \perp$ 
52  $((\sigma_1, rseed_1, h_1), (\sigma_2, rseed_2, h_2)) \leftarrow m\hat{s}g$ 
53  $acc_{\text{DFLP}} \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 
54  $acc_{\text{VIDPF}} \leftarrow \text{VIDPF.Verify}(h_1, h_2)$ 
55  $acc_0 \leftarrow jseed = \text{RO}_6(0^\kappa, (rseed_1, rseed_2))$ 
56 if  $acc_{\text{DFLP}}$  and  $acc_{\text{VIDPF}}$  and  $acc_0$ :
57 Out[ $\hat{i}, \hat{k}$ ]  $\leftarrow \vec{y}$ ; Batch0[ $\hat{i}, \hat{k}$ ]  $\leftarrow \alpha_0$ ; Batch1[ $\hat{i}, \hat{k}$ ]  $\leftarrow \alpha_1$ 
58 ret finished
59 ret failed

Agg( $\hat{i} \in \mathbb{N}$ ):
60 if Status[ $\hat{i}$ ]  $\neq \text{running}$ : ret  $\perp$ 
61  $st_{\text{Init}} \leftarrow \text{Setup}[\hat{i}]$ 
62 if  $F(st_{\text{Init}}, \text{Batch}_0[\hat{i}, \cdot]) \neq F(st_{\text{Init}}, \text{Batch}_1[\hat{i}, \cdot])$ : ret  $\perp$ 
63 Status[ $\hat{i}$ ]  $\leftarrow \text{finished}$ 
64 ret  $\sum_{\vec{y} \in \text{Out}[\hat{i}, \cdot]} \vec{y}$ 

```

Figure 27: Games $\underline{\mathsf{G}}_0$ and $\underline{\mathsf{G}}_1$ for the proof of Theorem 4.

<pre> Shard($\hat{k} \in \mathbb{N}, \alpha_0, \alpha_1 \in \mathcal{I}$): 1 if Used[$\hat{k}$] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 // Construct the VIDPF key shares. 4 $seed_1, seed_2 \leftarrow \{0, 1\}^\kappa$ 5 for $\ell \in [\eta]$: 6 $D[\hat{k}, \ell] \leftarrow \text{RO}_2 \text{ PO}_2(seed_1, n \parallel \ell \parallel 1)$ 7 + $\text{RO}_2 \text{ PO}_2(seed_2, n \parallel \ell \parallel 2)$ 8 $\vec{\beta}[\ell] \leftarrow \text{Encode}(D[\hat{k}, \ell], 1)$ 9 $(key_1, key_2, pub) \leftarrow \text{VIDPF.Gen}(\alpha_b, \vec{\beta})$ 10 // Prepare the joint randomness. 11 $rseed[1] \leftarrow \text{RO}_5 \text{ PO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 12 $rseed[2] \leftarrow \text{RO}_5 \text{ PO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 13 $jseed \leftarrow \text{RO}_6 \text{ PO}_6(0^\kappa, \vec{rseed})$ 14 // Generate the level proofs. 15 for $\ell \in [\eta]$: 16 $jr \leftarrow \text{RO}_1 \text{ PO}_1(jseed, n \parallel \ell)$ 17 $\pi \leftarrow \text{DFLP.Prove}(\{0, 1\}, D[\hat{k}, \ell], jr)$ 18 $\vec{pf}[\ell] \leftarrow \pi - \text{RO}_3 \text{ PO}_3(seed_2, n \parallel \ell)$ 19 // Prepare the initial message and input shares. 20 $x_1 \leftarrow (key_1, seed_1, \vec{pf})$ 21 $x_2 \leftarrow (key_2, seed_2)$ 22 $\text{In}[\hat{k}] \leftarrow x_z$ 23 $\text{Pub}[\hat{k}] \leftarrow (pub, \vec{rseed})$ 24 $\text{Used}[\hat{k}] \leftarrow (n, \alpha_0, \alpha_1)$ 25 ret $(n, \text{Pub}[\hat{k}], (x_z,))$ </pre>	<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; padding: 5px;"> <pre> RO_i(seed, cntxt): 26 $l \leftarrow (jl, el, m, ql)$ 27 if Rand[$i, seed, cntxt$] = \perp: 28 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 29 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 30 ret Rand[$i, seed, cntxt$] </pre> </td> <td style="width: 50%; text-align: center; padding: 5px;"> <div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_1</div> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_2</div> </div> </td> </tr> <tr> <td style="padding: 5px;"> <pre> PO_i(seed, cntxt): 31 $l \leftarrow (jl, el, m, ql)$ 32 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 33 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 34 ret Rand[$i, seed, cntxt$] </pre> </td> <td></td> </tr> </table>	<pre> RO_i(seed, cntxt): 26 $l \leftarrow (jl, el, m, ql)$ 27 if Rand[$i, seed, cntxt$] = \perp: 28 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 29 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 30 ret Rand[$i, seed, cntxt$] </pre>	<div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_1</div> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_2</div> </div>	<pre> PO_i(seed, cntxt): 31 $l \leftarrow (jl, el, m, ql)$ 32 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 33 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 34 ret Rand[$i, seed, cntxt$] </pre>	
<pre> RO_i(seed, cntxt): 26 $l \leftarrow (jl, el, m, ql)$ 27 if Rand[$i, seed, cntxt$] = \perp: 28 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 29 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 30 ret Rand[$i, seed, cntxt$] </pre>	<div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_1</div> <div style="border: 1px solid black; padding: 2px;">\mathbb{G}_2</div> </div>				
<pre> PO_i(seed, cntxt): 31 $l \leftarrow (jl, el, m, ql)$ 32 if $i \leq 4$: Rand[$i, seed, cntxt$] $\leftarrow \mathbb{F}^{l[i]}$ 33 else: Rand[$i, seed, cntxt$] $\leftarrow \{0, 1\}^\kappa$ 34 ret Rand[$i, seed, cntxt$] </pre>					

Figure 28: Game \mathbb{G}_2 for the proof of Theorem 4.

<pre> Shard($\hat{k} \in \mathbb{N}, \alpha_0, \alpha_1 \in \mathcal{I}$): 1 if Used[$\hat{k}$] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 // Construct the VIDPF key shares. 4 $seed_1, seed_2 \leftarrow \{0, 1\}^\kappa$ 5 for $\ell \in [\eta]$: 6 $D[\hat{k}, \ell] \leftarrow \text{PO}_2(seed_1, n \parallel \ell \parallel 1)$ 7 $+ \text{PO}_2(seed_2, n \parallel \ell \parallel 2)$ 8 $\vec{\beta}[\ell] \leftarrow \text{Encode}(D[\hat{k}, \ell], 1)$ 9 $(key_1, key_2, pub) \leftarrow \text{VIDPF.Gen}(\alpha_b, \vec{\beta})$ 10 $(T[\hat{k}], pub) \leftarrow S_{\text{VIDPF}}^1(\vec{z}); key_z \leftarrow T[\hat{k}]; key_z \leftarrow \perp$ 11 // Prepare the joint randomness. 12 $rseed[1] \leftarrow \text{PO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 13 $rseed[2] \leftarrow \text{PO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 14 $jseed \leftarrow \text{PO}_6(0^\kappa, rseed)$ 15 // Generate the level proofs. 16 for $\ell \in [\eta]$: 17 $jr \leftarrow \text{PO}_1(jseed, n \parallel \ell)$ 18 $\pi \leftarrow \text{DFLP.Prove}(\{0, 1\}, D[\hat{k}, \ell], jr)$ 19 $\vec{pf}[\ell] \leftarrow \pi - \text{PO}_3(seed_2, n \parallel \ell)$ 20 // Prepare the initial message and input shares. 21 $x_1 \leftarrow (key_1, seed_1, \vec{pf})$ 22 $x_2 \leftarrow (key_2, seed_2)$ 23 $\text{In}[\hat{k}] \leftarrow x_z$ 24 $\text{Pub}[\hat{k}] \leftarrow (pub, rseed)$ 25 $\text{Used}[\hat{k}] \leftarrow (n, \alpha_0, \alpha_1)$ 26 ret $(n, \text{Pub}[\hat{k}], (x_z,))$ </pre>	<div style="text-align: right; margin-bottom: 5px;"> \mathbb{G}_2 \mathbb{G}_3 </div> <pre> Prep($\hat{i} \in \mathbb{N}, \hat{k} \in \mathbb{N}, msg \in \mathcal{M}^*$): 27 if Status[$\hat{i}$] $\neq \text{running}$ or $\text{In}[\hat{k}] = \perp$: ret \perp 28 if $\text{St}[\hat{i}, \hat{k}] = \perp$: $\text{St}[\hat{i}, \hat{k}] \leftarrow \text{Setup}[\hat{i}]$ 29 $(n, \alpha_0, \alpha_1) \leftarrow \text{Used}[\hat{k}]$ 30 if $\text{St}[\hat{i}, \hat{k}] \in \mathcal{Q}_{\text{Init}}$: // Process initial message from client 31 $(\ell, \vec{pfx}) \leftarrow \text{St}[\hat{i}, \hat{k}]; u \leftarrow \vec{pfx}$ 32 $(pub, rseed) \leftarrow \text{Pub}[\hat{k}]$ 33 $(key_{\square}, seed, \pi) \leftarrow \text{Unpack}(z, \text{In}[\hat{k}], n, \ell)$ 34 $\Delta \leftarrow \text{RO}_2(seed, n \parallel \ell \parallel z)$ 35 $rseed[z] \leftarrow \text{RO}_5(seed, n \parallel \ell \parallel z \parallel pub \parallel key)$ 36 $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$ 37 $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell); qr \leftarrow \text{RO}_4(sk, n \parallel \ell)$ 38 $(h, \vec{y}) \leftarrow \text{VIDPF.VEval}(z, pub, key, \vec{pfx})$ 39 $inp \leftarrow \sum_{i \in [u]} \vec{y}[i]$ 40 $key_z \leftarrow T[\hat{k}]$ 41 $h \leftarrow S_{\text{VIDPF}}^2(\vec{z}, pub, key_z, \vec{pfx})$ 42 $(-, \vec{y}) \leftarrow \text{VIDPF.VEval}(\vec{z}, pub, key_z, \vec{pfx})$ 43 $x_b \leftarrow \{\vec{pfx}[i] : \vec{pfx}[i] \text{ prefixes } \alpha_b\}_{i \in [u]}$ 44 $inp_b \leftarrow \text{DFLP.Encode}(\Delta[\hat{k}, \ell], x_b)$ 45 $inp \leftarrow inp_b - \sum_{i \in [u]} \vec{y}[i]$ 46 $\sigma \leftarrow \text{DFLP.Query}(inp, \Delta, \pi, jr; qr)$ 47 $msg \leftarrow (\sigma, rseed[z], h)$ 48 $\text{St}[\hat{i}, \hat{k}] \leftarrow (jseed, (\text{DFLP.Decode}(\vec{y}[i])_{i \in [u]}))$ 49 ret (running, msg) 50 // Process broadcast messages from aggregators 51 $(jseed, \vec{y}) \leftarrow \text{St}[\hat{i}, \hat{k}]; \text{St}[\hat{i}, \hat{k}] \leftarrow \perp$ 52 $((\sigma_1, rseed_1, h_1), (\sigma_2, rseed_2, h_2)) \leftarrow msg$ 53 $acc_{\text{DFLP}} \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 54 $acc_{\text{VIDPF}} \leftarrow \text{VIDPF.Verify}(h_1, h_2)$ 55 $acc_0 \leftarrow jseed = \text{RO}_6(0^\kappa, (rseed_1, rseed_2))$ 56 if acc_{DFLP} and acc_{VIDPF} and acc_0: 57 $\text{Out}[\hat{i}, \hat{k}] \leftarrow \vec{y}; \text{Batch}_0[\hat{i}, \hat{k}] \leftarrow \alpha_0; \text{Batch}_1[\hat{i}, \hat{k}] \leftarrow \alpha_1$ 58 ret finished 59 ret failed </pre>
---	---

Figure 29: Game \mathbb{G}_3 for the proof of Theorem 4.

Shard ($\hat{k} \in \mathbb{N}, \alpha_0, \alpha_1 \in \mathcal{I}$):	\mathbb{G}_3 \mathbb{G}_4	Prep ($\hat{i} \in \mathbb{N}, \hat{k} \in \mathbb{N}, \vec{msg} \in \mathcal{M}^*$):	\mathbb{G}_4 \mathbb{G}_5
<pre> 1 if Used[\hat{k}] $\neq \perp$: ret \perp 2 $n \leftarrow \mathcal{N} \setminus \mathcal{N}^*$; $\mathcal{N}^* \leftarrow \mathcal{N}^* \cup \{n\}$ 3 // Construct the VIDPF key shares. 4 $seed_1, seed_2 \leftarrow \{0, 1\}^\kappa$ 5 for $\ell \in [\eta]$: 6 $D[\hat{k}, \ell] \leftarrow \text{PO}_2(seed_1, n \parallel \ell \parallel 1)$ 7 $+ \text{PO}_2(seed_2, n \parallel \ell \parallel 2)$ 8 $(T[\hat{k}], pub) \leftarrow S_{\text{VIDPF}}^1(\vec{z}); key_z \leftarrow T[\hat{k}]; key_z \leftarrow \perp$ 9 // Prepare the joint randomness. 10 $rseed[1] \leftarrow \text{PO}_5(seed_1, n \parallel \ell \parallel 1 \parallel pub \parallel key_1)$ 11 $rseed[2] \leftarrow \text{PO}_5(seed_2, n \parallel \ell \parallel 2 \parallel pub \parallel key_2)$ 12 $jseed \leftarrow \text{PO}_6(0^\kappa, rseed)$ 13 // Generate the level proofs. 14 for $\ell \in [\eta]$: 15 $jr \leftarrow \mathbb{F}^{jl}; qr \leftarrow \mathbb{F}^{ql}; D[\hat{k}, \ell], \vec{\Delta} \leftarrow \mathbb{F}^{el}$ 16 $\text{Rand}[2, seed_z, n \parallel \ell \parallel z] \leftarrow D[\hat{k}, \ell] - \vec{\Delta}$ 17 $\text{Rand}[2, seed_z, n \parallel \ell \parallel \vec{z}] \leftarrow \vec{\Delta}$ 18 $\text{Rand}[4, sk, n \parallel \ell] \leftarrow qr$ 19 $\text{Rand}[1, jseed, n \parallel \ell] \leftarrow jr$ 20 $P[\hat{k}, \ell] \leftarrow \text{DFLP.Prove}(\{0, 1\}, \Delta, jr)$ 21 $\vec{pf}[\ell] \leftarrow P[\hat{k}, \ell] - \text{PO}_3(seed_2, n \parallel \ell)$ 22 $jr \leftarrow \text{PO}_1(jseed, n \parallel \ell)$ 23 $\pi \leftarrow \text{DFLP.Prove}(\{0, 1\}, D[\hat{k}, \ell], jr)$ 24 $\vec{pf}[\ell] \leftarrow \pi - \text{PO}_3(seed_2, n \parallel \ell)$ 25 // Prepare the initial message and input shares. 26 $x_1 \leftarrow (key_1, seed_1, \vec{pf}); x_2 \leftarrow (key_2, seed_2)$ 27 $\text{In}[\hat{k}] \leftarrow x_2; \text{Pub}[\hat{k}] \leftarrow (pub, rseed)$ 28 $\text{Used}[\hat{k}] \leftarrow (n, \alpha_0, \alpha_1)$ 29 ret $(n, \text{Pub}[\hat{k}], (x_z,))$ </pre>		<pre> 1 if Status[\hat{i}] $\neq \text{running}$ or $\text{In}[\hat{k}] = \perp$: ret \perp 2 if $\text{St}[\hat{i}, \hat{k}] = \perp$: $\text{St}[\hat{i}, \hat{k}] \leftarrow \text{Setup}[\hat{i}]$ 3 $(n, \alpha_0, \alpha_1) \leftarrow \text{Used}[\hat{k}]$ 4 if $\text{St}[\hat{i}, \hat{k}] \in \mathcal{Q}_{\text{Init}}$: // Process initial message from client 5 $(\ell, \vec{pfx}) \leftarrow \text{St}[\hat{i}, \hat{k}]; u \leftarrow \vec{pfx}$ 6 $(pub, rseed) \leftarrow \text{Pub}[\hat{k}]$ 7 $(-, seed, \pi) \leftarrow \text{Unpack}(z, \text{In}[\hat{k}], n, \ell)$ 8 $\Delta \leftarrow \text{RO}_2(seed, n \parallel \ell \parallel z)$ 9 $rseed[z] \leftarrow \text{RO}_5(seed, n \parallel \ell \parallel z \parallel pub \parallel key)$ 10 $jseed \leftarrow \text{RO}_6(0^\kappa, rseed)$ 11 $jr \leftarrow \text{RO}_1(jseed, n \parallel \ell); qr \leftarrow \text{RO}_4(sk, n \parallel \ell)$ 12 $key_z \leftarrow T[\hat{k}]$ 13 $h \leftarrow S_{\text{VIDPF}}^2(\vec{z}, pub, key_z, \vec{pfx})$ 14 $(-, \vec{y}) \leftarrow \text{VIDPF.VEval}(\vec{z}, pub, key_z, \vec{pfx})$ 15 $x_b \leftarrow \{\vec{pfx}[i] : \vec{pfx}[i] \text{ prefixes } \alpha_b\}_{i \in [u]}$ 16 $inp_b \leftarrow \text{DFLP.Encode}(\Delta[\hat{k}, \ell], x_b)$ 17 $inp \leftarrow inp_b - \sum_{i \in [u]} \vec{y}[i]$ 18 $\sigma \leftarrow \text{DFLP.Query}(inp, \Delta, \pi, jr; qr)$ 19 $V[\hat{k}, \ell] \leftarrow \text{DFLP.Query}(inp_b, D[\hat{k}, \ell], P[\hat{k}, \ell], jr; qr)$ 20 $\sigma \leftarrow V[\hat{k}, \ell] - \text{DFLP.Query}(\sum_{i \in [u]} \vec{y}[i], \Delta, \pi, jr; qr)$ 21 $msg \leftarrow (\sigma, rseed[z], h)$ 22 $\text{St}[\hat{i}, \hat{k}] \leftarrow (jseed, (\text{DFLP.Decode}(\vec{y}[i])_{i \in [u]}))$ 23 ret $(\text{running}, msg)$ 24 // Process broadcast messages from aggregators 25 $(jseed, \vec{y}) \leftarrow \text{St}[\hat{i}, \hat{k}]; \text{St}[\hat{i}, \hat{k}] \leftarrow \perp$ 26 $((\sigma_1, rseed_1, h_1), (\sigma_2, rseed_2, h_2)) \leftarrow \vec{msg}$ 27 $acc_{\text{DFLP}} \leftarrow \text{DFLP.Decide}(\sigma_1 + \sigma_2)$ 28 $acc_{\text{VIDPF}} \leftarrow \text{VIDPF.Verify}(h_1, h_2)$ 29 $acc_0 \leftarrow jseed = \text{RO}_6(0^\kappa, (rseed_1, rseed_2))$ 30 if acc_{DFLP} and acc_{VIDPF} and acc_0: 31 $\text{Out}[\hat{i}, \hat{k}] \leftarrow \vec{y}; \text{Batch}_0[\hat{i}, \hat{k}] \leftarrow \alpha_0; \text{Batch}_1[\hat{i}, \hat{k}] \leftarrow \alpha_1$ 32 ret finished 33 ret failed </pre>	

Figure 30: Games \mathbb{G}_4 and \mathbb{G}_5 for the proof of Theorem 4.