

PLASMA: Private, Lightweight Aggregated Statistics against Malicious Adversaries

Dimitris Mouris^{1*}, Pratik Sarkar^{2*}, and Nektarios Georgios Tsoutsos³

¹ University of Delaware & Nillion
dimitris@nillion.com

² Supra Research
pratik93@bu.edu

³ University of Delaware
tsoutsos@udel.edu

Abstract. Private heavy-hitters is a data-collection task where multiple clients possess private bit strings, and data-collection servers aim to identify the most popular strings without learning anything about the clients’ inputs. In this work, we introduce PLASMA: a private analytics framework in the three-server setting that protects the privacy of honest clients and the correctness of the protocol against a coalition of malicious clients and a malicious server.

Our core primitives are a verifiable incremental distributed point function (VIDPF) and a batched consistency check, which are of independent interest. Our VIDPF introduces new methods to validate client inputs based on hashing. Meanwhile, our batched consistency check uses Merkle trees to validate multiple client sessions together in a batch. This drastically reduces server communication across multiple client sessions, resulting in significantly less communication compared to related works. Finally, we compare PLASMA with the recent works of Asharov et al. (CCS’22) and Poplar (S&P’21) and compare in terms of monetary cost for different input sizes.

Keywords: Function secret sharing, histograms, heavy hitters, privacy enhancing technologies, secure multiparty computation

1 Introduction

In today’s technology-driven world, companies are constantly collecting user data to perform analysis, compute statistics, expose patterns in user behaviors, and apply them to improve their products [EPK14, IKN+20, LPR+21, CJK+23, PAD+23]. Common analysis practices resort to histograms, where client data are aggregated together in predefined and non-overlapping buckets. Each bucket may represent a quantitative range (e.g., salary) or a categorical value (e.g., profession). The resulting histogram displays the frequencies of each bucket based on multiple aggregated participant responses.

Private Histograms. When computing histograms, it is crucial to maintain client privacy, such as preventing data collection servers from inferring additional information about the clients. Existing solutions for privacy-preserving histograms solve this problem efficiently [CGB17, BGG+22, BBC+21], given a relatively small number of buckets. Nevertheless, histograms are resource-intensive on the server side when the goal is to find popular entries among the clients’ inputs. For instance, assume clients that hold GPS coordinates of their location and servers aiming to discover crowded areas without compromising client privacy. The naive solution of creating a histogram over all possible inputs results in sparsely populated sets, which wastes server-side computational power due to sparse inputs. Conversely, in an optimal solution, the server computation should scale with the most popular inputs, instead of all possible ones.

* The first two authors have equal contribution and appear in alphabetical order.

¹ Research mainly conducted at the University of Delaware and completed at Nillion.

Private Heavy-Hitters. This problem is addressed by the concept of “heavy hitters”. \mathcal{T} -heavy hitters allow computing the \mathcal{T} most popular responses (for a given threshold \mathcal{T}) among clients’ inputs and have a broad range of applications: from finding popular websites that users visit or malicious URLs that cause browsers to crash [HKR12, BBC+21], to discovering commonly used passwords [NPR19], learning new words typed by users and identifying frequently used emojis [FPE16], to name a few. Private heavy-hitters allow computing these results while also preserving client privacy. Existing protocols (such as [NPR19, BBC+21, ACD+21, BK21]) only focus on the “popular” inputs and disregard other inputs that appear less than \mathcal{T} times (i.e., they are pruned by the protocol). This renders private heavy hitters a suitable candidate for finding the most common client entries, such as computing crowded areas using client-provided GPS coordinates.

Table 1: Threat model comparisons, client input validation, and server-to-server communication.

Protocol	Correctness & Privacy Against Malicious Corruption			Client Input Validation	Low Server-to-Server Communication	No. of Servers
	Clients	Server	Server & Clients			
DPF [BGI15, BGI16, GI14]	●	● [†]	○	○	○	2+
Poplar (IDPF) [BBC+21]	●	● [†]	○	●	○	2
Bucketization (DP) [ACD+21]	●	● [†]	○	●	○	2-3
MPC-based [BK21]	● [‡]	● [†]	○	○	○	3
Sorting-based [AHI+22, JKK+22]	●	●	●	●	○	3
PLASMA (this work)	●	●	●	●	●	3

[†] These works only preserve privacy against a malicious server but not correctness.

[‡] [BK21] is susceptible to data poisoning attacks by malicious clients or malicious servers. Privacy of honest clients is preserved.

Different Approaches. The literature considers the setting where two or more servers collect client inputs and run the private heavy-hitters protocol. A notable approach based on differential privacy (DP) is [ACD+21] (we discuss DP-based solutions in Section 1.2). While these protocols are computationally fast, they are limited to DP-based privacy guarantees for the client. Likewise, MPC-based solutions [BK21] employ general-purpose secure computation frameworks (e.g., MP-SPDZ [Kel20], SCALE-MAMBA [ACC+21], Sharemind [BLW08]), so these methods fall short in terms of practicality. Thus, recent works introduced custom MPC-based techniques for private heavy-hitters [AHI+22, JKK+22]. The underlying protocols perform secure sorting of client inputs under MPC [AHI+22, JKK+22] and then aggregate the sorted data, guaranteeing that private inputs remain hidden when a majority of the servers are honest. However, the communication of all aforementioned solutions is linearly dependent on the number of clients, resulting in high server-to-server communication costs.

Distributed point functions (DPFs) [BGI15] offer an alternative approach for private histograms. Informally, DPFs allow a client to send succinct shares of a point function corresponding to their private inputs to two or more servers. The servers then use these shares to locally evaluate the function over the entire input space and add the resulting outputs to obtain additive shares of a histogram.

Poplar [BBC+21] builds upon the DPF approach by introducing incremental DPFs (IDPF), detailed in Appendix A. It provides an IDPF-based solution for private heavy-hitters in the two-server setting, and their server-to-server communication depends on the input string length in semi-honest security. For security against malicious clients, the servers validate every client’s input so that malformed inputs are preemptively discarded from the computation. This is referred to as *client input validation* and it prevents malicious clients from causing an abort in the protocol. To do so, Poplar requires additional checks, which cause the server-to-server communication to scale linearly with the total number of clients. As a result, their concrete server-to-server communication is large. Sabre [VSH22] uses multi-verifier MPC-in-the-head that attests to the well-formedness of DPFs but does not focus on heavy hitters. The concurrent work of Doplar also introduced a “Verifiable IDPF (VIDPF)” similar to ours, which guarantees the same security properties. However, their constructions, namely Doplar and Prio3, rely on multiple Fully-Linear Proofs (FLPs) [DPRS23] to verify that

the client’s input is valid, resulting in significant communication overheads. Moreover, their approach does not consider malicious servers.

Motivation. Since all aforementioned solutions incur server-to-server communication that scales linearly with the number of clients (with large concrete communication costs), they are prohibitive for most real-world applications that require millions of clients for data collection. The concrete server-to-server communication should be low, even for a large number of clients. Likewise, neither Poplar nor the DP-based solutions [ACD⁺21] tolerate additive attacks from a malicious server, which results in incorrect outputs when one of the servers does not follow the protocol steps. More formally, they fail to provide both correctness and privacy against the collusion of a malicious server and malicious clients. In this regard, we ask the following motivating question:

Can we obtain a private heavy-hitters protocol with low concrete server-to-server communication that is secure against malicious clients and a malicious server?

1.1 Our Contributions

We answer the aforementioned question with PLASMA, a framework for private statistics that provides security against a malicious server and malicious clients. Our contributions are as follows:

Verifiable incremental DPF (VIDPF). First, we introduce our VIDPF primitive, which builds upon incremental DPFs (IDPF) [BBC⁺21] and verifiable DPFs (VDPF) [dCP22]. VIDPF allows us to verify that clients’ inputs are valid by relying on hashing while preserving the client’s input privacy. We also propose a novel way to verify that IDPF keys are “one-hot” - i.e., they have a single non-zero evaluation path (containing the same value along the path) by solely relying on hashing. This is of independent interest and can be used to improve earlier results in [BBC⁺21, DRPS22, DPRS23]. Previous protocols solved this problem using FLPs [BBC⁺19, DPRS23] or expensive sketching that involves information-theoretic MACs [BBC⁺21, DRPS22, BBC⁺23]. More specifically, [DPRS23] uses FLPs in each level to verify that the client’s input is one-hot, resulting in significant communication overhead as each FLP entails a large proof. Conversely, our checks for one-hot vectors do not require field multiplications, only additions and hashes which allow us to batch-verify multiple inputs together.

Batched Consistency Check. Next, we introduce a novel batched consistency check that allows us to drastically reduce server-to-server communication. At a high level, we validate the inputs of ℓ clients using a Merkle tree and identify the malformed ones using logarithmic (in the total number of clients denoted as ℓ) communication. This optimization reduces the dependency of our server-to-server communication on the total number of clients from $\mathcal{O}(\ell)$ to $\mathcal{O}(\ell(\log_2 \frac{\ell}{\ell'}))$ number of hashes where there are ℓ' malicious clients, yielding a concrete improvement over the state-of-the-art (as reported in our experiments), even in the presence of malicious clients. Here, ℓ' is the number of corrupt clients who provide malformed inputs during the protocol execution and it does not need to be a priori bounded. In case $\ell' = 0$, then our servers only exchange a pair of hashes. Our communication cost remains low even when a constant fraction (e.g., 10%) of the clients are malicious.

PLASMA framework. We combine these new primitives to construct PLASMA, a protocol for private histograms and heavy hitters in the three-server setting that guarantees security against a malicious server and malicious clients while maintaining low server-to-server communication. PLASMA relies only on efficient hashing and cheap field additions rather than expensive general-purpose MPC or field multiplications. Due to our novel VIDPF primitive, PLASMA outperforms Poplar with regard to runtime by a factor of 5 – 10 \times over WAN for $\mathcal{T} = 1\%$ of the clients. In the same setting, our batched consistency check optimization enables us to drastically outperform both Poplar and the sorting-based protocol of [AHI⁺22] in terms of server-to-server communication by a factor of 35 \times and 45 \times , respectively. For these conditions, we further analyzed the monetary cost of PLASMA, [AHI⁺22], and Poplar and report that PLASMA is more than 2.5 \times and 4 \times cheaper respectively.

Applications. We evaluate PLASMA for two applications: a) detecting frequently visited URLs, and b) identifying popular coordinates.

Popular URLs. A prominent application (discussed both in [AHI+22] and [BBC+21]) is identifying which URLs crash the clients’ browsers more frequently. Each client has a string of n bits that represents the last URL that crashed their browser. In our evaluations (Section 6), we consider $n = 256$ bits, which is sufficient for standard domain names, and compute the heavy hitter URLs that caused more than 1% of client browsers to crash. We perform the task over WAN in approximately 5 minutes for 10^6 clients, while incurring less than 1 GB of server-to-server communication (less than \$1 in total cost).

Popular GPS coordinates. We demonstrate a new application where PLASMA identifies popular geographic locations without sacrificing user privacy. This can be beneficial with traffic avoidance, restaurant recommendations, as well as advertising (e.g., businesses may identify crowded shopping areas and target their marketing efforts), while ensuring the GPS coordinates of the users remain private to the servers. Likewise, ride-sharing services can enhance vehicle distribution in busy areas and proactively dispatch more drivers during rush hour. This is possible by encoding GPS coordinates as 64-bit strings using *plus codes* [LLC19]. We compute the heavy hitter plus codes for a threshold $\mathcal{T} = 1\%$ in under 2 minutes over WAN across 10^6 clients, while incurring very minimal server-to-server communication with \$0.3 in total monetary costs.

Extensions. We also discuss how to extend PLASMA to obtain fairness against a malicious adversary that corrupts one server and an arbitrary number of clients. PLASMA is the first work to consider different thresholds for heavy hitters based on pre-agreed prefixes by the servers, allowing for more elaborate private statistics, such as the GPS application, where different coordinates (e.g. highways and suburban roads) have different congestion thresholds.

1.2 Related Work

We now discuss relevant works for private heavy hitters. They can be classified into four main groups: those based on DPFs, those based on differential privacy (DP), those based on MPC sorting, and finally those based on general-purpose MPC. A comparison of our protocol with related works can be found in Table 1.

DPF-based. Distributed point functions [BGI15] offer a straightforward solution for private histograms but they fail for heavy hitters due to the blowup in key size, as the client would need to send new DPF keys for each level, resulting in $\mathcal{O}(n)$ DPF keys for n levels. This was addressed by Poplar [BBC+21], which uses two non-colluding servers and introduces the notion of IDPFs to allow efficient evaluation of strings based on prefixes by reusing the same DPF key. Poplar’s threat model is robust against malicious clients but remains susceptible to additive attacks by a malicious server. Therefore, as the servers reconstruct the output, a malicious server can add arbitrary noise to the result without the honest server realizing it. The recent works of [DPRS23, MPD+24] propose a framework for secure data aggregation and they improve the clients’ consistency checks in Poplar and Prio [CGB17]. However, their threat model does not address additive attacks from a malicious server either. Adding such security using zero-knowledge [YSWW21, CSW22] is interesting future work. In contrast, PLASMA provides security against both a malicious server and malicious clients by adding one additional server. Also, Poplar still leaks some information about the heavy hitter prefixes to the servers as they reconstruct the roots of the paths before they prune them. PLASMA performs a secure comparison and either keeps the node with its subtree if $\mathcal{T} > count$, or prunes the subtree.

DP-based. There is also a body of work based on local DP and randomized responses for heavy hitters [QYY+16, BNSGT17, ZKM+20]. These techniques use a single server to collect data from clients. Therefore, this method introduces a trade-off between utility and privacy, as it leaks some information about clients’ private data to the server. In contrast, other methods that provide stronger privacy guarantees require at least two non-colluding servers. Notably, secure computation-based solutions can be modified to achieve DP either by using local DP or by adding a smaller amount of noise in MPC and achieving higher data utility while maintaining privacy.

Likewise, bucketization [ACD+21] computes approximate statistics on a permuted version of the clients’ data combined with dummy data that are sampled as differentially private noise. Bucketization ensures

security against malicious clients, and similarly to Poplar, it can only guarantee privacy *without correctness* in the presence of a malicious server. In contrast, PLASMA focuses on exact statistics and provides both correctness and privacy against both malicious clients and one malicious server. Note that PLASMA is compatible with DP as we describe in Appendix F.

Sorting-based. Recent works that rely on secure sorting algorithms construct private heavy-hitter protocols [AHI+22, JKK+22] or private ad attribution measurement [CJK+23] based on the sorted data. They provide security against malicious servers and clients in the three-server setting, where one of the servers can be malicious. These protocols are computationally fast over LAN. However, they perform secure sorting under MPC, and as a result, they incur heavy communication overheads and their performance degrades significantly over realistic WAN networks. Notably, PLASMA achieves a $45\times$ improvement in server-to-server communication compared to [AHI+22] as shown in Fig. 13 for $\mathcal{T} = 1\%$. Moreover, our PLASMA protocol allows different thresholds for heavy hitters based on pre-agreed prefixes (allowing for more elaborate statistics), this is not possible for sorting-based heavy-hitter protocols.

General MPC-based. One could use generic MPC in the honest majority [FLNW17, CGH+18] or dishonest majority setting [Kel20, WSK+23] to compute heavy hitters, but *an efficient representation* of the heavy-hitters problem in terms of addition and multiplication gates is not known. In fact, the work by Böhler and Kerschbaum [BK21] provides a generic MPC-based protocol for computing differentially private heavy hitters. They use MPC frameworks like MP-SPDZ [Kel20] and SCALE-MAMBA [ACC+21] to achieve semi-honest and malicious security, but their solution suffers from high communication and slow runtime.

3-Party Computation based. Multiple customized 3-party protocols [AHI+22, JKK+22] aim to solve the problem of heavy-hitters. These works consider a third server to exploit the faster computation guarantees in the honest majority. Using a third server is a realistic setup and it is widely considered both in the industry and academia as it ensures practical deployments with malicious security. Notable examples include the Interoperable Private Attribution (IPA) proposal by Meta and Mozilla [CJK+23], JP Morgan’s PrimeMatch [PAD+23], NTT’s heavy-hitters protocol [AHI+22], protocols for private advertisement measurement [MMT+24], Duoram [VHG23], Sabre [VSH22], and others. The servers are meant to run across different organizations; for example, they can be hosted by companies and non-profit organizations as mentioned in Google-Apple’s Covid Exposure system [AG21]. Table 1 compares our work with state-of-the-art results.

2 Preliminaries

2.1 Threat Model

Our threat model assumes three non-colluding servers ($\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$) that run the histogram/heavy-hitters protocol, as well as ℓ clients. The clients provide inputs to the servers and the servers do not have any private input. We assume that an adversary \mathcal{A} maliciously corrupts one of the servers and $\ell' < \ell$ clients.

Clients. Malicious clients may try to deviate from the protocol to disproportionately influence the result or even corrupt the output of the protocol. PLASMA is robust against malicious clients and PLASMA servers preemptively reject any malformed client input before incorporating it into the computation. PLASMA preserves the privacy of honest clients when one of the servers is corrupt along with any number of clients.

Servers. Similarly, a malicious server may try to deviate from the protocol and attempt to learn private user inputs; PLASMA always protects input *privacy* against one malicious server. Another possible attack for a malicious server would be to over-influence or corrupt the protocol result. The semi-honest model does not protect correctness against a malicious server, which is problematic in real-world applications, like advertisement measurements [CJK+23] between two companies, where one company may benefit from reporting inflated measurements by introducing undetectable errors. Malicious security ensures that such behaviors are caught and parties are forced to behave honestly, fostering a transparent environment for

computation. Poplar has this limitation while PLASMA protects *correctness*. Hence, PLASMA is *robust* against a malicious server, since it protects both correctness and privacy. Note that in all DPF-based approaches, the servers learn the heavy prefixes, which can be beneficial in some cases (e.g., for detection of a heavy-hitting web domain that contains multiple non-heavy hitting URL errors) but can also be viewed as leakage. However, PLASMA preserves the exact counts of the prefixes.

2.2 Notation

We denote the computational and statistical security parameters by κ and μ , respectively. Let $\text{PRG} : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2(\kappa+1)}$ be a pseudorandom generator and $\text{Convert} : \{0, 1\}^\kappa \rightarrow \mathbb{G}$ be a map converting a κ -bit string to a pseudorandom group element of additive group \mathbb{G} (where $|\mathbb{G}| > \ell$). We use $:=$ for assignment, $\overset{\mathcal{D}}{\leftarrow}$ for sampling from distribution \mathcal{D} , $=$ for checking equality, and \parallel for concatenation. For histograms, we define a public set \mathbf{X} with m n -bit strings as $\mathbf{X} := \{x_1, x_2, \dots, x_m\}$ where the i th string is denoted as x_i for $i \in [m]$ and the j th bit in $x_i \in \{0, 1\}^n$ is denoted as $x_{i,j}$ for $j \in [n]$. We denote the first L bits of x_i as $x_{i,\leq L} := (x_{i,1}, x_{i,2}, \dots, x_{i,L})$ for $L \leq n$. Let \mathcal{S}_b denote the b th server, for $b \in \{0, 1, 2\}$; we consider $b+1 := (b+1) \bmod 3$ and $b+2 := (b+2) \bmod 3$. We assume ℓ clients, each denoted as \mathcal{C}_i for $i \in [\ell]$. For an n -bit string a we represent its bit decomposition as $a_1, \dots, a_n \in \{0, 1\}$. In histograms, each client \mathcal{C}_i has an n -bit input string $\alpha_i \in \mathbf{X}$, for $i \in [\ell]$, while $\alpha_i \in \{0, 1\}^n$ in the case of heavy-hitters. We use $\alpha_{i,1}, \dots, \alpha_{i,n} \in \{0, 1\}$ to denote the bit representation of the client's input α_i .

2.3 Distributed Point Functions (DPF)

Function secret sharing (FSS) [BGI15] enables splitting the output of a function f into additive shares, where each share of the function is represented by a separate key. Each key allows the owner to efficiently generate an additive share of the output $f(x)$ on a given input x . DPFs are a special case of FSS where f is a point function $f_{\alpha,\beta}(x) := \beta$ if $x = \alpha$, or 0 otherwise. A DPF consists of two algorithms: Gen and Eval. The Gen algorithm takes as input the function $f_{\alpha,\beta}$ and outputs two keys key_0 and key_1 . The Eval algorithm evaluates an input x such that $\text{Eval}(0, \text{key}_0, x) + \text{Eval}(1, \text{key}_1, x) = \beta$ for $x = \alpha$, and 0 for $x \neq \alpha$. Privacy ensures (α, β) remains hidden from an adversary in possession of one of the keys (but not both). We discuss DPF, IDPF [BBC+21] and VDPF [dCP22] in Appendix A for completeness.

3 Technical Overview

We recall the histogram and heavy-hitters protocol by Poplar [BBC+21] in Section 3.1. Then, we briefly describe our histogram protocol in Section 3.2 as a stepping stone to our heavy-hitters protocol, which we describe in Sections 3.3 and 3.4.

3.1 Histogram Protocol of Poplar

Poplar first considers the problem of computing private subset histograms. Each client holds an n -bit string α and the servers \mathcal{S}_0 and \mathcal{S}_1 have a small set $\mathbf{X} := \{x_1, x_2, \dots, x_m\}$ of m n -bit strings. Each client secret shares their input $\alpha \in \mathbf{X}$ using a DPF as $(\text{key}_0, \text{key}_1) := \text{DPF.Gen}(1^\kappa, \alpha, 1, \mathbb{G})$. The client sends key_0 to \mathcal{S}_0 and key_1 to \mathcal{S}_1 . Upon receiving the client key, each server \mathcal{S}_b evaluates the DPF on all m strings of \mathbf{X} as $y_b := \{\text{DPF.Eval}(b, \text{key}_b, x_i)\}_{x_i \in \mathbf{X}}$ and computes a vector of output shares $y_b \in \mathbb{F}^m$, for some large enough finite field \mathbb{F} and $m = |\mathbf{X}|$. The servers repeat this for multiple clients and aggregate the y_b vectors in a counter vector Y_b . Finally, the servers exchange Y_0 and Y_1 to compute the output histogram as $Y := Y_0 + Y_1$. This protocol requires the client to communicate one key to each server and the server-to-server communication is independent of the number of clients since Y_0 and Y_1 are aggregated values. This protocol preserves client privacy.

However, a malicious client can double vote by generating the DPF keys maliciously such that it contains more than one non-zero point or the DPF output at α is greater than 1. To tackle this, Poplar introduces a

malicious sketching protocol to ensure that the client inputs are well-formed. It also preserves the client’s privacy against a malicious server. However, Poplar allows a malicious server to add an error to its shares of the output without the honest server realizing it. For instance, say \mathcal{S}_0 is malicious and introduces additive errors (e.g., $\delta \in \mathbb{F}^m$) in $Y'_0 := Y_0 + \delta$. That way, the output Y of the histogram would be biased by δ as $Y := Y'_0 + Y_1 = Y_0 + Y_1 + \delta$. The honest server \mathcal{S}_1 cannot detect such an additive attack, leading to an error in the correctness of the protocol. Moreover, Poplar’s server-to-server communication scales linearly with $\mathcal{O}(\ell)$ due to the malicious sketching protocol.

3.2 Our Basic Histogram Protocol

We address Poplar’s limitations by (1) introducing one additional server, (2) building upon the primitive of verifiable DPF [dCP22] (Appendix A), and (3) introducing novel consistency checks in the three-party setting. We claim the following benefits over Poplar:

- (a) Robustness against a collusion of a malicious server and malicious clients,
- (b) Lightweight consistency checks for malicious behavior (using only symmetric operations and field additions),
- (c) Server-to-server communication depends logarithmically on the total number of clients.

Our work provides the first maliciously secure protocol whose server-to-server communication is logarithmic in the total number of clients ℓ . Our servers communicate $\mathcal{O}(\ell' \log_2 \frac{\ell}{\ell'})$ hashes for consistency checks, where ℓ' is the number of corrupt clients. Similar to Poplar, we ensure input validation against malicious clients (i.e., honest servers preemptively detect inconsistent inputs and discard them). We present the ideas of our histogram protocol, which are crucial for our heavy-hitters protocol in Section 3.4.

Robustness Against a Malicious Server. The histogram protocol of Poplar is not robust against a malicious server. We consider a third server \mathcal{S}_2 to allow an honest majority to obtain security against one malicious server with improved efficiency. Each client runs three DPF sessions, one between each pair of servers, with independent randomness, but the same input α (i.e., the pairwise evaluation of the DPF keys on point α outputs secret shares of one).

However, adding a third server significantly complicates things as we need to ensure consistency between the three sessions. For instance, we need to check that a malicious client submitted the same input α to all three sessions without revealing it. The client sends the DPF keys for the sessions to the servers and each server obtains two keys. Upon obtaining the DPF keys, each server evaluates the DPF on all input points in \mathbf{X} . It is ensured that if the client behaved honestly then at least one of the three sessions will be evaluated honestly since two of the servers are honest. After aggregating all the clients’ inputs, the output histogram is reconstructed across the three sessions. If the output is the same between each pair of servers then the servers behaved honestly and that is considered as the output. If the output is inconsistent across a pair of servers then one of the servers behaves maliciously (by launching an additive attack) and the honest servers abort, which provides robustness against the malicious server.

Reducing Server-to-Server Latency. We empirically observed that the server-to-server latency increases if there is pairwise communication between the three servers for consistency checks. There are three server-to-server sessions for each client, and the third server \mathcal{S}_2 is involved in two of the three sessions: specifically, sessions $\mathcal{S}_1 - \mathcal{S}_2$ and $\mathcal{S}_2 - \mathcal{S}_0$. The client generates $(\text{key}_{(0,1)}, \text{key}_{(1,0)})$ for session $\mathcal{S}_0 - \mathcal{S}_1$, $(\text{key}_{(1,2)}, \text{key}_{(2,1)})$ for session $\mathcal{S}_1 - \mathcal{S}_2$, and $(\text{key}_{(0,2)}, \text{key}_{(2,0)})$ for session $\mathcal{S}_2 - \mathcal{S}_0$. \mathcal{S}_0 receives $\text{key}_{(0,1)}$ and $\text{key}_{(0,2)}$ from the client for sessions $\mathcal{S}_0 - \mathcal{S}_1$ and $\mathcal{S}_2 - \mathcal{S}_0$, respectively. \mathcal{S}_1 receives $\text{key}_{(1,0)}$ for session $\mathcal{S}_0 - \mathcal{S}_1$ and $\text{key}_{(1,2)}$ for $\mathcal{S}_1 - \mathcal{S}_2$, while \mathcal{S}_2 receives $\text{key}_{(2,1)}$ and $\text{key}_{(2,0)}$ for sessions $\mathcal{S}_1 - \mathcal{S}_2$ and $\mathcal{S}_2 - \mathcal{S}_0$, respectively.

In our optimization, instead of running two sessions in each server, we run all three sessions between \mathcal{S}_0 and \mathcal{S}_1 and use \mathcal{S}_2 as the attestation server. By doing that, we significantly reduce the latency due to the synchronization overhead of the three servers. To enable that, our protocol instructs the client to send $\text{key}_{(2,1)}$ to server \mathcal{S}_0 and $\text{key}_{(2,0)}$ to server \mathcal{S}_1 respectively. The key distribution process by the client is illustrated in Fig. 1.




Fig. 1: Distribution of session keys by client C_i .

Our optimization allows S_0 to replicate the computation of S_2 in session $S_1 - S_2$ (because they both have $\text{key}_{(2,1)}$) and S_2 acts as an attestator by just sending hashes to S_1 for the same messages that S_0 should send. These hashes prevent S_0 from acting maliciously. Similar protocol steps are run by S_2 to attest the $S_2 - S_0$ session and prevent S_1 (who is replicating S_2) from acting maliciously. This optimization, shown in Fig. 2, allows us to batch-verify all three sessions as a single session between S_0 and S_1 using hashes.




Fig. 2: Session keys and attestation by S_2 .

Client Input Validation. The above protocol assumes that the client computes the DPF evaluation keys honestly and sends them to the servers. A malicious client could construct malformed DPF keys such that the client’s input gets counted more than once. To prevent this class of attacks, we propose a novel consistency check that only relies on inexpensive symmetric operations, like hashing.

We first ensure that the DPF output is non-zero only at a single point. The work of [dCP22] introduces the primitive of verifiable DPF (VDPF), which we summarize in Appendix A. This is a stronger notion of DPF, where the servers obtain a correctness proof π upon evaluating a pair of DPF keys on a given input point. The two servers obtain the same proof π if the client generates the DPF keys honestly (i.e., the DPF output is non-zero only at a single point α). Multiple proofs corresponding to different evaluation points are batch-verified. Next, we ensure that the DPF output value at the non-zero point is indeed 1. Our protocol instructs the servers to sum up all the output shares (corresponding to each point in \mathbf{X}) of the client and reconstruct the output. If the reconstructed output is not well-formed (i.e., is not 1), then the client’s input is discarded. If the output is 1 (i.e., the client behaved honestly), then the DPF output shares are aggregated by the server in the histogram share.

Client Input Consistency Across Sessions. A malicious client can provide inconsistent inputs across the three server sessions by providing DPF keys for different points α_1, α_2 , and α_3 in each session respectively. The verifiability of the VDPF fails to detect this attack since each individual VDPF in each session is valid.

To address the challenge, we propose a novel consistency check that relies on a single hash verification. Let us denote $\mathbf{Y}_{(0,1)}$, $\mathbf{Y}_{(0,2)}$, and $\mathbf{Y}_{(2,1)}$ be the output of the VDPF evaluation by S_0 on keys $\text{key}_{(0,1)}$, $\text{key}_{(0,2)}$, and $\text{key}_{(2,1)}$ corresponding to sessions $S_0 - S_1$, $S_0 - S_2$, and $S_2 - S_1$, respectively. Similarly, let us denote $\mathbf{Y}_{(1,0)}$, $\mathbf{Y}_{(2,0)}$, and $\mathbf{Y}_{(1,2)}$ be the output of the VDPF evaluation by S_1 on keys $\text{key}_{(1,0)}$, $\text{key}_{(2,0)}$, and $\text{key}_{(1,2)}$ corresponding to sessions $S_0 - S_1$, $S_0 - S_2$, and $S_2 - S_1$, respectively. By definition, reconstructing each pair of secret shared outputs (e.g., $\mathbf{Y}_{(0,1)}$, $\mathbf{Y}_{(1,0)}$) results in a vector of zeros except a single location. Note that the client has also sent $\text{key}_{(2,1)}$ to S_0 and $\text{key}_{(2,0)}$ to S_1 respectively. Server S_0 sends hash $h := H(\mathbf{Y}_{(0,1)} - \mathbf{Y}_{(0,2)} \parallel \mathbf{Y}_{(0,2)} - \mathbf{Y}_{(2,1)})$ to S_1 , who verifies that $h = H(\mathbf{Y}_{(2,0)} - \mathbf{Y}_{(1,0)} \parallel \mathbf{Y}_{(1,2)} - \mathbf{Y}_{(2,0)})$. The verification of the hash h ensures that the client’s input is consistent between: (1) the sessions $S_0 - S_1$ and

$\mathcal{S}_0 - \mathcal{S}_2$, as well as (2) the sessions $\mathcal{S}_0 - \mathcal{S}_2$ and $\mathcal{S}_2 - \mathcal{S}_1$. By transitivity, all three sessions are consistent if the hash verification succeeds. Observe that if the servers acted honestly, $\mathbf{Y}_{(0,1)} + \mathbf{Y}_{(1,0)} = \mathbf{Y}_{(0,2)} + \mathbf{Y}_{(2,0)} = \mathbf{Y}_{(1,2)} + \mathbf{Y}_{(2,1)}$ and thus, $\mathbf{Y}_{(0,1)} - \mathbf{Y}_{(0,2)} = \mathbf{Y}_{(2,0)} - \mathbf{Y}_{(1,0)}$ and $\mathbf{Y}_{(0,2)} - \mathbf{Y}_{(2,1)} = \mathbf{Y}_{(1,2)} - \mathbf{Y}_{(2,0)}$. Our novel check requires additions (without any multiplications) and a cheap hash computation. The communication cost is one hash of size κ bits. This leads to $\mathcal{O}(\kappa\ell)$ server-server communication for ℓ clients, but it is optimized to logarithmic communication by applying batched client verification, described in Section 5. We present the histogram protocol in Appendix H.

3.3 Heavy-Hitters from \mathcal{T} -Prefix Count

Poplar reduced the problem of computing heavy hitters to the problem of computing prefix count queries for a prefix $p \in \{0, 1\}^*$ over client inputs. Then, they implemented prefix count queries by relying on IDPFs (summarized in Appendix A). However, they leak the count of strings that contain the \mathcal{T} heavy-hitting prefix p due to the reliance on a prefix-count query oracle that outputs the count. To mitigate this leakage, we introduce the notion of \mathcal{T} -threshold prefix-count queries that return 1 if at least \mathcal{T} of clients' input strings contain p , otherwise, it returns 0. We define it as:

Definition 1 (\mathcal{T} -Prefix-count Query Oracle $\Omega_{\alpha_1, \dots, \alpha_\ell}(p, \mathcal{T})$). Return 1 (on input prefix $p \in \{0, 1\}^*$) if prefix p appears at least \mathcal{T} times in the clients' input strings $\alpha_1, \alpha_2, \dots, \alpha_\ell \in \{0, 1\}^*$ where client \mathcal{C}_i has input string α_i for $i \in [\ell]$, otherwise, return 0.

\mathcal{T} -Heavy hitters. The \mathcal{T} -heavy hitters algorithm (for threshold \mathcal{T}) is provided with oracle $\Omega_{\alpha_1, \dots, \alpha_\ell}(p, \mathcal{T})$ for computing \mathcal{T} -prefix count for prefix p over the client input strings $\alpha_1, \dots, \alpha_\ell$. The initial prefix is the empty string ϵ . At each level k , it considers the heavy-hitter prefixes $p \in \{0, 1\}^k$ of length k in set HH^k , which contains the list of k -bit strings that appear at least \mathcal{T} times. The algorithm performs a breadth-first search of the prefix tree. It includes $k + 1$ bit length strings $p \parallel 0$ in HH^{k+1} if $p \parallel 0$ occurs at least \mathcal{T} times in the input strings $(\alpha_1, \dots, \alpha_\ell)$, otherwise it gets pruned along its subtree. This is performed by querying the oracle $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel 0, \mathcal{T})$. The same process is repeated for $p \parallel 1$. The algorithm repeats this for all k -bit strings in HH^k (which updates HH^{k+1} based on the search and pruning of set HH^k). At the end of the breadth-first search and pruning, the algorithm outputs the set of strings that are \mathcal{T} -heavy hitters. Our formal algorithm is presented in Fig. 3.

<p>Parameters: Threshold $\mathcal{T} \in \mathbb{N}$ and string length $n \in \mathbb{N}$.</p> <p>Inputs: The algorithm has no explicit input. It has access to t-prefix count query oracle $\Omega_{\alpha_1, \dots, \alpha_\ell}(p, t)$ for securely computing t-prefix-count queries over prefix p for strings $\alpha_1, \dots, \alpha_\ell$.</p> <p>Outputs: The set of \mathcal{T}-heavy-hitter strings in $\alpha_1, \alpha_2, \dots, \alpha_\ell$.</p> <p>Algorithm:</p> <ul style="list-style-type: none"> – Init. $\text{HH}^{\leq n} := \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\} := \{\{\epsilon\}, \emptyset, \dots, \emptyset\}$, where HH^0 contains empty string ϵ and $\text{HH}^1, \dots, \text{HH}^n$ are empty sets. – For each prefix $p \in \text{HH}^k$ of length k-bits in set HH^k (where $k = 0, 1, 2, \dots, n - 1$) and $b \in \{0, 1\}$: <ul style="list-style-type: none"> – If $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel b, \mathcal{T}) = 1$, then $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \{p \parallel b\}$. – Output \mathcal{T}-heavy hitters $\text{HH}^{\leq n} = \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\}$.
--

Fig. 3: Algorithm for computing \mathcal{T} -heavy hitters.

Cost Analysis. There are ℓ input strings in total. For any string of length k , there are at most ℓ/\mathcal{T} candidate heavy hitter strings. At each level k , the algorithm makes at most one oracle query per heavy hitter string. Hence, the algorithm makes at most $n\ell/\mathcal{T}$ prefix-count oracle queries for n levels. If we set the threshold to be a constant fraction of all input strings (e.g., $\mathcal{T} = 0.01\ell$), then the number of prefix-count queries are independent of the number of input strings (e.g., $n\ell/\mathcal{T} = n\ell/0.01\ell = 100n$).

3.4 \mathcal{T} -Prefix Count Queries Oracle from VIDPF

We realize the \mathcal{T} -Prefix Count Query Oracle $\Omega(\cdot, \mathcal{T})$ from Def. 1 by relying on a new verifiable incremental DPF (VIDPF) primitive and using an ideal functionality \mathcal{F}_{CMP} (Fig. 8) for secure comparison.

Verifiable Incremental DPF (VIDPF). A DPF allows a client to succinctly share a vector of size 2^n with a single non-zero point. Meanwhile, an incremental DPF (introduced by Poplar and denoted as IDPF) allows the client to succinctly secret share a path in the binary tree (used for representing 2^n leaves in binary format) and each node in the path can hold non-zero values. Our novel VIDPF primitive offers strong integrity guarantees over IDPFs since the evaluation of the client keys also provides proofs (π_1, \dots, π_n) to the servers ensuring that the VIDPF output is non-zero along a single path in the binary tree. It also allows incremental evaluation of the VIDPF over an input $x \in \{0, 1\}^k$, given state st_b^{k-1} and proof π_b^{k-1} , corresponding to VIDPF evaluation of the first $k-1$ bits of x . The incremental evaluation enables the party possessing key_b to process one level and obtain the secret sharing of output $f(x)$, a new state st_b^k , and a new proof π_b^k corresponding to the VIDPF evaluation of the path involving x . More formally, we capture the high-level ideas of VIDPF using the following two algorithms:

- $\text{Gen}(1^\kappa, 1^n, \alpha, (\beta^1, \beta^2, \dots, \beta^n), \mathbb{G}) \rightarrow (\text{key}_0, \text{key}_1)$: Given security parameter κ , input size n , input string $\alpha \in \{0, 1\}^n$, and values β^1, \dots, β^n , the key generation algorithm outputs two VIDPF keys key_0 and key_1 .
- $\text{EvalPref}(b, \text{key}_b, x, \text{st}_b^{k-1}, \pi_b^{k-1}) \rightarrow (\text{st}_b^k, y_b, \pi_b^k)$: Given a VIDPF key key_b and an input string $x \in \{0, 1\}^k$ of length $k \leq n$ bits, this algorithm outputs an internal state st_b^k , secret-shared value $y_b \in \mathbb{G}$, and a proof $\pi_b^k \in \{0, 1\}^*$.

Correctness of the VIDPF ensures that for all input points $\alpha \in \{0, 1\}^n$, output values $\beta^1, \dots, \beta^n \in \mathbb{G}$, VIDPF keys generated as $(\text{key}_0, \text{key}_1) \leftarrow \text{Gen}(\alpha, \beta^1, \beta^2, \dots, \beta^n, \mathbb{G})$ and all values $x \in \{0, 1\}^k$, where $k \leq n$, the following holds for all $k \leq n$:

$$\pi_0^k = \pi_1^k \text{ and } y = (y_0 + y_1) = \begin{cases} \beta^k, & \text{if } x \text{ is a prefix of } \alpha, \\ 0, & \text{otherwise,} \end{cases}$$

where $(\text{st}_0^k, y_0, \pi_0^k) := \text{EvalPref}(0, \text{key}_0, x, \text{st}_0^{k-1}, \pi_0^{k-1})$ and $(\text{st}_1^k, y_1, \pi_1^k) := \text{EvalPref}(1, \text{key}_1, x, \text{st}_1^{k-1}, \pi_1^{k-1})$. For security guarantees, we require two additional properties from the VIDPF primitive:

- *Input Privacy.* The security of VIDPF guarantees that an adversarial evaluator in possession of either key_0 or key_1 (but not both), does not learn anything about the input α or the outputs β^1, \dots, β^n of the client.
- *Verifiability.* This property states that if two proofs (e.g., π_0^k and π_1^k) are the same, then there is at most one path of length k in the binary tree whose evaluation with $(\text{key}_0, \text{key}_1)$ outputs $(\beta^1, \beta^2, \dots, \beta^k)$. More formally, for any $k \in [n]$ there exists a single k -bit string $\tilde{x} \in \{0, 1\}^k$ such that if $\pi_0^k = \pi_1^k$, then the following holds:

$$\begin{aligned} (\text{st}_0^k, y_0, \pi_0^k) &:= \text{EvalPref}(0, \text{key}_0, z, \text{st}_0^{k-1}, \pi_0^{k-1}) \\ (\text{st}_1^k, y_1, \pi_1^k) &:= \text{EvalPref}(1, \text{key}_1, z, \text{st}_1^{k-1}, \pi_1^{k-1}) \\ y_0 + y_1 &= \begin{cases} \beta^k, & \text{if } z = \tilde{x}, \\ 0, & \text{if } z = \{0, 1\}^k \setminus \{\tilde{x}\}, \end{cases} \end{aligned}$$

where $\text{st}_0^{k-1}, \pi_0^{k-1}$ and $\text{st}_1^{k-1}, \pi_1^{k-1}$ are obtained by recursively running the EvalPref algorithm on $k-1$ bits of z . The evaluators initialize $\text{st}_0^0 := \text{st}_1^0 := 0$ and $\pi_0^0 := \pi_1^0 := 0$. It also implicitly captures the requirement that $\tilde{x} \in \{0, 1\}^{k-1}$ is a prefix of $\tilde{x} \in \{0, 1\}^k$ for $k \in [n]$.

We provide a construction of VIDPF in Figs. 15 and 16 (Appendix B) based on length doubling PRG in the random oracle model. Next, we outline our protocol for securely implementing \mathcal{T} -prefix count queries using VIDPF and the comparison functionality \mathcal{F}_{CMP} .

Implementing \mathcal{T} -Prefix Count Queries. Each client generates three pairs of VIDPF keys, one for each pair of servers, with independent randomness but the same input point α and output values $(1, \dots, 1)$. The client sends the keys for the sessions to the respective servers (Fig. 1) as in our histogram protocol.

Basic Protocol. As depicted in Fig. 2, \mathcal{S}_1 replicates \mathcal{S}_2 in the $\mathcal{S}_2 - \mathcal{S}_0$ session and \mathcal{S}_2 behaves as an attestator for \mathcal{S}_1 by sending hashes of the messages that \mathcal{S}_1 should send. The hash prevents \mathcal{S}_1 from acting maliciously corresponding to the $\mathcal{S}_2 - \mathcal{S}_0$ session. Similar protocol steps are run by \mathcal{S}_2 for the session $\mathcal{S}_1 - \mathcal{S}_2$, where \mathcal{S}_2 sends hashes to \mathcal{S}_1 . Hence, \mathcal{S}_0 and \mathcal{S}_1 run three sessions, and \mathcal{S}_2 runs two of those sessions in parallel. Next, we describe the protocol to compute a \mathcal{T} -prefix count query on a string $p \parallel 0 \in \{0, 1\}^k$ (note, the same process can be repeated for $p \parallel 1$). \mathcal{S}_0 and \mathcal{S}_1 evaluate the VIDPF keys for the three sessions on $p \parallel 0$ and obtain a secret share of the output $y^{p \parallel 0}$ and proof π . For an honest client, $y^{p \parallel 0}$ should be $\beta^k = 1$. However, a malicious client can construct malformed keys such that the client’s input gets counted more than once.

Client Input Validation. We introduce the following consistency checks to validate a client’s input. Checks 1-3 ensure that the VIDPF keys are “one-hot”, i.e., they have a single non-zero evaluation path (containing 1 in this case, along the path), and check 4 ensures that the client input is consistent across the sessions:

Check 1: The servers \mathcal{S}_0 and \mathcal{S}_1 first verify that the proofs π are the same for all three sessions. This ensures that there is at most one path in the binary tree that is non-zero.

Check 2: For the root level (i.e., $k = 0$), the servers evaluate the VIDPF keys on the empty string ϵ and verify it is 1.

Check 3: Finally, at the k^{th} level, the servers need to verify that $y^{p \parallel 0}$ is either 0 or 1, without reconstructing the output. We perform this check by observing that the output of the parent p should be the sum of the outputs of $p \parallel 0$ and $p \parallel 1$. The servers evaluate the VIDPF keys on the parent string p and sibling (of $p \parallel 0$) string $p \parallel 1$ to obtain secret shares of the output of y^p and $y^{p \parallel 1}$ respectively. The servers reconstruct $y^p - (y^{p \parallel 0} + y^{p \parallel 1})$ and verify that it is 0. The first check ensures that at most one of $y^{p \parallel 0}$ or $y^{p \parallel 1}$ is non-zero. Combining the two checks, we can conclude that either $(y^{p \parallel 0} = 0, y^{p \parallel 1} = 1)$ or $(y^{p \parallel 0} = 1, y^{p \parallel 1} = 0)$, since at most one child can equal 1 when the parent holds a value of 1. Iterating this for all k levels ensures that $y^{p \parallel 0} = 1$ iff $y^p = 1$ and $y^{p \parallel 1} = 0$, else $y^{p \parallel 0} = 0$. The servers also verify (using check 1) the corresponding proofs π generated during the VIDPF evaluation along the path, to ensure there is at most one non-zero path in the entire binary tree.

Check 4: The servers also need to ensure that the client input is consistent across the sessions. This is ensured by computing the difference of the reconstructed outputs across the sessions and verifying that they are equal to 0 by matching their hash values. For more details, we defer to Section 4.

Output Phase. Once the client’s VIDPF output $y^{p \parallel 0}$ is verified, the secret shares of $y^{p \parallel 0}$ are aggregated into counter $\text{cnt}^{p \parallel 0}$. The servers repeat the above steps for all the clients in parallel to obtain secret shares of $\text{cnt}^{p \parallel 0}$. The servers invoke the comparison functionality \mathcal{F}_{CMP} (Fig. 8) with the secret shares of cnt and threshold \mathcal{T} . \mathcal{F}_{CMP} reconstructs cnt and it outputs 1 if $\text{cnt} \geq \mathcal{T}$, otherwise, it outputs 0. This is returned by the servers as the output of the \mathcal{T} -prefix count oracle query response to the string $p \parallel 0$. Similar steps are run for $p \parallel 1$. The comparison functionality \mathcal{F}_{CMP} is securely implemented using the state-of-the-art protocol of Rabbit [MRVW21].

Robustness Against a Malicious Server. Note that the above validation check assumes that both servers are honest. Otherwise, malicious behaviour is detected as described next. The third server ensures that if the client behaves honestly then at least one of the three sessions will be evaluated correctly since two of the servers are honest. After aggregating all the client’s inputs, cnt is reconstructed across the three sessions by \mathcal{F}_{CMP} . If cnt is inconsistent across any pair of servers then \mathcal{F}_{CMP} returns \perp indicating that one of the servers behaved maliciously by launching an additive attack. This causes the honest servers to abort, providing robustness against the malicious server. We observe that our protocol satisfies *fairness* (which is a stronger security notion than selective abort) if \mathcal{F}_{CMP} is implemented using a fair protocol. We discuss this in Sec. 7.

Batched Client Verification. In our final protocol, we verify multiple client inputs at each level in one batch. We batch all the clients’ VIDPF evaluations using a Merkle tree that has ℓ leaves for ℓ clients. First, the servers check the equality of ℓ leaves by asserting that the Merkle roots are the same. If the roots match then the leaves are the same, while if they differ then the servers recursively repeat the same process for each

of the two children of the parent node. Proceeding this way, the servers identify the malformed leaves on which the two trees differ. This reduces the dependency of our server-to-server communication to $\mathcal{O}(\ell'(\log_2 \frac{\ell}{\ell'}))$, for ℓ' malicious clients, instead of $\mathcal{O}(\ell)$, while when $\ell' = 0$ our communication is down to $\mathcal{O}(1)$. Formal details can be found in Section 5.

4 Private Heavy Hitters

We provide the ideal functionality \mathcal{F}_{HH} for heavy-hitters between three servers and ℓ clients in Fig. 4. Adversary \mathcal{A} maliciously corrupts any one of the servers and multiple clients. Note that this corruption can easily happen; if \mathcal{A} has maliciously corrupted a server, then \mathcal{A} can spawn multiple malicious clients. Additionally, if \mathcal{A} controls a server, it can instruct \mathcal{F}_{HH} to discard an honest client's input. It can also instruct the functionality to abort at a particular level $k + 1$. In this case, \mathcal{A} and the honest servers receive the set of all (that have not been discarded by \mathcal{A}) k -bit heavy-hitting prefixes as output, and the functionality instructs the honest servers to abort. We remark that \mathcal{F}_{HH} never leaks an honest client's inputs.

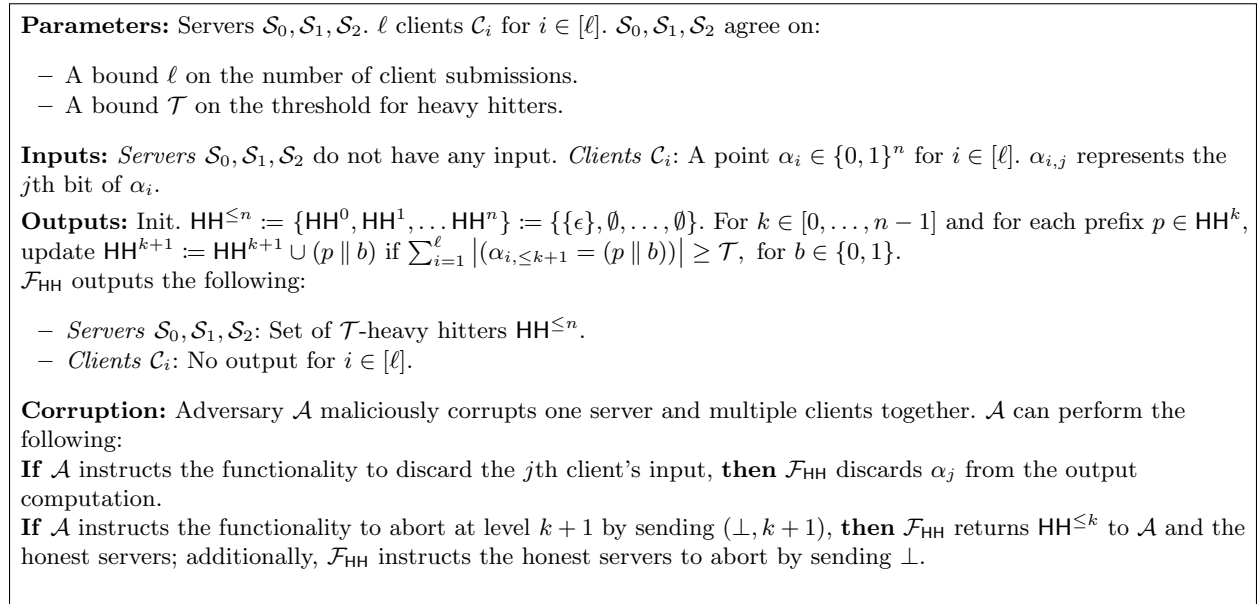


Fig. 4: The ideal \mathcal{F}_{HH} functionality for \mathcal{T} -heavy hitters.

Our detailed protocol π_{HH} that implements \mathcal{F}_{HH} appears in Figs. 5 and 6, while high-level ideas of our protocol can be found in Sections 3.3 and 3.4. Our π_{HH} protocol privately computes all the \mathcal{T} -heavy-hitting strings (and their heavy-hitting prefixes) given the input data of ℓ clients, while protecting the privacy of the individual data points. π_{HH} runs on three servers ($\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$) that utilize our verifiable incremental DPF (VIDPF) protocol to privately aggregate the clients' data points. Specifically, π_{HH} runs three VIDPF sessions, which guarantees security against a malicious server. Our protocol proceeds in three phases: a client computation phase, a server computation phase, and an output phase.

Client Computation. During the client computation phase, each client \mathcal{C} prepares three pairs of VIDPF keys for their private data point $\alpha \in \{0, 1\}^n$, and output value $(\beta^1, \dots, \beta^n) := (1, \dots, 1)$ along the path to α , using independent randomness for each key generation. Employing three pairs of keys essentially allows us to run three separate VIDPF sessions. \mathcal{S}_0 and \mathcal{S}_1 each have one key for each of the three sessions, while \mathcal{S}_2 acts as a consistency checking server and shares one key with each of the other two servers. More specifically, the client generates $(\text{key}_{(0,1)}, \text{key}_{(0,2)})$ for \mathcal{S}_0 , $(\text{key}_{(1,0)}, \text{key}_{(1,2)})$ for \mathcal{S}_1 , and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ for \mathcal{S}_2 . The client

Input: Each client \mathcal{C}_i has an input point $\alpha_i \in \mathbf{X}$ for $i \in [\ell]$.

Output: The servers \mathcal{S}_b (for $b \in \{0, 1, 2\}$) output the set of \mathcal{T} -heavy hitters $\text{HH}^{\leq n} := \mathcal{F}_{\text{HH}}(\ell, \mathcal{T}, \{\alpha_i\}_{i \in [\ell]})$.

Primitive: VIDPF := (Gen, EvalPref, EvalNext) is a verifiable incremental DPF. $H_1, H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ are random oracles.

Client \mathcal{C} Computation. (Repeated for ℓ clients, each of which has their own private input α)

- Client \mathcal{C} with input α prepares three pairs DPF keys with independent randomness $u, v, w \xleftarrow{r} \{0, 1\}^\kappa$, as follows:

$$\begin{aligned} (\text{key}_{(0,1)}, \text{key}_{(1,0)}) &:= \text{Gen}(1^\kappa, 1^n, \alpha, (1, \dots, 1), \mathbb{G}), & (\text{key}_{(1,2)}, \text{key}_{(2,1)}) &:= \text{Gen}(1^\kappa, 1^n, \alpha, (1, \dots, 1), \mathbb{G}), \\ (\text{key}_{(2,0)}, \text{key}_{(0,2)}) &:= \text{Gen}(1^\kappa, 1^n, \alpha, (1, \dots, 1), \mathbb{G}) \end{aligned}$$

- The client sends $(\text{key}_{(0,1)}, \text{key}_{(0,2)}, \text{key}_{(2,1)})$ to \mathcal{S}_0 , $(\text{key}_{(1,0)}, \text{key}_{(1,2)}, \text{key}_{(2,0)})$ to \mathcal{S}_1 and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ to \mathcal{S}_2 .

Server Computation.

Each server \mathcal{S}_b initializes $\text{HH}_b^{\leq n} = \{\text{HH}_b^0, \text{HH}_b^1, \dots, \text{HH}_b^n\} := \{\{\epsilon\}, \emptyset, \dots, \emptyset\}$. Repeat the following steps for length of k bits, where $k \in [0, \dots, n-1]$:

- Initialization.** For prefix $p \in \text{HH}_b^k$, servers initialize the aggregation variables for prefixes $\gamma \in \{p \parallel 0, p \parallel 1\}$ as follows:

$$\begin{aligned} \mathcal{S}_0 \text{ sets } \text{cnt}_{(0,1)}^\gamma &:= \text{cnt}_{(0,2)}^\gamma := \text{cnt}_{(2,1)}^\gamma := 0, \\ \mathcal{S}_1 \text{ sets } \text{cnt}_{(1,2)}^\gamma &:= \text{cnt}_{(1,0)}^\gamma := \text{cnt}_{(2,0)}^\gamma := 0, & \mathcal{S}_2 \text{ sets } \text{cnt}_{(2,0)}^\gamma &:= \text{cnt}_{(2,1)}^\gamma := 0 \end{aligned}$$

- VIDPF Evaluation.** For prefix $p \in \text{HH}_b^{\leq k}$, Server \mathcal{S}_b computes: (Repeated for ℓ clients)

- If $(p = \emptyset)$: then \mathcal{S}_0 sets $\text{st}_{(0,1)}^\emptyset := \pi_{(0,1)}^\emptyset := \text{st}_{(0,2)}^\emptyset := \pi_{(0,2)}^\emptyset := \text{st}_{(2,1)}^\emptyset := \pi_{(2,1)}^\emptyset := \emptyset$, \mathcal{S}_1 sets $\text{st}_{(1,2)}^\emptyset := \pi_{(1,2)}^\emptyset := \text{st}_{(1,0)}^\emptyset := \pi_{(1,0)}^\emptyset := \text{st}_{(2,0)}^\emptyset := \pi_{(2,0)}^\emptyset := \emptyset$. \mathcal{S}_2 sets $\text{st}_{(2,0)}^\emptyset := \pi_{(2,0)}^\emptyset := \text{st}_{(2,1)}^\emptyset := \pi_{(2,1)}^\emptyset := \emptyset$.

If $(p \neq \emptyset)$: then \mathcal{S}_b retrieves the state from memory corresponding to the internal states of π_{VIDPF} for prefix p : \mathcal{S}_0 retrieves $(\text{st}_{(0,1)}^p, y_{(0,1)}^p, \pi_{(0,1)}^p)$, $(\text{st}_{(0,2)}^p, y_{(0,2)}^p, \pi_{(0,2)}^p)$ and $(\text{st}_{(2,1)}^p, y_{(2,1)}^p, \pi_{(2,1)}^p)$. \mathcal{S}_1 retrieves $(\text{st}_{(1,2)}^p, y_{(1,2)}^p, \pi_{(1,2)}^p)$, $(\text{st}_{(1,0)}^p, y_{(1,0)}^p, \pi_{(1,0)}^p)$ and $(\text{st}_{(2,0)}^p, y_{(2,0)}^p, \pi_{(2,0)}^p)$. \mathcal{S}_2 retrieves $(\text{st}_{(2,0)}^p, y_{(2,0)}^p, \pi_{(2,0)}^p)$ and $(\text{st}_{(2,1)}^p, y_{(2,1)}^p, \pi_{(2,1)}^p)$.

- Each server \mathcal{S}_b evaluates the VIDPF on the prefixes $\gamma \in \{p \parallel 0, p \parallel 1\}$ as follows and stores them in memory:

$$\begin{aligned} \mathcal{S}_0 \text{ sets } (\text{st}_{(0,1)}^\gamma, y_{(0,1)}^\gamma, \pi_{(0,1)}^\gamma) &:= \text{EvalPref}(0, \text{key}_{(0,1)}, \gamma, \text{st}_{(0,1)}^p, k, \pi_{(0,1)}^p), \\ (\text{st}_{(0,2)}^\gamma, y_{(0,2)}^\gamma, \pi_{(0,2)}^\gamma) &:= \text{EvalPref}(1, \text{key}_{(0,2)}, \gamma, \text{st}_{(0,2)}^p, k, \pi_{(0,2)}^p) \text{ and stores them in memory.} \\ \mathcal{S}_1 \text{ sets } (\text{st}_{(1,2)}^\gamma, y_{(1,2)}^\gamma, \pi_{(1,2)}^\gamma) &:= \text{EvalPref}(0, \text{key}_{(1,2)}, \gamma, \text{st}_{(1,2)}^p, k, \pi_{(1,2)}^p), \\ (\text{st}_{(1,0)}^\gamma, y_{(1,0)}^\gamma, \pi_{(1,0)}^\gamma) &:= \text{EvalPref}(1, \text{key}_{(1,0)}, \gamma, \text{st}_{(1,0)}^p, k, \pi_{(1,0)}^p) \text{ and stores them in memory.} \\ \mathcal{S}_2 \text{ and } \mathcal{S}_1 \text{ set } (\text{st}_{(2,0)}^\gamma, y_{(2,0)}^\gamma, \pi_{(2,0)}^\gamma) &:= \text{EvalPref}(0, \text{key}_{(2,0)}, \gamma, \text{st}_{(2,0)}^p, k, \pi_{(2,0)}^p) \text{ and store them in memory.} \\ \mathcal{S}_2 \text{ and } \mathcal{S}_0 \text{ set } (\text{st}_{(2,1)}^\gamma, y_{(2,1)}^\gamma, \pi_{(2,1)}^\gamma) &:= \text{EvalPref}(1, \text{key}_{(2,1)}, \gamma, \text{st}_{(2,1)}^p, k, \pi_{(2,1)}^p) \text{ and store them in memory.} \end{aligned}$$

- If $k = 1$: Servers compute the proof that the VIDPF evaluation at the root layer sums up to 1:

$$\begin{aligned} \mathcal{S}_0 \text{ sets } h_{(0,1)}^0 &:= H_1(\emptyset, 1 - y_{(0,1)}^0 - y_{(0,1)}^1) \text{ and } h_{(0,2)}^0 := H_1(\emptyset, y_{(0,2)}^0 + y_{(0,2)}^1), \\ \mathcal{S}_1 \text{ sets } h_{(1,2)}^0 &:= H_1(\emptyset, 1 - y_{(1,2)}^0 - y_{(1,2)}^1) \text{ and } h_{(1,0)}^0 := H_1(\emptyset, y_{(1,0)}^0 + y_{(1,0)}^1), \\ \mathcal{S}_2 \text{ and } \mathcal{S}_1 \text{ set } h_{(2,0)}^0 &:= H_1(\emptyset, 1 - y_{(2,0)}^0 - y_{(2,0)}^1), & \mathcal{S}_2 \text{ and } \mathcal{S}_0 \text{ set } h_{(2,1)}^0 &:= H_1(\emptyset, y_{(2,1)}^0 - y_{(2,1)}^1). \end{aligned}$$

Fig. 5: Private \mathcal{T} -Heavy Hitters Protocol π_{HH} (continues in Fig. 6).

Server Computation (Continued from Fig. 5) Repeat the following steps for length of k bits, where $k \in [n]$:

2. **VIDPF Evaluation (Cont.)**. For prefix $p \in \text{HH}_b^{\leq k}$, Server \mathcal{S}_b computes: (Repeated for ℓ clients)

(d) If $k \neq 1$: Servers compute proof that (VIDPF output on prefix p) = (VIDPF output on prefix $p \parallel 0$) + (VIDPF output on prefix $p \parallel 1$):

$$\mathcal{S}_0 \text{ sets } h_{(0,1)}^p := \text{H}_1(p, y_{(0,1)}^p - y_{(0,1)}^{p\parallel 0} - y_{(0,1)}^{p\parallel 1}) \text{ and } h_{(0,2)}^p := \text{H}_1(p, -(y_{(0,2)}^p - y_{(0,2)}^{p\parallel 0} - y_{(0,2)}^{p\parallel 1}))$$

$$\mathcal{S}_1 \text{ sets } h_{(1,2)}^p := \text{H}_1(p, y_{(1,2)}^p - y_{(1,2)}^{p\parallel 0} - y_{(1,2)}^{p\parallel 1}) \text{ and } h_{(1,0)}^p := \text{H}_1(p, -(y_{(1,0)}^p - y_{(1,0)}^{p\parallel 0} - y_{(1,0)}^{p\parallel 1}))$$

$$\mathcal{S}_2 \text{ and } \mathcal{S}_1 \text{ set } h_{(2,0)}^p := \text{H}_1(p, y_{(2,0)}^p - y_{(2,0)}^{p\parallel 0} - y_{(2,0)}^{p\parallel 1}), \quad \mathcal{S}_2 \text{ and } \mathcal{S}_0 \text{ set } h_{(2,1)}^p := \text{H}_1(p, -(y_{(2,1)}^p - y_{(2,1)}^{p\parallel 0} - y_{(2,1)}^{p\parallel 1})).$$

(e) \mathcal{S}_0 and \mathcal{S}_1 ensure that the client input is consistent across the three sessions by computing the following hashes.

$$\mathcal{S}_0 \text{ computes } \widehat{h^{p\parallel 0}} = \text{H}_1(y_{(0,1)}^{p\parallel 0} - y_{(0,2)}^{p\parallel 0}, y_{(0,2)}^{p\parallel 0} - y_{(2,1)}^{p\parallel 0}) \text{ and } \widehat{h^{p\parallel 1}} = \text{H}_1(y_{(0,1)}^{p\parallel 1} - y_{(0,2)}^{p\parallel 1}, y_{(0,2)}^{p\parallel 1} - y_{(2,1)}^{p\parallel 1}).$$

$$\mathcal{S}_1 \text{ computes } \overline{h^{p\parallel 0}} := \text{H}_1(y_{(2,0)}^{p\parallel 0} - y_{(1,0)}^{p\parallel 0}, y_{(1,2)}^{p\parallel 0} - y_{(2,0)}^{p\parallel 0}) \text{ and } \overline{h^{p\parallel 1}} := \text{H}_1(y_{(2,0)}^{p\parallel 1} - y_{(1,0)}^{p\parallel 1}, y_{(1,2)}^{p\parallel 1} - y_{(2,0)}^{p\parallel 1})$$

(f) *Client State Accumulation*: The servers accumulate their local state for each client session as follows:

$$\mathcal{S}_0 \text{ sets } R_{(0,1)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(0,1)}^p, \pi_{(0,1)}^{p\parallel 0}, \pi_{(0,1)}^{p\parallel 1}) \right. \right) \text{ and } R_{(0,2)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(0,2)}^p, \pi_{(0,2)}^{p\parallel 0}, \pi_{(0,2)}^{p\parallel 1}) \right. \right)$$

$$\mathcal{S}_1 \text{ sets } R_{(1,2)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(1,2)}^p, \pi_{(1,2)}^{p\parallel 0}, \pi_{(1,2)}^{p\parallel 1}) \right. \right) \text{ and } R_{(1,0)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(1,0)}^p, \pi_{(1,0)}^{p\parallel 0}, \pi_{(1,0)}^{p\parallel 1}) \right. \right)$$

$$\mathcal{S}_2, \mathcal{S}_1 \text{ set } R_{(2,0)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(2,0)}^p, \pi_{(2,0)}^{p\parallel 0}, \pi_{(2,0)}^{p\parallel 1}) \right. \right),$$

$$\text{and } \mathcal{S}_2, \mathcal{S}_0 \text{ set } R_{(2,1)}^k := \text{H}_2(\left| \left|_{p \in \text{HH}^k} (p, h_{(2,1)}^p, \pi_{(2,1)}^{p\parallel 0}, \pi_{(2,1)}^{p\parallel 1}) \right. \right)$$

3. **Batch-Verification**. The servers batch-verify the client inputs for all three sessions and across the three sessions by invoking π_{check} (Fig. 9):

(a) \mathcal{S}_0 sets $u_i := \{(R_{(0,1)}^k, R_{(0,2)}^k, R_{(2,1)}^k, \widehat{h^{p\parallel 0}}, \widehat{h^{p\parallel 1}})\}$ values for client $i \in [\ell]$. \mathcal{S}_1 sets

$v_i := \{(R_{(1,0)}^k, R_{(2,0)}^k, R_{(1,2)}^k, \overline{h^{p\parallel 0}}, \overline{h^{p\parallel 1}})\}$ values for client $i \in [\ell]$. \mathcal{S}_0 sets $\mathbf{u} := \{u_i\}_{i \in [\ell]}$ and \mathcal{S}_1 sets

$\mathbf{v} := \{v_i\}_{i \in [\ell]}$. \mathcal{S}_0 and \mathcal{S}_1 batch-verify all the client inputs by computing the bit ver and list \mathbf{L} (comprising of invalid client inputs) by running π_{check} with inputs \mathbf{u} and \mathbf{v} respectively: $(\text{ver}, \mathbf{L}) := \pi_{\text{check}}(\mathbf{u}, \mathbf{v})$:

$\text{ver} := 0$ if \exists a client whose $(R_{(0,1)}^k \neq R_{(1,0)}^k) \vee (R_{(0,2)}^k \neq R_{(2,0)}^k) \vee (R_{(2,1)}^k \neq R_{(1,2)}^k) \vee (\widehat{h^{p\parallel 0}} \neq$

$\overline{h^{p\parallel 0}}) \vee (\widehat{h^{p\parallel 1}} \neq \overline{h^{p\parallel 1}})$, and $\text{List } \mathbf{L} := \{\text{list of invalid clients' since they failed to pass the above check}\}$. If

$\text{ver} = 1$, then all the clients' inputs are valid.

(b) \mathcal{S}_2 possesses $R_{(2,0)}^k, R_{(2,1)}^k$ values for each client. \mathcal{S}_2 verifies that \mathcal{S}_2 's version of $R_{(2,1)}^k$ matches with \mathcal{S}_0 's version of $R_{(2,1)}^k$. \mathcal{S}_2 also attests that \mathcal{S}_2 's version of $R_{(2,0)}^k$ matches with \mathcal{S}_0 's version of $R_{(2,0)}^k$ by computing $(\text{ver}', \mathbf{L}') := \pi_{\text{check}}(\{R_{(2,1)}^k, R_{(2,0)}^k\}_{\ell \text{ clients of } \mathcal{S}_2}, \{R_{(2,1)}^k, R_{(2,0)}^k\}_{\ell \text{ clients of } \mathcal{S}_0})$.

(c) \mathcal{S}_2 verifies that \mathcal{S}_2 's version of $R_{(2,0)}^k$ matches with \mathcal{S}_1 's version of $R_{(2,0)}^k$. \mathcal{S}_2 also attests that \mathcal{S}_2 's version of $R_{(2,1)}^k$ matches with \mathcal{S}_1 's version of $R_{(1,2)}^k$ by computing

$$(\text{ver}'', \mathbf{L}'') := \pi_{\text{check}}(\{R_{(2,0)}^k, R_{(2,1)}^k\}_{\ell \text{ clients of } \mathcal{S}_2}, \{R_{(2,0)}^k, R_{(1,2)}^k\}_{\ell \text{ clients of } \mathcal{S}_0}).$$

After batch verification, the servers identify the list of bad clients as $\mathbf{L} := \mathbf{L} \cup \mathbf{L}' \cup \mathbf{L}''$. The servers ignore the inputs of all clients in \mathbf{L} .

Fig. 6: **Private \mathcal{T} -Heavy Hitters Protocol π_{HH}** (continues in Fig. 7).

Server Computation (Continued from Fig. 6) Repeat the following steps for length of k bits, where $k \in [n]$:

4. **Aggregation.** Aggregate the VIDPF outputs for prefixes $\gamma \in \{p \parallel 0, p \parallel 1\}$ as follows: **(Repeated for all validated clients in $[\ell] \setminus \mathbf{L}$)**

$$\mathcal{S}_0 \text{ sets } \text{cnt}_{(0,1)}^\gamma := \text{cnt}_{(0,1)}^\gamma + y_{(0,1)}^\gamma, \text{cnt}_{(0,2)}^\gamma := \text{cnt}_{(0,2)}^\gamma + y_{(0,2)}^\gamma, \text{ and } \text{cnt}_{(2,1)}^\gamma := \text{cnt}_{(2,1)}^\gamma + y_{(2,1)}^\gamma$$

$$\mathcal{S}_1 \text{ sets } \text{cnt}_{(1,2)}^\gamma := \text{cnt}_{(1,2)}^\gamma + y_{(1,2)}^\gamma, \text{cnt}_{(1,0)}^\gamma := \text{cnt}_{(1,0)}^\gamma + y_{(1,0)}^\gamma, \text{ and } \text{cnt}_{(2,0)}^\gamma := \text{cnt}_{(2,0)}^\gamma + y_{(2,0)}^\gamma$$

$$\mathcal{S}_2 \text{ sets } \text{cnt}_{(2,0)}^\gamma := \text{cnt}_{(2,0)}^\gamma + y_{(2,0)}^\gamma \text{ and } \text{cnt}_{(2,1)}^\gamma := \text{cnt}_{(2,1)}^\gamma + y_{(2,1)}^\gamma$$

The servers have aggregated the VIDPF evaluations (over all the ℓ clients) for all candidate $(k+1)$ -bit strings.

5. **Pruning.** For every $(k+1)$ -bit string γ , the servers invoke \mathcal{F}_{CMP} functionality (Fig. 8) with the additive shares of the node frequency.

$$\mathcal{S}_0 \text{ invokes } \mathcal{F}_{\text{CMP}}(\text{cnt}_{(0,1)}^\gamma, 0, \text{cnt}_{(0,2)}^\gamma, \text{cnt}_{(2,1)}^\gamma, \text{cnt}_{(0,2)}^\gamma, \mathcal{T}),$$

$$\mathcal{S}_1 \text{ invokes } \mathcal{F}_{\text{CMP}}(\text{cnt}_{(1,0)}^\gamma, \text{cnt}_{(1,2)}^\gamma, 0, \text{cnt}_{(1,2)}^\gamma, \text{cnt}_{(2,0)}^\gamma, \mathcal{T}),$$

$$\mathcal{S}_2 \text{ invokes } \mathcal{F}_{\text{CMP}}(0, \text{cnt}_{(2,1)}^\gamma, \text{cnt}_{(2,0)}^\gamma, 0, 0, \mathcal{T})$$

The servers abort if \mathcal{F}_{CMP} aborts. If \mathcal{F}_{CMP} outputs 1 set $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \gamma$. Otherwise, the servers ignore γ since it is non-heavy hitter.

Servers have successfully computed the HH^{k+1} set. Servers repeat “*Server Computation*” steps (starting from Step 2b) on $k+1$ bit prefixes.

Output Phase. The servers output $\text{HH}^{\leq n}$ as the set of \mathcal{T} -heavy hitter strings.

Fig. 7: **Private \mathcal{T} -Heavy Hitters Protocol π_{HH}** (continuing from Fig. 6).

sends $(\text{key}_{(0,1)}, \text{key}_{(0,2)}, \text{key}_{(2,1)})$ to \mathcal{S}_0 , $(\text{key}_{(1,0)}, \text{key}_{(1,2)}, \text{key}_{(2,0)})$ to \mathcal{S}_1 , and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ to \mathcal{S}_2 as shown in Fig. 1.

Server Computation. Each server initializes a set of sets for heavy-hitters as $\text{HH}^{\leq n} := \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\} := \{\{\epsilon\}, \emptyset, \dots, \emptyset\}$, where HH^0 is a set with the empty string ϵ , $\text{HH}^1, \dots, \text{HH}^n$ are empty sets and HH^k corresponds to the k th level. The servers start accepting VIDPF keys from the clients. As in our histogram protocol, \mathcal{S}_2 acts as an attesting server for the sessions involving keys $\text{key}_{(2,0)}$ and $\text{key}_{(2,1)}$ by sending hashes (depicted in Fig. 2). Next, for $k \in [n]$ the servers perform the following:

- (a) *Initialization.* For each k -bit heavy-hitting prefix $p \in \text{HH}^k$, the servers initialize to 0 a $\text{cnt}^{p \parallel 0}$ (resp. $\text{cnt}^{p \parallel 1}$) variable for each session to count the frequency of prefix $p \parallel 0$ (resp. $p \parallel 1$). Each server aggregates for each of the three sessions their additive shares of each frequency in their local cnt variables and uses them for pruning.
- (b) *VIDPF Evaluation.* Next, the servers retrieve from memory the states for VIDPF evaluation in all three sessions corresponding to prefix $p \in \{0, 1\}^k$ for each client. These states are used to incrementally evaluate the VIDPF on prefix strings $\gamma \in \{p \parallel 0, p \parallel 1\}$ for every client in all three sessions. For each client, the servers obtain new evaluation states (corresponding to prefix γ), VIDPF output for prefix string γ , and proof strings. The states are stored in memory for future VIDPF evaluations on $\gamma \parallel 0$ and $\gamma \parallel 1$ in the $(k+1)^{\text{th}}$ level. More formally, the servers compute a secret shared vector $y_{(b_1, b_2)}^\gamma$ and a hash $\pi_{(b_1, b_2)}^\gamma$ that is used for consistency checking by relying on the verifiability property of the VIDPF. Next, the servers validate the client’s input. If $k = 1$, then the servers reconstruct $y^0 + y^1$ for each client to verify that $y^0 + y^1 = 1$ (i.e., the non-zero root value is 1). If $k \neq 1$, then the servers reconstruct $y^p - (y^{p \parallel 0} + y^{p \parallel 1})$ and verify that it is 0, asserting that the parent value is propagated to the children correctly. Note that in either of the above cases, nothing is leaked about the client’s input, apart from the fact that it is a valid submission (i.e., 1 at the root layer and correct propagation). This ensures that the subtrees involving $p \parallel 0$ and $p \parallel 1$ are valid. The servers also need to ensure that the client has

provided consistent input across the three sessions. This is ensured by computing the difference of the reconstructed outputs across the sessions and verifying that they equal 0 by matching their hash values with the other servers' hash in Step 2e of Fig. 5.

- (c) *Batch-Verification.* The servers need to check: (1) that the hashes they possess for a client are equal, and (2) that $y^p = (y^{p||0} + y^{p||1})$. Both these checks are reduced to checking the equality of a string (corresponding to each client) held by servers. Let \mathbf{u} (resp. \mathbf{v}) be the list of ℓ (one for each client) strings held by the first (resp. second) server. Then, the servers perform a batch verification of \mathbf{u} and \mathbf{v} strings by invoking the subprotocol $\pi_{\text{check}}(\mathbf{u}, \mathbf{v})$ in Fig. 9. If the two lists \mathbf{u} and \mathbf{v} are equal then π_{check} returns $\text{ver} = 1$, else it returns $\text{ver} = 0$ and a list \mathbf{L} containing the indices of elements where the lists differ. This is performed for all three sessions. \mathcal{S}_2 also attests to the sessions that it is involved in. This is performed using batch-verification, yielding output lists \mathbf{L}' and \mathbf{L}'' . Finally, the servers identify the list of bad clients as $\mathbf{L} = \mathbf{L} \cup \mathbf{L}' \cup \mathbf{L}''$ and their VIDPF output is ignored. The servers consider the rest of the clients as “validated” and they are moved to the aggregation phase.
- (d) *Aggregation.* Once a client's VIDPF output y^γ is validated for $\gamma \in \{p || 0, p || 1\}$, it is aggregated into $\text{cnt}^\gamma := \text{cnt}^\gamma + y^\gamma$. This is locally performed by each server (for all three sessions) using the secret shares of y^γ since it only involves addition. The servers perform this over every validated client output, and at the end of this phase, the servers possess a secret share of the frequency of $p || 0$ and $p || 1$ as $\text{cnt}^{p||0}$ and $\text{cnt}^{p||1}$.

Inputs: Party P_0 has input $(a_0, b_0, c_0, d_0, e_0, \mathcal{T}_0)$, Party P_1 has input $(a_1, b_1, c_1, d_1, e_1, \mathcal{T}_1)$, and Party P_2 has input $(a_2, b_2, c_2, d_2, e_2, \mathcal{T}_2)$.

Outputs: Compute $a := a_0 + a_1, b := b_1 + b_2, c := c_0 + c_2, d := d_0 + d_1, e := e_1 + e_2$, and proceed as follows:

(a) **If** not $\mathcal{T}_0 = \mathcal{T}_1 = \mathcal{T}_2$, **then** \mathcal{F}_{CMP} aborts. **Else**, set $\mathcal{T} := \mathcal{T}_0$.

(b) **If** $a = b = c = d = e$ and $a < \mathcal{T}$ **then** output 0.

(c) **If** $a = b = c = d = e$ and $a \geq \mathcal{T}$ **then** output 1.

(d) **Else**, \mathcal{F}_{CMP} aborts (i.e. a, b, c, d , or e strings are not equal).

Corruption: Adversary \mathcal{A} maliciously corrupts one server. If \mathcal{A} instructs the functionality to abort, the functionality instructs the honest servers to abort.

Fig. 8: The ideal \mathcal{F}_{CMP} functionality for comparison.

- (e) *Pruning.* The servers proceed to pruning and invoke \mathcal{F}_{CMP} (Fig. 8) on the secret shares of cnt^γ (for $\gamma \in \{p || 0, p || 1\}$) for all sessions and threshold \mathcal{T} . Based on the output of \mathcal{F}_{CMP} the following occurs:
- \mathcal{F}_{CMP} returns 1 if $\text{cnt}^\gamma \geq \mathcal{T}$ (i.e., γ is a heavy-hitter string). In this case, the prefix γ is added to the list of $k + 1$ -bit heavy-hitter set (i.e., $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \gamma$).
 - \mathcal{F}_{CMP} returns 0 if $\text{cnt}^\gamma < \mathcal{T}$ (i.e., γ is a non heavy-hitter string). In this case, the prefix γ is ignored.
 - If \mathcal{F}_{CMP} returns \perp , then one of the servers behaved maliciously and the honest servers abort. This occurs if the malicious server has provided an incorrect threshold as input (condition 1 in \mathcal{F}_{CMP}) or it provided incorrect client output shares as input (condition 4 in \mathcal{F}_{CMP}).

This computation is performed in parallel for all $(k + 1)$ -bit prefixes in consideration, and after the pruning phase, HH^{k+1} contains the list of $(k + 1)$ -bit heavy hitter strings. Next, the above computation is repeated for $(k + 1)$ -bit strings to compute $(k + 2)$ -bit heavy hitters, until we reach $k = n - 1$. As already mentioned, \mathcal{F}_{CMP} is securely implemented using the state-of-the-art protocol of Rabbit [MRVW21].

Output Phase. At the end, the servers output $\text{HH}^{\leq n} = \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\}$ as the set of \mathcal{T} -heavy hitter strings. This completes the description of π_{HH} (Figs. 5, 6).

Theorem 1. *Assuming VIDPF is a verifiable incremental DPF and H_1, H_2 are random oracles, \mathcal{F}_{CMP} is a secure comparison functionality (Fig. 8), and H (in π_{check}) is collision-resistant, then π_{HH} (Figs. 5 and 6) implements \mathcal{F}_{HH} in the (random oracle, \mathcal{F}_{CMP})-model against malicious corruption of one server and $\ell' \leq \ell$ clients.*

Proof Sketch. Security of our protocol is captured in Theorem 1 and proven in Appendix C. Below we provide a security sketch. The adversary is allowed to corrupt $\ell' \leq \ell$ clients and one of the servers. The other two servers are honest. **A malicious client** attempts to inject an error and is detected in the following ways:

Client VIDPF keys are malformed. A malicious client can provide malformed VIDPF keys that are non-zero in more than one path in the tree. This gets detected in the session involving the honest servers due to the verifiable property of the VIDPF at each level when the servers verify the VIDPF proofs. If the checks pass, then it is ensured that the VIDPF keys provided by the client are valid.

Client VIDPF input is malformed. Next, a malicious client can try to double-vote on an input point, say $p \parallel 0 \in \{0, 1\}^{k+1}$ by constructing the VIDPF on $(p \parallel 0, \widetilde{\beta}^k)$, i.e., $f(p \parallel 0) = \widetilde{\beta}^k$, where $\widetilde{\beta}^k > 1$, instead of $(p \parallel 0, 1)$. This is detected by the honest servers since they perform a local subtree verification by reconstructing the value $y^p - (y^{p \parallel 0} - y^{p \parallel 1})$ and verifying that it equals 0 for all $k > 0$. For $k = 0$, the servers verify that $y^\epsilon = 1$.

VIDPF input is inconsistent across sessions. Finally, a malicious client can try to provide different VIDPF keys in different sessions. For example, it constructs VIDPF keys for $(\alpha_1, 1)$ for session $\mathcal{S}_0 - \mathcal{S}_1$, $(\alpha_2, 1)$ for session $\mathcal{S}_1 - \mathcal{S}_2$, and $(\alpha_3, 1)$ for session $\mathcal{S}_2 - \mathcal{S}_0$, where $\alpha_1, \alpha_2, \alpha_3 \in \{0, 1\}^n$ and might be different. To ensure the input is consistent across sessions, the servers match the difference of the reconstructed output of $\mathcal{S}_0 - \mathcal{S}_1$ and $\mathcal{S}_2 - \mathcal{S}_0$ session, and the difference of the reconstructed output of $\mathcal{S}_2 - \mathcal{S}_0$ and $\mathcal{S}_1 - \mathcal{S}_2$ session, to verify that they are all 0. By transitivity, it is ensured that the VIDPF evaluation is the same across the sessions if and only if the checks pass, ensuring that $\alpha_1 = \alpha_2 = \alpha_3$. This is performed by computing $\widehat{h^{p \parallel 0}}$ and $\widehat{h^{p \parallel 1}}$ hashes.

A malicious server can collude with malicious clients. Observe that the honest clients' inputs are always hidden from the adversary due to input privacy of VIDPF. Next, a malicious server could incorporate an erroneous VIDPF evaluation (from a malformed client input key) or inject additive errors into the output. We show how this is tackled in the protocol based on the server corruption:

\mathcal{S}_0 is corrupt. In this case, the session between $\mathcal{S}_1 - \mathcal{S}_2$ is honest. \mathcal{S}_0 runs this session with \mathcal{S}_1 since it obtained $\text{key}_{(2,1)}$ from the client. However, \mathcal{S}_2 behaves as an attestator by sending hashes of the messages that \mathcal{S}_0 is supposed to send. This forces \mathcal{S}_0 to act honestly in the $\mathcal{S}_1 - \mathcal{S}_2$, otherwise, it leads to an abort. Another way a malicious \mathcal{S}_0 can behave badly is by colluding with a malicious client. The client can provide malformed inputs in $\mathcal{S}_0 - \mathcal{S}_1 / \mathcal{S}_2 - \mathcal{S}_0$ session or inconsistent inputs across the three sessions. In such a case, a malicious \mathcal{S}_0 could compute an incorrect hash $\widehat{h^{p \parallel 0}} := \text{H}_1(y_{(0,1)}^{p \parallel 0'} - y_{(0,2)}^{p \parallel 0'}, y_{(0,2)}^{p \parallel 0'} - y_{(2,1)}^{p \parallel 0'})$ and $\widehat{h^{p \parallel 1}} := \text{H}_1(y_{(0,1)}^{p \parallel 1'} - y_{(0,2)}^{p \parallel 1'}, y_{(0,2)}^{p \parallel 1'} - y_{(2,1)}^{p \parallel 1'})$ where $y_{(0,1)}^{p \parallel 0'}, y_{(0,2)}^{p \parallel 0'}, y_{(0,1)}^{p \parallel 1'}, y_{(0,2)}^{p \parallel 1}'$ are incorrect. This allows \mathcal{S}_0 to introduce an additive error into the frequency for $p \parallel 0$ and $p \parallel 1$ (for the $\mathcal{S}_0 - \mathcal{S}_1$ and $\mathcal{S}_2 - \mathcal{S}_0$ sessions) by incorporating the client's malformed input. However, this gets detected when the output count is secretly reconstructed by \mathcal{F}_{CMP} for all three sessions. The reconstructed count will not match and the ideal functionality would return a \perp message detecting that one of the servers behaved maliciously, leading to an abort in the π_{HH} . The case where \mathcal{S}_1 is corrupt is symmetrical.

\mathcal{S}_2 is corrupt. In this case, the session between $\mathcal{S}_0 - \mathcal{S}_1$ is honest. If \mathcal{S}_2 behaves as a malicious attestator by sending incorrect hashes for the $\mathcal{S}_1 - \mathcal{S}_2$ or $\mathcal{S}_2 - \mathcal{S}_0$ sessions then the honest servers abort. A malicious \mathcal{S}_2 can also collude with a malicious client, and the latter can provide malformed inputs in the three sessions. If this happens in the $\mathcal{S}_0 - \mathcal{S}_1$ session then it gets detected due to verifiability of the VIDPF and the local subtree verification, since both \mathcal{S}_0 and \mathcal{S}_1 are honest. If the client provides malformed (e.g., double voting) VIDPF keys $\text{key}'_{(2,0)}$ and $\text{key}'_{(2,1)}$ to \mathcal{S}_1 and \mathcal{S}_0 for the sessions involving \mathcal{S}_2 , it again gets detected since \mathcal{S}_0 computes the hashes $\widehat{h^{p \parallel 0}}$ and $\widehat{h^{p \parallel 1}}$ honestly and \mathcal{S}_1 verifies them honestly.

Round Complexity. Next, we analyze the round complexity of our heavy-hitters protocol. The Server computation is performed for n levels ($k \in [0, \dots, n - 1]$), where each level involves “VIDPF Evaluation”, “Batch Verification”, “Aggregation”, and “Pruning” phases. The VIDPF evaluation and aggregation steps are performed locally by each server. Each batch-verification step requires $\lceil \log_2 \ell \rceil + 1$ rounds in the worst case (when there are malformed client inputs at each level) and a single round in the best case (when all the clients are honest). However, all verification steps for level k are performed in parallel and are batched.

We further elaborate on this in Section 5. In the pruning phase, the servers run a protocol that implements \mathcal{F}_{CMP} for each prefix, which is performed in parallel for all prefixes at the same level. Instantiating \mathcal{F}_{CMP} with Rabbit [MRVW21] involves $\log_2 |\mathbb{G}|$ rounds, where the frequency count is performed over \mathbb{G} . Summing up, the best case round complexity of PLASMA is $n \cdot (1 + \log_2 |\mathbb{G}|)$ and the worst case round complexity is $n \cdot (\lceil \log_2 \ell \rceil + 1 + \log_2 |\mathbb{G}|)$. For benchmarking, we implement group \mathbb{G} using a 64-bit ring to exploit native CPU ring optimizations.




Fig. 9: π_{check} for equality verification of ℓ strings between two parties and identification of unequal strings.

5 Batched Consistency Check

We now present our batched consistency check π_{check} that enables two parties, P_0 and P_1 , to verify the equality of lists \mathbf{u} and \mathbf{v} containing ℓ strings using Merkle trees. If the two lists are equal then π_{check} returns

$\text{ver} = 1$, else it returns $\text{ver} = 0$ and a list L containing the indices of elements where the lists differ. Correctness follows from the collision resistance property of the hash function H .

As summarized in Fig. 9, π_{check} requires $K + 1$ rounds of communication, where $K = \lceil \log_2 \ell \rceil$. The total communicated hashes are roughly $4\ell'(\log_2 \frac{\ell}{\ell'} + 2)$, where \mathbf{u} and \mathbf{v} differ on ℓ' elements. It can be further optimized to $2\ell'(\log_2 \frac{\ell}{\ell'} + 2)$, where only one of the parties sends its hashes instead of both. We provide a detailed analysis of the protocol in Appendix D. In case $\ell' = 0$, then our communication is a pair of hashes.

6 Experimental Evaluations

We implement our heavy-hitters protocol π_{HH} in Rust and use the `tarpc` framework by Google for asynchronous Remote Procedure Calls (RPC).⁴ PLASMA is fully parallelized: all sessions in each server run in parallel and we employ parallel iterators to process multiple client requests concurrently. (We apply the same parallelization for benchmarking Poplar.) We instantiate the PRG for VIDPF using the AES-NI hardware instructions for AES encryption with a seed of $\kappa = 128$ bits. We used rings in PLASMA (instead of fields) since our checks rely on the security of VIDPF (i.e., XOR-collision resistant property that is provided by the random oracle). Conversely, the security of Poplar relies on a statistical check for the client’s input validation. This check relies on the underlying group size and needs 62 bits for the statistical failure probability to be 2^{-60} for intermediate levels; for the leaves, we use the default size of a finite field of $2\kappa = 256$ bits as mentioned in Poplar.

Experiment Details. Our experiments vary the number of clients between $\ell = 10^3$ and $\ell = 10^6$ with two different bit-string sizes, $n = 64$ and $n = 256$ bits. We configured the threshold \mathcal{T} to be 1% of the clients’ strings, and we report the client and server costs, while empirically comparing with Poplar. Then, we compute the total monetary costs (due to runtime and communication) incurred by PLASMA servers, and we compare it with [AHI+22] (since the code of [AHI+22] is not open-source) based on the monetary cost.

Experimental Setup. We performed both LAN and WAN⁵ experiments on AWS EC2 machines (c5.9xlarge) each with 36 vCPUs at 3.60 GHz. PLASMA is compiled using Rust 1.74, and client-side experiments are carried out using a standard laptop with an Intel i7-8650U CPU (1.90 GHz).

Performance Evaluation. In our experiments, our goal is to answer the following questions:

- How efficient is PLASMA for each client and server?
- How does PLASMA compare with similar works (such as Poplar) that leverage DPFs?
- How does PLASMA compare with the related works that provide similar security guarantees, such as [AHI+22]?

Client costs. The PLASMA client generates three pairs of DPF keys. Meanwhile, the Poplar client generates two pairs of DPF keys but also computes a malicious sketching operation. As a result, both PLASMA and Poplar clients are extremely fast, running in the order of 20 – 24 microseconds on 256-bit inputs. A detailed comparison of client runtime can be found in Fig. 10 (a).

In terms of client communication, PLASMA transmits eight DPF keys, whereas Poplar transmits four DPF keys plus the correlated randomness for the sketching operation. As shown in Fig. 10 (b), we observed that the clients in both protocols incur the same communication overhead, roughly around 55 KB for 256 bits.

Server costs. In this experiment, we run PLASMA with randomly distributed malicious clients and compare it with Poplar. We set the malicious clients ℓ' to be a 0, 0.01, 0.1, and 0.3 fraction of the total clients ℓ . We observe that running with $\ell' = 0.01\ell$ has slightly faster performance than 0.1ℓ , while 0.3ℓ exhibits slightly worse performance than 0.1ℓ . Still, these differences are marginal compared to the total runtime,⁶ so we opt for reporting the 0 and 0.1ℓ to make the figures more clear.

⁴ Our code is available at <https://github.com/TrustworthyComputing/plasma>.

⁵ We used one server in Oregon, one in Ohio, and one in N. Virginia. For Poplar, we used one in Oregon and the other one in N. Virginia.

⁶ Performance is impacted by expanding the Merkle tree which happens if there is at least one malicious client.




Fig. 10: Comparisons of client costs for PLASMA and Poplar (KB is Kilobytes and μs is microseconds).

LAN Server Runtime. PLASMA outperforms Poplar in terms of server runtime by $2.7\times$ (64 bits) and $5\times$ (256 bits) for $\ell = 10^6$ clients and $\mathcal{T} = 1\%$ of the clients. This improvement is largely attributed to our efficient VIDPF-based client input validation. Although the presence of malicious clients has an impact on PLASMA’s performance, it still remains significantly faster than Poplar as presented in Fig. 11. Meanwhile, Poplar servers validate clients’ inputs using an expensive malicious secure sketching protocol.

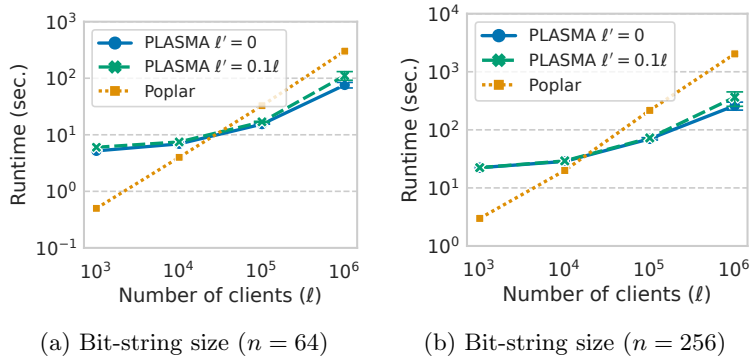


Fig. 11: Server runtime over LAN.

WAN Server Runtime. We benchmark PLASMA and Poplar over WAN for $n = 64$ and 256 bits and report our findings in Fig. 12. While the total latency is increased for both frameworks, we observe that the server WAN runtime for PLASMA increased by roughly 5-10% compared to server LAN runtime, whereas for Poplar the runtime increases by roughly 50%. We observe almost 5 – 10 \times improvement in terms of server WAN runtime for PLASMA compared to Poplar since PLASMA incurs significantly less communication for $\mathcal{T} = 1\%$.

Server-to-Server Communication. We compare the total communication costs incurred by all servers for an increasing number of clients, $\mathcal{T} = 1\%$, and $n = 256$ in Fig. 13. Poplar servers incur 35 GB of communication, whereas, PLASMA servers communicate less than 1 GB of data when considering $\ell' = 0$ and 0.1ℓ corrupt clients, hence yielding a 35 \times improvement over Poplar. The implementation of [AHI+22] is not open-source so we estimate the communication cost of [AHI+22] in Appendix G. The protocol of [AHI+22] communicates 45 GB of data to compute heavy-hitters over 10^6 client submitted 256-bit inputs. This yields a 45 \times improvement of PLASMA over [AHI+22].

Server Monetary Cost. To obtain fair comparisons between Poplar, [AHI+22], and PLASMA, we perform cumulative monetary cost analysis for a varying number of clients, assuming \$0.05/GB and \$1.53/hour. To estimate monetary costs, we run PLASMA and Poplar in a similar setup as [AHI+22] and compare it with the runtime provided in [AHI+22, Table 7.3] (which only considers 100k-400k clients over LAN). Note that




Fig. 12: Server runtime over WAN.

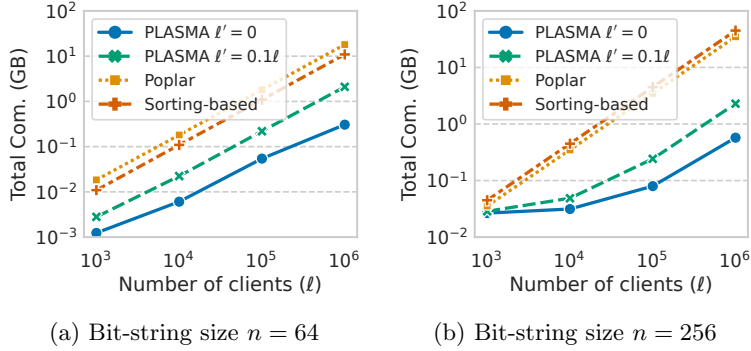


Fig. 13: Comparisons with Poplar [BBC⁺21] and the sorting-based approach of [AHI⁺22] in terms of total server-to-server communication (in GB).

Poplar runs two servers while PLASMA runs three. The monetary cost incurred by Poplar is two times the cost incurred by a single Poplar server, while for PLASMA it’s three times a single PLASMA server.

We present our findings in Fig. 14 for $\mathcal{T} = 1\%$ of the clients. Computing the \mathcal{T} most popular strings among 1 million clients with $n = 256$ bit strings, costs \$4.7 with Poplar, while PLASMA incurs \$0.6-\$0.9 costs for 0 to 0.1ℓ malicious clients. Meanwhile, [AHI⁺22] costs at least \$2.2 to perform the same task, so PLASMA yields a 2.5 – 3.5 \times improvement over [AHI⁺22] despite having a 15 \times runtime slowdown. This is largely due to the communication incurred by [AHI⁺22] for performing secure sorting under MPC. When considering input strings of smaller size, like $n = 64$, PLASMA is 4 \times cheaper than Poplar and 2 \times cheaper than [AHI⁺22]. Lastly, Vogue [JKK⁺22] is not open source and it benchmarks 100k-400k clients over LAN. It claims a 6 \times improvement over Poplar, whereas [AHI⁺22] claims an improvement of over 100 \times ; therefore we focused on comparing with [AHI⁺22].

Applications. We discuss two realistic applications:

Popular URLs. Each URL is represented as a 256-bit string and 10000 most popular URLs are computed among 1 million client-submitted URLs, assuming $\mathcal{T} = 1\%$. Server runtimes of PLASMA and Poplar are reported in Figs. 11 (b) and 12 (b), while the client communication costs in Figs. 10 (a) and (b) for $n = 256$. This benchmark is completed in under 5 minutes with less than 1 GB of data of communication for PLASMA, while Poplar servers incur more than 5 \times additional runtime costs and communicate 35 GB.

Popular GPS coordinates. We employ *plus codes* [LLC19] to efficiently encode the client GPS coordinates using 64 bits. This approach uses a grid system aligned on top of the world map, assigning specific codes to each area. Areas with similar codes are located in proximity to each other and a code that is a prefix of another encompasses the area of the latter. For instance, code 87 represents the North East US region, while code 87G8 represents a part of New York City. PLASMA uses plus codes to compute the most popular

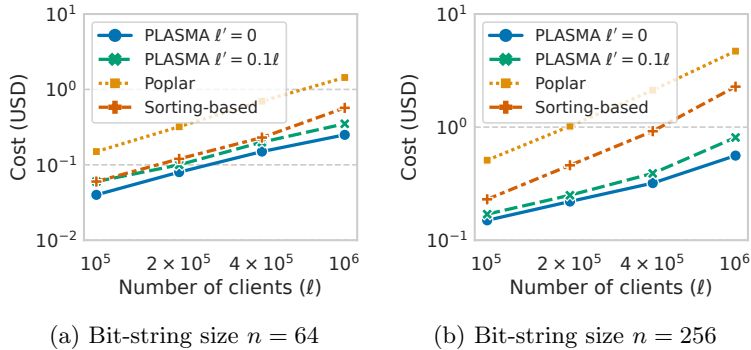


Fig. 14: Comparisons with Poplar and the sorting-based approach of [AHI+22] in terms of total monetary cost (in USD).

locations (submitted by more than $\mathcal{T} = 1\%$ of the clients) among a set of client-provided inputs using 64-bit strings in roughly 2 minutes for 10^6 clients, as shown in Fig. 12 (a). Client cost is shown in Fig. 10.

7 Further Extensions

We discuss two interesting extensions of PLASMA and compare them with the state-of-the-art protocol of [AHI+22]:

Fairness: The notion of fairness ensures that if an adversary receives an output then the honest parties also receive the correct output. If the adversary aborts then the honest parties also abort. In our case, we observe that the count is secret shared between the servers and based on the output of \mathcal{F}_{CMP} in the pruning phase, the servers compute the heavy-hitting prefix set. As a result, PLASMA is fair if the pruning phase is fair. This happens if \mathcal{F}_{CMP} functionality is implemented using a three-party subprotocol [CCPS19] that guarantees fairness against one malicious party. Hence, PLASMA can satisfy a stronger notion of security as compared to Poplar or [AHI+22], which only satisfies security with selective abort.

Heavy-Hitters over Multiple Thresholds: PLASMA enables computing heavy-hitters over multiple thresholds ($\mathcal{T}_1, \mathcal{T}_2, \dots$) based on some pre-agreed strings by the servers. This enables new applications like traffic avoidance, since different roads may have different traffic densities (e.g., highways are busier than smaller suburban roads). The servers consider that during evaluation and use higher values of \mathcal{T} for highways with more vehicles and lower values for smaller roads. Conversely, it is unclear how to extend [AHI+22] to support this feature. Protocol details are in Appendix E.

8 Concluding Remarks

In this work, we present PLASMA: a framework to privately identify the most popular strings – or heavy hitters – among a set of client inputs without revealing the client data points. Previous works for private heavy hitters, such as Poplar, consider security against malicious clients and were prone to additive attacks by a malicious server, compromising the correctness of the protocol. To address this challenge, PLASMA introduces a novel hash-based primitive, called verifiable incremental distributed point functions, which allows the servers to validate client inputs using inexpensive operations. Additionally, we introduce a new batched consistency check that uses Merkle trees to validate multiple client sessions in a batch. This drastically reduces the concrete server-to-server communication, incurred during the heavy-hitters computation.

Acknowledgments

D. Mouris and N.G. Tsoutsos would like to acknowledge the support of the National Science Foundation (Award 2239334), and the Electrical & Computer Engineering department at the University of Delaware. P.

Sarkar conducted the research at Boston University and he was supported by NSF Awards 1931714, 1414119, and the DARPA SIEVE program.

References

- ACC⁺21. Abdelrahman Aly, K Cong, D Cozzo, M Keller, E Orsini, D Rotaru, O Scherer, P Scholl, N Smart, T Tanguy, et al. Scale-mamba v1. 12: Documentation, 2021.
- ACD⁺21. Erik Anderson, Melissa Chase, F. Betul Durak, Esha Ghosh, Kim Laine, and Chenkai Weng. Aggregate measurement via oblivious shuffling. Cryptology ePrint Archive, Report 2021/1490, 2021. <https://eprint.iacr.org/2021/1490>.
- AG21. Apple and Google. Exposure Notification Privacy-preserving Analytics (ENPA) white paper, 2021.
- AHI⁺22. Gilad Asharov, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Ariel Nof, Benny Pinkas, Katsumi Takahashi, and Junichi Tomida. Efficient secure three-party sorting with applications to data analysis and heavy hitters. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 125–138, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.
- BBC⁺19. Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Zero-knowledge proofs on secret-shared data via fully linear PCPs. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 67–97, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.
- BBC⁺21. Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Lightweight techniques for private heavy hitters. In *2021 IEEE Symposium on Security and Privacy*, pages 762–776, San Francisco, CA, USA, May 24–27, 2021. IEEE Computer Society Press.
- BBC⁺23. Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. Arithmetic sketching. In *Advances in Cryptology – CRYPTO 2023, Part I*, Lecture Notes in Computer Science, pages 171–202, Santa Barbara, CA, USA, August 2023. Springer, Heidelberg, Germany.
- BGG⁺22. James Bell, Adrià Gascón, Badih Ghazi, Ravi Kumar, Pasin Manurangsi, Mariana Raykova, and Phillip Schoppmann. Distributed, private, sparse histograms in the two-server model. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 307–321, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.
- BGI15. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 337–367, Sofia, Bulgaria, April 26–30, 2015. Springer, Heidelberg, Germany.
- BGI16. Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 1292–1303, Vienna, Austria, October 24–28, 2016. ACM Press.
- BK21. Jonas Böhrer and Florian Kerschbaum. Secure multi-party computation of differentially private heavy hitters. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2361–2377, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- BLW08. Dan Bogdanov, Sven Laur, and Jan Willemsen. Sharemind: A framework for fast privacy-preserving computations. In Sushil Jajodia and Javier López, editors, *ESORICS 2008: 13th European Symposium on Research in Computer Security*, volume 5283 of *Lecture Notes in Computer Science*, pages 192–206, Málaga, Spain, October 6–8, 2008. Springer, Heidelberg, Germany.
- BNSGT17. Raef Bassily, Kobbi Nissim, Uri Stemmer, and Abhradeep Guha Thakurta. Practical locally private heavy hitters. *Advances in Neural Information Processing Systems*, 30:1–32, 2017.
- Can00. Ran Canetti. Security and composition of multiparty cryptographic protocols. *Journal of Cryptology*, 13(1):143–202, January 2000.
- CCPS19. Harsh Chaudhari, Ashish Choudhury, Arpita Patra, and Ajith Suresh. ASTRA: high throughput 3pc over rings with application to secure prediction. In *ACM SIGSAC CCSW@CCS 2019*, pages 81–92, London, UK, 2019. ACM.
- CGB17. Henry Corrigan-Gibbs and Dan Boneh. Prio: Private, robust, and scalable computation of aggregate statistics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation*, NSDI’17, page 259–282, USA, 2017. USENIX Association.

- CGH⁺18. Koji Chida, Daniel Genkin, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Yehuda Lindell, and Ariel Nof. Fast large-scale honest-majority MPC for malicious adversaries. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part III*, volume 10993 of *Lecture Notes in Computer Science*, pages 34–64, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- CJK⁺23. Benjamin Case, Richa Jain, Alex Koshelev, Andy Leiserson, Daniel Masny, Ben Savage, Erik Taubeneck, Martin Thomson, and Taiki Yamaguchi. Interoperable Private Attribution: A Distributed Attribution and Aggregation Protocol. *Cryptology ePrint Archive*, Report 2023/437, 2023. <https://eprint.iacr.org/2023/437>.
- CSW22. Ran Canetti, Pratik Sarkar, and Xiao Wang. Triply adaptive UC NIZK. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part II*, volume 13792 of *Lecture Notes in Computer Science*, pages 466–495, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany.
- dCP22. Leo de Castro and Antigoni Polychroniadou. Lightweight, maliciously secure verifiable function secret sharing. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part I*, volume 13275 of *Lecture Notes in Computer Science*, pages 150–179, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany.
- DKM⁺06. Cynthia Dwork, Krishnamurthy Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: Privacy via distributed noise generation. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 486–503, St. Petersburg, Russia, May 28 – June 1, 2006. Springer, Heidelberg, Germany.
- DMNS06. Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284, New York, NY, USA, March 4–7, 2006. Springer, Heidelberg, Germany.
- DPRS23. Hannah Davis, Christopher Patton, Mike Rosulek, and Phillipp Schoppmann. Verifiable Distributed Aggregation Functions. *Proceedings on Privacy Enhancing Technologies*, 2023(4):578–592, July 2023.
- DRPS22. Emma Dauterman, Mayank Rathee, Raluca Ada Popa, and Ion Stoica. Waldo: A private time-series database from function secret sharing. In *2022 IEEE Symposium on Security and Privacy*, pages 2450–2468, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.
- EDG14. Tariq Elahi, George Danezis, and Ian Goldberg. PrivEx: Private collection of traffic statistics for anonymous communication networks. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014: 21st Conference on Computer and Communications Security*, pages 1068–1079, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.
- EPK14. Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *ACM CCS 2014: 21st Conference on Computer and Communications Security*, pages 1054–1067, Scottsdale, AZ, USA, November 3–7, 2014. ACM Press.
- FLNW17. Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. High-throughput secure three-party computation for malicious adversaries and an honest majority. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 225–255, Paris, France, April 30 – May 4, 2017. Springer, Heidelberg, Germany.
- FPE16. Giulia Fanti, Vasyl Pihur, and Úlfar Erlingsson. Building a RAPPOR with the Unknown: Privacy-Preserving Learning of Associations and Data Dictionaries. *Proc. Priv. Enhancing Technol.*, 2016(3):41–61, 2016.
- GI14. Niv Gilboa and Yuval Ishai. Distributed point functions and their applications. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 640–658, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.
- HKR12. Justin Hsu, Sanjeev Khanna, and Aaron Roth. Distributed Private Heavy Hitters. In *Proceedings of the 39th International Colloquium Conference on Automata, Languages, and Programming - Volume Part I*, ICALP’12, page 461–472, Berlin, Heidelberg, 2012. Springer-Verlag.
- IKN⁺20. Mihaela Ion, Ben Kreuter, Ahmet Erhan Nergiz, Sarvar Patel, Shobhit Saxena, Karn Seth, Mariana Raykova, David Shanahan, and Moti Yung. On Deploying Secure Computing: Private Intersection-Sum-with-Cardinality. In *EuroS&P*, pages 370–389, Genoa, Italy, 2020. IEEE.
- JKK⁺22. Pranav Jangir, Nishat Koti, Varsha Bhat Kukkala, Arpita Patra, Bhavish Raj Gopal, and Somya Sangal. Vogue: Faster computation of private heavy hitters. *Cryptology ePrint Archive*, Report 2022/1561, 2022. <https://eprint.iacr.org/2022/1561>.

- Kel20. Marcel Keller. MP-SPDZ: A versatile framework for multi-party computation. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1575–1590, Virtual Event, USA, November 9–13, 2020. ACM Press.
- LLC19. Google LLC. Open Location Code. <https://github.com/google/open-location-code>, 2019.
- LPR⁺21. Tancrède Lepoint, Sarvar Patel, Mariana Raykova, Karn Seth, and Ni Trieu. Private join and compute from PIR with default. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part II*, volume 13091 of *Lecture Notes in Computer Science*, pages 605–634, Singapore, December 6–10, 2021. Springer, Heidelberg, Germany.
- MMT⁺24. Dimitris Mouris, Daniel Masny, Ni Trieu, Shubho Sengupta, Prasad Buddhavarapu, and Benjamin M Case. Delegated Private Matching for Compute. *Proceedings on Privacy Enhancing Technologies*, 2024(2):1–24, July 2024.
- MPD⁺24. Dimitris Mouris, Christopher Patton, Hannah Davis, Pratik Sarkar, and Nektarios Georgios Tsoutsos. Mastic: Private weighted heavy-hitters and attribute-based metrics. Cryptology ePrint Archive, Paper 2024/221, 2024. <https://eprint.iacr.org/2024/221>.
- MRVW21. Eleftheria Makri, Dragos Rotaru, Frederik Vercauteren, and Sameer Wagh. Rabbit: Efficient comparison for secure multi-party computation. In Nikita Borisov and Claudia Díaz, editors, *FC 2021: 25th International Conference on Financial Cryptography and Data Security, Part I*, volume 12674 of *Lecture Notes in Computer Science*, pages 249–270, Virtual Event, March 1–5, 2021. Springer, Heidelberg, Germany.
- NPR19. Moni Naor, Benny Pinkas, and Eyal Ronen. How to (not) share a password: Privacy preserving protocols for finding heavy hitters with adversarial behavior. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019: 26th Conference on Computer and Communications Security*, pages 1369–1386, London, UK, November 11–15, 2019. ACM Press.
- PAD⁺23. Antigoni Polychroniadou, Gilad Asharov, Benjamin E. Diamond, Tucker Balch, Hans Buehler, Richard Hua, Suwen Gu, Greg Gimler, and Manuela Veloso. Prime Match: A Privacy-Preserving Inventory Matching System, 2023.
- QYY⁺16. Zhan Qin, Yin Yang, Ting Yu, Issa Khalil, Xiaokui Xiao, and Kui Ren. Heavy hitter estimation over set-valued data with local differential privacy. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 192–203, Vienna, Austria, October 24–28, 2016. ACM Press.
- VHG23. Adithya Vadapalli, Ryan Henry, and Ian Goldberg. Duoram: A Bandwidth-Efficient Distributed ORAM for 2- and 3-Party Computation. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 3907–3924, Anaheim, CA, August 2023. USENIX Association.
- VSH22. Adithya Vadapalli, Kyle Storrier, and Ryan Henry. Sabre: Sender-anonymous messaging with fast audits. In *2022 IEEE Symposium on Security and Privacy*, pages 1953–1970, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.
- WSK⁺23. Yongqin Wang, Pratik Sarkar, Nishat Koti, Arpita Patra, and Murali Annavaram. Compacttag: Minimizing computation overheads in actively-secure MPC for deep neural networks. *IACR Cryptol. ePrint Arch.*, page 1729, 2023.
- YSWW21. Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. QuickSilver: Efficient and affordable zero-knowledge proofs for circuits and polynomials over any field. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2986–3001, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- ZKM⁺20. Wennan Zhu, Peter Kairouz, Brendan McMahan, Haicheng Sun, and Wei Li. Federated Heavy Hitters Discovery with Differential Privacy. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3837–3847, Online, 26–28 Aug 2020. PMLR.

Appendix

A Variants of Distributed Point Functions

Incremental and Verifiable DPF (IDPF and VDPF). The IDPF [BBC⁺21] and VDPF [dCP22] build on standard DPFs to secret share the weights of a tree w.r.t. a single non-zero path. IDPFs perform this task with linear cost in the number of bits n for strings that share common prefixes [BBC⁺21], whereas using standard DPFs this cost would grow to $\mathcal{O}(n^2)$. IDPFs rely on expensive malicious secure sketching checks to

ensure that an IDPF key is not malformed. Meanwhile, the work of [dCP22] considers efficient hashing-based verifiable properties to ensure that a DPF (not IDPF) key is well-formed. Moreover, [dCP22] enables a batched verification procedure with communication proportional to the security parameter. However, VDPFs work only for DPF and not IDPF. We present the VDPF algorithms below:

- $\text{VDPF.Gen}(1^\kappa, f_{\alpha,\beta}) \rightarrow (\text{key}_0, \text{key}_1)$. Given the security parameter 1^κ and a function f , output keys $\text{key}_0, \text{key}_1$.
- $\text{VDPF.BatchEval}(b, \text{key}_b, \mathbf{X}) \rightarrow (\mathbf{Y}_b, \pi_b)$: For $b \in \{0, 1\}$, batch verifiable evaluation takes a set $\mathbf{X} := \{x_1, x_2, \dots, x_m\}$, where each $x_i \in \{0, 1\}^n$. Outputs $\mathbf{Y}_b := \{y_{b,1}, y_{b,2}, \dots, y_{b,m}\}$.

Correctness ensures that $\mathbf{Y}_0 + \mathbf{Y}_1 = f_{\alpha,\beta}(\mathbf{X})$. Privacy ensures that an adversary in possession of one of the keys (but not both) does not obtain any information about the function f . The verifiability property of VDPF ensures that the proofs π_0 and π_1 are the same if and only if they have been generated from valid keys key_0 and key_1 of a point function.

B Verifiable Incremental DPF

We present the VIDPF construction, denoted as π_{VIDPF} , in Figs. 15 and 16. Our VIDPF construction is obtained by adding verifiability (steps 15-17 from Fig. 15) on top of the IDPF construction of Poplar. We have underlined the lines that focus on verifiability in these two figures. The `Convert` takes the corrected seed $\tilde{s}_b^{(i)}$ for level i , runs PRG'' and outputs κ bit seed $s_b^{(i)}$ for level i and value $W_b^{(i)}$. This occurs at the intermediate levels and is performed by executing the “else” part of `Convert`. $W_b^{(i)}$ comes from \mathbb{G} since it generates the output W_{cw} based on intermediate β_i . At the leaves, the “if” part of `Convert` is executed where only $W_b^{(i)}$ is generated. The security of our protocol is summarized in Theorem 2.

Theorem 2. *Assuming $(\text{PRG}, \text{PRG}', \text{PRG}'')$ are pseudorandom generators, PRG is κ -collision resistant and (H_1, H_2) are random oracles then $\pi_{\text{VIDPF}} = (\text{Gen}, \text{EvalPref})$ in Figs. 15 and 16 is a VIDPF.*

We define κ -collision resistant PRG as follows:

Definition 2 (κ -Collision Resistant PRG). *We say that a $\text{PRG} : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2\kappa+2}$ is κ -collision resistant if a PPT adversary cannot output s_0 and s_1 such that*

$$\begin{aligned} (A_0 \parallel T_0 \parallel B_0 \parallel T'_0) &:= \text{PRG}(s_0), \\ (A_1 \parallel T_1 \parallel B_1 \parallel T'_1) &:= \text{PRG}(s_1), \\ \text{and } B_0 &= B_1, \end{aligned}$$

where $A_0, A_1, B_0, B_1 \in \{0, 1\}^\kappa$ and $T_0, T'_0, T_1, T'_1 \in \{0, 1\}$.

We recall the notion of XOR-collision resistance from [dCP22] below for our security proof.

Definition 3 (XOR-Collision Resistance [dCP22]). *We say a function family \mathcal{F} is XOR-collision resistant if no PPT adversary given a randomly sampled $f \in \mathcal{F}$ can find four values $x_0, x_1, x_2, x_3 \in \{0, 1\}^\kappa$ such that $(x_0, x_1) \neq (x_2, x_3)$, $(x_0, x_1) \neq (x_3, x_2)$, and $f(x_0) \oplus f(x_1) = f(x_2) \oplus f(x_3) \neq 0$, except with negligible probability in security parameter κ .*

It can be implemented by assuming the function f is a random oracle. Next, we proceed to the proof of Thm. 2.

Proof. Input privacy of our VIDPF follows from the input privacy of the underlying IDPF protocol from Poplar, which in turn relies on the pseudorandomness of PRG. Adding $\text{cs}^{(i)}$ in steps 16-17 (Fig. 15) does not affect the input privacy of the client in the random oracle model since $\text{cs}^{(i)} = \tilde{\pi}_0^{(i)} \oplus \tilde{\pi}_1^{(i)}$ is an XOR of two random oracle outputs. Each server will know the preimage of either $\tilde{\pi}_0^{(i)}$ or the preimage of $\tilde{\pi}_1^{(i)}$ by evaluating the given VIDPF key. The server breaks input privacy if it computes both preimages. However, to compute

the other preimage it needs to invert the random oracle on $\tilde{\pi}_{1-b'}^{(i)}$ (assuming it obtained the preimage of $\tilde{\pi}_b^{(i)}$ by evaluating the VIDPF key).

A malicious client breaks the verifiability property if there are two non-zero paths, say u and v in the evaluation tree such that the client still passes the verification check performed by the servers on $\text{cs}^{(i)}$ for $i \in [n]$. We prove this via two steps:

- *At most one non-zero value at each level i* : We prove this via contradiction. Assume a client generates VIDPF keys that evaluate to two non-zero values at level i . It means the servers obtain $s_0^i(u)$, $s_1^i(u)$, $s_0^i(v)$ and $s_1^i(v)$ from Step 11 of EvalNext (Fig. 16) by evaluating on u and v such that the following holds:

$$s_0^i(u) \neq s_1^i(u) \text{ and } s_0^i(v) \neq s_1^i(v)$$

$$\text{cs}^{(i)} = \tilde{\pi}_0^{(i)}(u) \oplus \tilde{\pi}_1^{(i)}(u) = \tilde{\pi}_0^{(i)}(v) \oplus \tilde{\pi}_1^{(i)}(v),$$

where $\tilde{\pi}_b^{(i)}(u) := H_1(u, s_b^i(u))$ and $\tilde{\pi}_b^{(i)}(v) := H_1(v, s_b^i(v))$ for $b \in \{0, 1\}$. However, this is not possible in the random oracle model since it breaks the XOR-collision-resistance property of the random oracle H_1 . The adversary cannot find such a set of $s_0^i(u)$, $s_1^i(u)$, $s_0^i(v)$ and $s_1^i(v)$ values. Lemma 3 of [dCP22] captures the formal details of the reduction. In addition, the servers also check multiple proofs by iteratively hashing them together using H_2 in step 12 of the EvalNext algorithm. So, we also rely on the collision resistance property of H_2 to argue that it suffices to check the equality of the hash values computed using H_2 to ensure that the preimages are equal.

- *Non-zero value at level $i+1$ is a child of non-zero value at level i* : We prove this via contradiction. Assume a client generates VIDPF keys that evaluate to a non-zero value at level i on prefix $u \in \{0, 1\}^i$ and a non-zero value at level $i+1$ on prefix $v \in \{0, 1\}^{i+1}$ such that the non-zero node at level i is not the parent of the non-zero value at level $i+1$, i.e., $u \neq v^{\leq i}$. This means that $s_0^i(v) = s_1^i(v)$ and $s_0^i(u) \neq s_1^i(u)$ since there can be at most one pair of non-zero s values at each level. Next, consider the inputs to the EvalNext algorithms for evaluation on input prefix v in Fig. 16. We consider the following two cases:
 - *st^i is same for both servers*: In this case both $s_0^i(v) = s_1^i(v)$ and $t_0^i(v) = t_1^i(v)$. Here the input of the server to EvalNext is the same except for the value b . Hence, the evaluation algorithm of the servers on prefix v will be identical except in step 10 where server b obtains $y_b^{(i+1)}$ values such that $y_0^{(i+1)} + y_1^{(i+1)} = 0$. So, the output cannot be non-zero in this case.
 - *st^i is different for both servers*: In this case, $s_0^i(v) = s_1^i(v)$ but $t_0^i(v) \neq t_1^i(v)$. For this to happen there exists $s_0^{i-1}(v)$, $t_0^{i-1}(v)$, $s_1^{i-1}(v)$, $t_1^{i-1}(v)$ such that $(s_0^i(v), t_0^i(v))$ and $(s_1^i(v), t_1^i(v))$ are obtained by computing PRG on $s_0^{i-1}(v)$ and $s_1^{i-1}(v)$ respectively and applying Step 4 of EvalNext based on $t_0^{i-1}(v)$ and $t_1^{i-1}(v)$ values.
 - *If $v^{\leq i-1} = u^{\leq i-1}$* : then $s_0^{i-1}(v) = s_0^{i-1}(u)$ and $s_1^{i-1}(v) = s_1^{i-1}(u)$. But we know that $s_0^i(u) \neq s_1^i(u)$. We also know that $s_b^i(u)$ and $s_b^i(v)$ is generated from the same state st_b^{i-1} by server b where only $u_i \neq v_i$, which is x_i in EvalNext. In this case, steps 1-5 are the same for the evaluation of u_i and v_i . Assume $u_i = 0$ and $v_i = 1$ without loss of generality. This means that $s_0^L \neq s_1^L$ and $s_0^R = s_1^R$, where s_b^L and s_b^R are computed by server b by evaluating the PRG on s_b^{i-1} and then XORing the output with $t^{(i-1)} \cdot s_{\text{cw}}$. This breaks the collision resistance of PRG since $s_0^{i-1} \neq s_1^{i-1}$ but $(A_0 \parallel T_0 \parallel \mathbf{B} \parallel T_0') := \text{PRG}(s_0^{i-1})$ and $(A_0 \parallel T_0 \parallel \mathbf{B} \parallel T_0') := \text{PRG}(s_1^{i-1})$ where $A_0, A_1, B, B' \in \{0, 1\}^\kappa$ and $T_0, T_0', T_1, T_1' \in \{0, 1\}$.
 - *If $v^{\leq i-1} \neq u^{\leq i-1}$* : then $s_0^{i-1}(v) = s_1^{i-1}(v)$ and $t_0^{i-1}(v) \neq t_1^{i-1}(v)$ and we apply our argument recursively for $i-2$ and so on until we get the previous case where $u^\ell = v^\ell$ for $\ell \in [i-1]$.

□

We note that we do not need collision resistance from the PRG since we do not require that the non-zero values lie on the same path. We only need that each level contains a single non-zero node and for that the XOR-collision resistance property suffices. This property is implemented by assuming that (H_1, H_2) are random oracles.

Notation: We denote the private n -bit string α and its bit decomposition as $\alpha_1, \dots, \alpha_n \in \{0, 1\}^n$.

Primitives: $\text{PRG} : \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2\kappa+2}$ is a pseudorandom generator. $\text{H}_1 : \{0, 1\}^* \times \{0, 1\}^\kappa \rightarrow \{0, 1\}^{2\kappa}$ and $\text{H}_2 : \{0, 1\}^{2\kappa} \rightarrow \{0, 1\}^{2\kappa}$ are random oracles.

Gen($1^\kappa, 1^n, \alpha, (\beta_1, \beta_2, \dots, \beta_n), \mathbb{G}$): ▷ Generate DPF keys.

- 1: Sample $s_b^{(0)} \xleftarrow{r} \{0, 1\}^\kappa$ for $b \in \{0, 1\}$ ▷ Secret seeds.
- 2: Let $t_0^{(0)} := 0$ and $t_1^{(0)} := 1$
- 3: **for** $i := 1$ to n **do** ▷ For each bit of α .
- 4: $s_b^L \parallel t_b^L \parallel s_b^R \parallel t_b^R := \text{PRG}(s_b^{(i-1)})$ for $b \in \{0, 1\}$ ▷ Parse the output of PRG as a sequence of $(\kappa \parallel 1 \parallel \kappa \parallel 1)$ bits.
- 5: **if** $\alpha_i = 0$ **then** $\text{Diff} := L, \text{Same} := R$ ▷ Set right children to be equal.
- 6: **else** $\text{Diff} := R, \text{Same} := L$ ▷ Set left children to be equal.
- 7: $s_{\text{cw}} := s_0^{\text{Same}} \oplus s_1^{\text{Same}}$
- 8: $t_{\text{cw}}^L := t_0^L \oplus t_1^L \oplus \alpha_i \oplus 1$ ▷ Left control bits not equal if $\alpha_i = 0$.
- 9: $t_{\text{cw}}^R := t_0^R \oplus t_1^R \oplus \alpha_i$ ▷ Right control bits not equal if $\alpha_i = 1$.
- 10: $\tilde{s}_b^{(i)} := s_b^{\text{Diff}} \oplus t_b^{(i-1)} \cdot s_{\text{cw}}$ for $b \in \{0, 1\}$ ▷ Correction.
- 11: $t_b^{(i)} := t_b^{\text{Diff}} \oplus t_b^{(i-1)} \cdot t_{\text{cw}}^{\text{Diff}}$ for $b \in \{0, 1\}$ ▷ Correction.
- 12: $s_b^{(i)} \parallel W_b^{(i)} := \text{Convert}(\tilde{s}_b^{(i)})$ for $b \in \{0, 1\}$
- 13: $W_{\text{cw}}^{(i)} := (-1)^{t_1^{(i)}} \cdot [\beta_i - W_0^{(i)} + W_1^{(i)}]$ ▷ Output correction.
- 14: $\text{cw}^{(i)} := s_{\text{cw}} \parallel t_{\text{cw}}^L \parallel t_{\text{cw}}^R \parallel W_{\text{cw}}^{(i)}$ ▷ Correction word for level i .
- 15: $\tilde{\pi}_b^{(i)} = \text{H}_1(\alpha_{\leq i} \parallel s_b^{(i)})$
- 16: $\text{cs}^{(i)} = \tilde{\pi}_0^{(i)} \oplus \tilde{\pi}_1^{(i)}$.
- 17: $\text{key}_b := (s_b^{(0)} \parallel \text{cw}^{(1)} \parallel \dots \parallel \text{cw}^{(n)} \parallel \text{cs}^{(1)} \parallel \dots \parallel \text{cs}^{(n)})$ for $b \in \{0, 1\}$ ▷ Key for party b .
- 18: **return** key_b for $b \in \{0, 1\}$

Convert $_{\mathbb{G}}(s)$:

- 1: Let $u \leftarrow |\mathbb{G}|$.
- 2: **if** $u = 2^m$ for an integer m **then**:
- 3: Return the group element represented by $\text{PRG}'(s) \bmod u$,
- 4: where $\text{PRG}' : \{0, 1\}^\kappa \rightarrow \{0, 1\}^m$.
- 5: **else**:
- 6: Let $n = \lceil \log_2 u \rceil + \kappa$.
- 7: Return the group element represented by $\text{PRG}''(s) \bmod u$,
- 8: where $\text{PRG}'' : \{0, 1\}^\kappa \rightarrow \{0, 1\}^n$.

Fig. 15: Protocol π_{VDPF} for Verifiable Incremental DPF (continues in Fig. 16).

EvalNext ($b, i, \text{st}^{(i-1)}, \text{cw}^{(i)}, \text{cs}^{(i)}, x^{\leq i}, \pi$):	▷ Evaluate x_i .
1: Parse $\text{st}^{(i-1)}$ as $(s^{i-1} \parallel t^{i-1})$.	
2: $s_{\text{cw}} \parallel t_{\text{cw}}^L \parallel t_{\text{cw}}^R \parallel W_{\text{cw}}^{(i)} := \text{cw}^i$	▷ Parse correction word.
3: $\tilde{s}^L \parallel \tilde{t}^L \parallel \tilde{s}^R \parallel \tilde{t}^R := \text{PRG}(s^{(i-1)})$	▷ Parse the output of PRG as a sequence of $(\kappa \parallel 1 \parallel \kappa \parallel 1)$ bits.
4: $\tau^{(i)} := (\tilde{s}^L \parallel \tilde{t}^L \parallel \tilde{s}^R \parallel \tilde{t}^R) \oplus (t^{(i-1)} \cdot [s_{\text{cw}} \parallel t_{\text{cw}}^L \parallel s_{\text{cw}} \parallel t_{\text{cw}}^R])$	
5: $s^L \parallel t^L \parallel s^R \parallel t^R := \tau^{(i)}$	▷ Parse $\tau^{(i)}$.
6: if $x_i = 0$ then $\tilde{s}^{(i)} := s^L, \quad t^{(i)} := t^L$	▷ Keep left path.
7: else $\tilde{s}^{(i)} := s^R, \quad t^{(i)} := t^R$	▷ Keep right path.
8: $s^{(i)} \parallel W^{(i)} := \text{Convert}(\tilde{s}^{(i)})$	▷ New seed and output for level i .
9: $\text{st}^{(i)} := s^{(i)} \parallel t^{(i)}$	▷ Save the state.
10: $y^{(i)} := (-1)^b \cdot [W^{(i)} + t^{(i)} \cdot W_{\text{cw}}]$	▷ Compute output at level i .
11: $\tilde{\pi}^{(i)} = \text{H}_1(x^{\leq i} \parallel s^{(i)})$.	
12: $\pi = \pi \oplus \text{H}_2(\pi \oplus (\tilde{\pi}^{(i)} \oplus t^{(i)} \cdot \text{cs}^{(i)}))$.	
13: return $(\text{st}^{(i)}, y^{(i)}, \pi)$	
EvalPref ($b, \text{key}, x \in \{0, 1\}^n, \text{st}^{(d-1)}, d, \pi$):	▷ Evaluate one public bitstring x on all its bits x_i for $i \in [n]$.
1: Parse key as $s^{(0)} \parallel \text{cw}^{(1)} \parallel \dots \parallel \text{cw}^{(n)} \parallel \text{cs}^{(1)} \parallel \dots \parallel \text{cs}^{(n)}$.	▷ Parse key for party b .
2: if $d \neq 1$ then parse $\text{st}^{(d-1)}$ as $(s^{(d-1)} \parallel t^{(d-1)})$,	
3: else $t^{(0)} := b, \quad \text{st}^{(0)} := s^{(0)} \parallel t^{(0)}$.	
4: for $i := d$ to n do	▷ For each bit of x .
5: $(\text{st}^{(i)}, y^{(i)}, \pi) := \text{EvalNext}(b, i, \text{st}^{(i-1)}, \text{cw}^i, \text{cs}^i, x^{\leq i}, \pi)$.	
6: return $(\text{st}^{(n)}, y^{(n)}, \pi)$	

Fig. 16: Protocol π_{VIDPF} for Verifiable Incremental DPF (continuing from Fig. 15).

C Heavy-Hitters Protocol π_{HH} Proof

In this section, we formally prove Theorem 1. Security of our protocol relies on the correctness of π_{check} . π_{check} is a protocol where two honest parties commit to their inputs using Merkle-tree-based commitments and then they decommit based on whether the root commitments match or not. Correctness of π_{check} follows straightforwardly from the binding property of the Merkle-tree commitment, which in turn follows from the collision-resistance property of the hash function used in π_{check} .

Next, we prove the security of our protocol in the real-ideal world paradigm of Canetti [Can00]. Let \mathcal{A} denote the real-world adversary corrupting one of the servers and ℓ' clients maliciously in the real-world execution of the protocol. Let $\text{REAL}_{\mathcal{A}, \pi_{\text{HH}}}$ denote \mathcal{A} 's view after participating in the real-world execution. Let simulator Sim be the ideal-world adversary, which given access to the algorithm of \mathcal{A} and functionality \mathcal{F}_{HH} , produces the ideal world adversarial view as $\text{IDEAL}_{\text{Sim}, \mathcal{F}_{\text{HH}}}$.

We prove that our protocol π_{HH} securely implements \mathcal{F}_{HH} functionality by providing an ideal world PPT simulator Sim for all PPT adversaries \mathcal{A} , and show that the real and ideal world view are indistinguishable, i.e., $\text{REAL}_{\mathcal{A}, \pi_{\text{HH}}} \stackrel{c}{\approx} \text{IDEAL}_{\text{Sim}, \mathcal{F}_{\text{HH}}}$. We use a sequence of hybrids (i.e., $\text{HYB}_0 - \text{HYB}_4$) to prove indistinguishability.

Proof. We first consider the case where \mathcal{A} corrupts server \mathcal{S}_2 along with ℓ' clients. Then, we consider the case where \mathcal{A} corrupts either \mathcal{S}_0 or \mathcal{S}_1 along with ℓ' clients.

\mathcal{S}_2 is corrupt. We provide the formal simulator in Fig. 17, 18 and argue indistinguishability as follows.

HYB₀ : The real world execution of the protocol.

HYB₁ : Same as HYB_0 , except Sim aborts if a malicious client i has provided inconsistent u_i and v_i inputs to \mathcal{S}_0 and \mathcal{S}_1 and yet passed the batched consistency check π_{check} . The two hybrids are indistinguishable due to the correctness of π_{check} .




Fig. 17: Simulation Algorithm against malicious corruption of server \mathcal{S}_2 and ℓ' clients (Continues in Fig. 18).

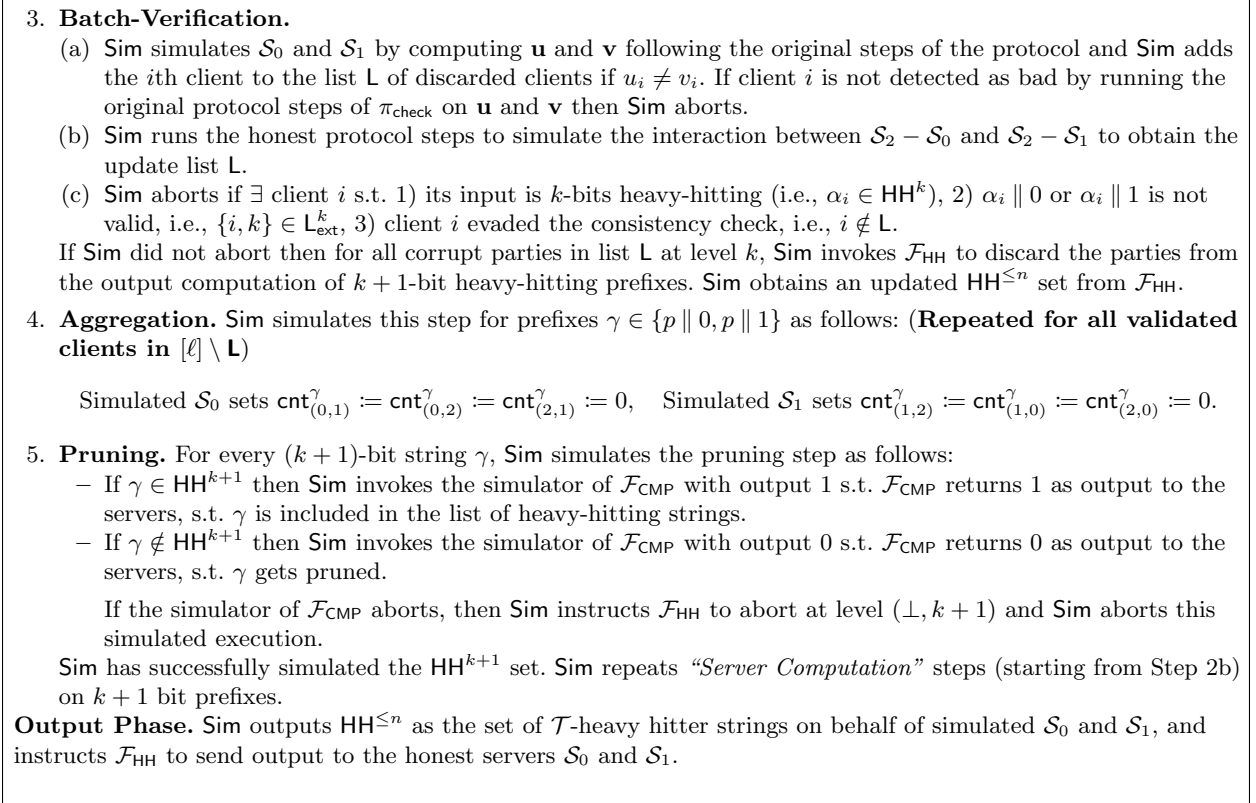


Fig. 18: Simulation Algorithm against malicious corruption of server \mathcal{S}_2 and ℓ' clients.

HYB₂ : Same as HYB₁, except the Sim extracts the corrupt client’s inputs using the three pairs of DPF keys. Then Sim runs Step 3c of simulated Batch-Verification, i.e., Sim aborts if 1) the client’s input α_i is k -bits heavy-hitting, 2) $\alpha_i \parallel 0$ or $\alpha_1 \parallel 1$ is invalid, and 3) client evaded the Batch-Verification check for the sessions run between honest servers. The two hybrids are indistinguishable due to the verifiability property of VIDPF in the random oracle model. This occurs when the client successfully evades the input extraction process of VIDPF by providing malformed VIDPF keys and yet passes the batch verification checks.

HYB₃ : Same as HYB₂, except Sim invokes \mathcal{F}_{HH} with the extracted inputs to obtain the $\text{HH}^{\leq n}$ set and simulates \mathcal{F}_{CMP} based on whether a prefix γ is in $\text{HH}^{\leq n}$ or not. The two hybrids are indistinguishable against a corrupt server \mathcal{S}_2 in the \mathcal{F}_{CMP} -model.

HYB₄ : Same as HYB₃, except Sim simulates the DPF key generation for the honest clients with input $(\alpha, (\beta_1, \dots, \beta_n)) = (1, (1, \dots, 1))$ and sets the counters to 0s in the aggregation step. Indistinguishable due to VIDPF input privacy. The 0-valued counters in the aggregation step are identically distributed to the actual aggregation counters since HYB₃ and HYB₄ are in the \mathcal{F}_{CMP} -model. This is the ideal world execution of the protocol, completing our simulation.

Either \mathcal{S}_0 or \mathcal{S}_1 is corrupt. Next, we consider the case where either \mathcal{S}_0 or \mathcal{S}_1 is corrupt along with ℓ' clients. We provide the simulator in Fig. 19, 20 and argue indistinguishability as follows. (This case is similar to the case where \mathcal{S}_1 is corrupt along with ℓ' clients.)

HYB₀ : The real world execution of the protocol.

HYB₁ : Same as HYB₀, except Sim aborts if a malicious client i provided values $(R_{(2,0)}^k, R_{(2,1)}^k)$ to \mathcal{S}_2 and values $(R_{(2,0)}^k, R_{(1,2)}^k)$ to \mathcal{S}_1 such that they are not equal, and yet client i passed the batched consistency check π_{check} . The two hybrids are indistinguishable due to the correctness of π_{check} .

HYB₂ : Same as HYB₁, except Sim extracts the corrupt client’s inputs following the extraction algorithm using the pair of DPF keys. Then Sim runs Step 3d of simulated Batch-Verification, i.e., Sim aborts if 1) the client’s input α_i is k -bits heavy-hitting, 2) $\alpha_i \parallel 0$ or $\alpha_1 \parallel 1$ is invalid, and 3) client evaded the Batch-Verification check for the sessions run between honest servers. The two hybrids are indistinguishable due to the verifiability property of VIDPF in the random oracle model. This occurs when a malicious client successfully evades the input extraction process of VIDPF by providing malformed VIDPF keys and yet passes the batch verification checks performed on the VIDPF proofs.

HYB₃ : Same as HYB₂, except Sim invokes \mathcal{F}_{HH} with the extracted inputs to obtain $\text{HH}^{\leq n}$ set and simulates the \mathcal{F}_{CMP} functionality based on whether a prefix γ is in $\text{HH}^{\leq n}$ or not. The two hybrids are indistinguishable against a corrupt server \mathcal{S}_0 in the \mathcal{F}_{CMP} -model.

HYB₄ : Same as HYB₃, except Sim simulates the key generation for the honest clients with $(\alpha, (\beta_1, \dots, \beta_n)) = (1, (1, \dots, 1))$ as input and sets the counters to 0s in the aggregation step. Indistinguishable due to VIDPF input privacy. The 0-valued counters in the aggregation step are identically distributed to the actual aggregation counters since HYB₃ and HYB₄ are in the \mathcal{F}_{CMP} -model. This is the ideal world execution of the protocol, completing our simulation algorithm. □

D Analysis of Batched Consistency check

We recall the batched consistency check in Fig. 9. \mathcal{P}_0 and \mathcal{P}_1 hash their leaves and verify the equality of their Merkle tree roots R_0 and R_1 . If the roots are equal then all the leaves are equal. Otherwise, the parties verify the equality of the left and the right children of the root node. If the left (resp. right) children are equal across the parties then the left (resp. right) subtrees are equal. If the left (resp. right) children are different, then the parties apply the above algorithm to the left (resp. right) subtree. Proceeding iteratively down the tree, the parties identify the malformed leaves as \overline{N}_0^K and \overline{N}_1^K where the two trees differ. Then they match them with their initial lists of input sets \mathbf{u} and \mathbf{v} to identify the indices where they differ and then store those indices in L .

Simulator Sim for maliciously corrupt ℓ' number of clients and server \mathcal{S}_0	
<p>Corruption: ℓ' number of clients and server \mathcal{S}_0 are maliciously corrupt. The rest $\ell - \ell'$ clients and servers ($\mathcal{S}_1, \mathcal{S}_2$) are simulated by simulator Sim. Without loss of generality, we will assume that \mathcal{S}_0 is corrupt; the case where \mathcal{S}_1 is corrupt is symmetric.</p> <p>Primitive: VIDPF := (Gen, EvalPref, EvalNext) is a verifiable incremental DPF. $H_1, H_2 : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ are random oracles.</p>	
<p>Client \mathcal{C} Computation. (Repeated for ℓ clients)</p> <p>1. <i>If the client is honest:</i> Sim simulates the client by preparing three pairs of DPF keys with input 1 and output values $(1, \dots, 1)$.</p> $(\text{key}_{(0,1)}, \text{key}_{(1,0)}) := \text{Gen}(1^\kappa, 1^n, 1, (1, \dots, 1), \mathbb{G}), \quad (\text{key}_{(1,2)}, \text{key}_{(2,1)}) := \text{Gen}(1^\kappa, 1^n, 1, (1, \dots, 1), \mathbb{G}),$ $(\text{key}_{(2,0)}, \text{key}_{(0,2)}) := \text{Gen}(1^\kappa, 1^n, 1, (1, \dots, 1), \mathbb{G})$ <p>Sim sends $(\text{key}_{(0,1)}, \text{key}_{(0,2)}, \text{key}_{(2,1)})$ to \mathcal{S}_0, $(\text{key}_{(1,0)}, \text{key}_{(1,2)}, \text{key}_{(2,0)})$ to \mathcal{S}_1 and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ to \mathcal{S}_2 on behalf of the client.</p> <p>2. <i>If the client is corrupt:</i> Client sends $(\text{key}_{(0,1)}, \text{key}_{(0,2)}, \text{key}_{(2,1)})$ to \mathcal{S}_0, $(\text{key}_{(1,0)}, \text{key}_{(1,2)}, \text{key}_{(2,0)})$ to \mathcal{S}_1 and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ to \mathcal{S}_2.</p>	
<p>Server Computation. (<i>Simulator Sim initializes a list $L_{\text{ext}} = \{\}$ and $L_{\text{inp}} = \{\}$, and simulates \mathcal{S}_1 and \mathcal{S}_2)</i></p> <p>For each corrupt client i, the simulator performs the following for input extraction: (Repeated for ℓ' corrupt clients)</p> <ol style="list-style-type: none"> 1. Sim extracts the corrupt client's input $(\alpha'_i, \beta'_{i,1}, \dots, \beta'_{i,n})$ from the pair of DPF keys - $\text{key}_{(1,2)}$ and $\text{key}_{(2,1)}$, provided by client i. 2. If the corrupt client has not provided a valid input at level j, i.e., $\exists j \in [n]$ s.t. $\beta'_j \neq 1$ (for the smallest j), then Sim truncates the extracted input of client i to the first j bits of α_i as $\alpha_i := \alpha_{i, \leq j-1}$. Sim sets $L_{\text{ext}}^{j-1} = L_{\text{ext}}^{j-1} \cup \{i, j-1\}$ and updates $L_{\text{ext}} = L_{\text{ext}} \cup L_{\text{ext}}^{j-1}$ to denote that the ith client's input is valid only till level $j-1$. 3. Sim stores the extracted input (after necessary truncation) α_i for client i in a list L_{inp} as $L_{\text{inp}} := L_{\text{inp}} \cup \{i, \alpha_i\}$. <p>After running the above extraction process for all corrupt clients, Sim invokes \mathcal{F}_{HH} with the input list L_{inp} to obtain the output set of \mathcal{T}-heavy hitting prefixes as $\text{HH}^{\leq n}$. The functionality \mathcal{F}_{HH} waits for further instructions from the ideal world adversary Sim.</p> <p>Repeat the following steps for length of k bits, where $k \in [0, \dots, n-1]$:</p> <ol style="list-style-type: none"> 1. Initialization. For prefix $p \in \text{HH}^k$, Sim initialize server \mathcal{S}_1's and \mathcal{S}_2's aggregation variables for prefixes $\gamma \in \{p \parallel 0, p \parallel 1\}$ as follows: $\text{Simulated } \mathcal{S}_1 \text{ sets } \text{cnt}_{(1,2)}^\gamma := \text{cnt}_{(1,0)}^\gamma := \text{cnt}_{(2,0)}^\gamma := 0, \quad \text{Simulated } \mathcal{S}_2 \text{ sets } \text{cnt}_{(2,0)}^\gamma := \text{cnt}_{(2,1)}^\gamma := 0.$ 2. VIDPF Evaluation. For prefix $p \in \text{HH}^{\leq k}$, Sim simulates \mathcal{S}_1 and \mathcal{S}_2 by running the original protocol steps. (Repeated for ℓ clients) 	

Fig. 19: Simulation Algorithm against malicious corruption of server \mathcal{S}_0 and ℓ' clients (Continues in Fig. 20).

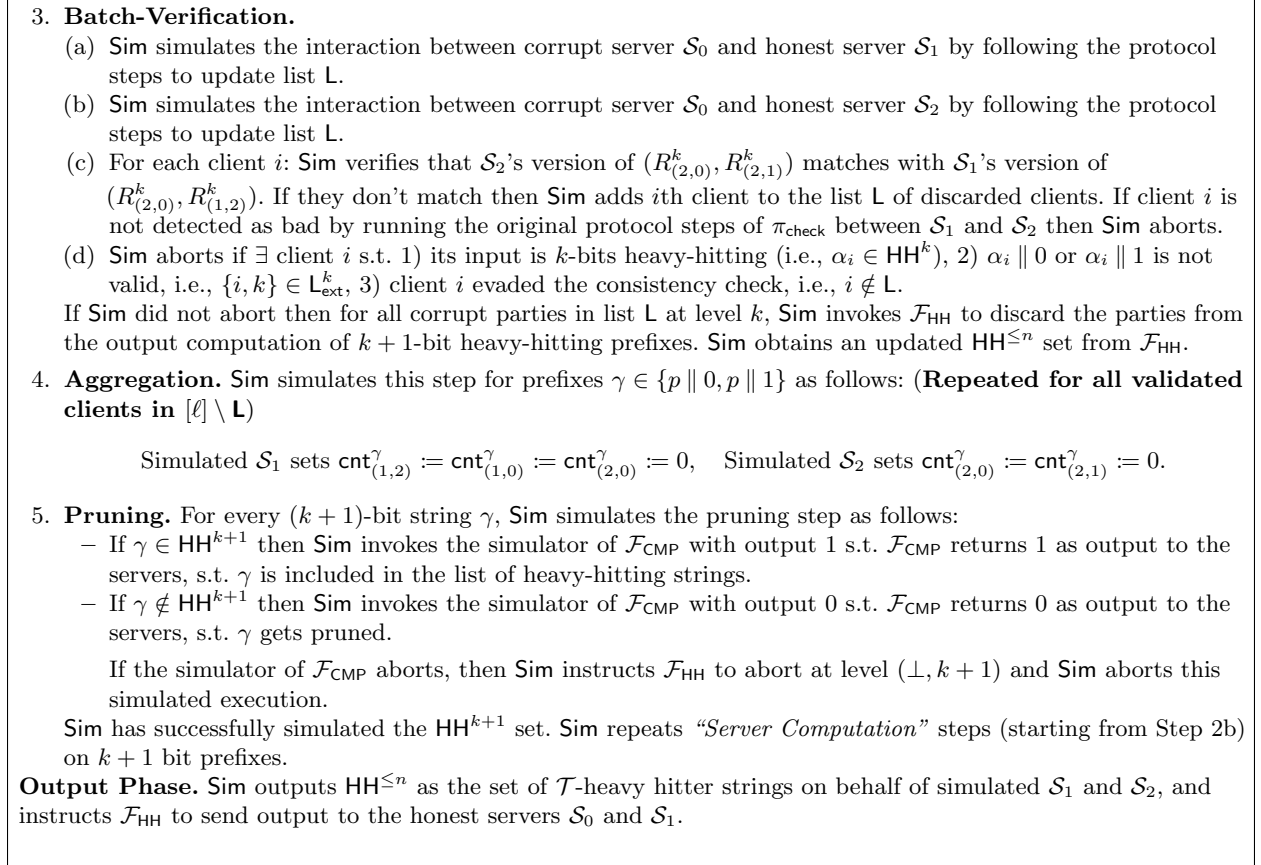


Fig. 20: Simulation Algorithm against malicious corruption of server \mathcal{S}_0 and ℓ' clients.

π_{check} requires $\mathbf{K} + 1$ rounds of communication, where $\mathbf{K} = \lceil \log_2 \ell \rceil$. Next, we demonstrate that if ℓ' out of ℓ leaves differ, then the total communication is $\mathcal{O}(\ell'(\log_2 \frac{\ell}{\ell'}))$ hashes. The *Root Computation* is local and *Root Verification* communicates two hashes. During *Leaf Identification*, the parties communicate 4 hashes for each unequal node. At the root layer, only the roots are different. At the next layer, both children can differ. More generally, at layer $k \in [\mathbf{K}]$, there can be at most $\min(2^k, \ell')$ unequal nodes.

The total communicated hashes are as follows:

$$\begin{aligned}
& 2 + 4 \times (\min(2^0, \ell') + \dots + \min(2^{\lceil \log_2 \ell \rceil}, \ell')) \\
&= 2 + 4 \times (1 + 2 + \dots + 2^{\lceil \log_2 \ell' \rceil} + \ell' + \ell' + \dots + \ell') \\
&\leq 2 + 4 \times (2^{\ell'} + \ell' \times (\lceil \log_2 \ell \rceil - \lceil \log_2 \ell' \rceil)) \\
&\approx 8\ell' + 4\ell'(\log_2 \ell - \log_2 \ell') = 4\ell'(\log_2 \frac{\ell}{\ell'} + 2).
\end{aligned}$$

We observe that the current version of π_{check} communicates roughly $4\ell'(\log_2 \frac{\ell}{\ell'} + 2)$ hashes. This can be further optimized to $2\ell'(\log_2 \frac{\ell}{\ell'} + 2)$ where only one server communicates at each level.

E Heavy Hitters with different Thresholds

Our protocol allows us to consider different heavy hitter thresholds \mathcal{T}_i based on some pre-agreed strings $x_i \in \mathbf{X}$ by the servers. This can be beneficial for traffic avoidance since different roads may have different traffic densities. For example, highways are busier than smaller suburban roads. The servers can take that

Different Threshold Heavy Hitters from \mathcal{T} -prefix count queries

Parameters: Threshold $\mathcal{T}_i \in \mathbb{N}$, for string $x_i \in \mathbf{X}$ where $|\mathbf{X}| = m$, and string length $n \in \mathbb{N}$.

Inputs: The algorithm has no explicit input. It has access to t -prefix count query oracle $\Omega_{\alpha_1, \dots, \alpha_\ell}(p, t)$ for securely computing t -prefix-count queries over prefix p for strings $\alpha_1, \dots, \alpha_\ell$.

Outputs: The set of heavy-hitter strings in $\alpha_1, \alpha_2, \dots, \alpha_\ell$.

Algorithm:

- Initialize $\text{HH}^{\leq n} = \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\} := \{\{\epsilon\}, \emptyset, \dots, \emptyset\}$, where HH^0 contains empty string ϵ and $\text{HH}^1, \dots, \text{HH}^n$ are empty sets.
- Set $\mathcal{T} := \min(\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_m)$.
- For each prefix $p \in \text{HH}^k$ of length k -bits in set HH^k (where $k = 0, 1, 2, \dots, n-2$):
 - If** $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel 0, \mathcal{T}) = 1$, **then** $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \{p \parallel 0\}$.
 - If** $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel 1, \mathcal{T}) = 1$, **then** $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \{p \parallel 1\}$.
- For each prefix $p \in \text{HH}^{n-1}$, perform the following:
 - If** $\exists x_i \in \mathbf{X}$ such that $(p \parallel 0) = x_i$ and $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel 0, \mathcal{T}_i) = 1$, **then** $\text{HH}^n := \text{HH}^n \cup \{p \parallel 0\}$.
 - If** $\exists x_i \in \mathbf{X}$ such that $(p \parallel 1) = x_i$ and $\Omega_{\alpha_1, \dots, \alpha_\ell}(p \parallel 1, \mathcal{T}_i) = 1$, **then** $\text{HH}^n := \text{HH}^n \cup \{p \parallel 1\}$.
- Output \mathcal{T} -heavy hitters $\text{HH}^{\leq n} := \{\text{HH}^0, \text{HH}^1, \dots, \text{HH}^n\}$.

Fig. 21: Algorithm for computing heavy hitters with different thresholds from \mathcal{T} -prefix count queries.

into consideration during evaluation, and use higher \mathcal{T} s for highways (since there are more vehicles), and lower thresholds for smaller roads.

We present our algorithm to compute heavy-hitters with different thresholds \mathcal{T}_i for string $x_i \in \mathbf{X}$ from \mathcal{T} -prefix oracle query in Fig. 21. The prefix oracle query with different thresholds can be computed using a simple modification to protocol π_{HH} , where the pruning at the leaf layer is performed based on the threshold \mathcal{T}_i for a given string $x_i \in \mathbf{X}$ instead of a fixed threshold \mathcal{T} .

F Supporting Differential Privacy

It is straightforward to complement PLASMA with ϵ -differential privacy techniques and ensure that the presence or absence of a single client does not reveal anything about their data [DMNS06]. In this case, running two instances of PLASMA, one with $\ell - 1$ clients and another just by adding client \mathcal{C} , should protect the private data of the new client from anyone observing the outputs of the two protocols. Additionally, honest clients should not be able to be identified when a malicious server attempts to ignore honest client data to infer their inputs based on the protocol output. Therefore, PLASMA is directly compatible with the well-studied techniques from [DKM⁺06, EDG14] and can adopt a similar approach as Poplar to bound the amount of information that an adversary \mathcal{A} can deduce from PLASMA’s output. Like Poplar, we need to ensure that the outputs of these prefix-count oracle queries are differentially private, which can be achieved by introducing noise on the oracle’s output with parameter $1/\epsilon$ from a Laplace distribution.

G Communication Cost of Asharov et al. [AHI⁺22]

We analyze the total server-to-server communication cost for the sorting-based protocol of Asharov et al. [AHI⁺22] (considering that its implementation is not open-source). We start from the optimized semi-honest communication cost from Appendix A.3 of [AHI⁺22], shown below: $mn(\frac{7}{3} + \frac{32}{9}||R||) + 3m||R|| + 2m||R' ||$ bits.

We ignore the R' term since it is a payload. For malicious security, the protocol requires two times the semi-honest protocol, and additionally, the ring needs to be a field of size 2^κ size for $2^{-\kappa}$ failure probability. This leads us to the optimized malicious sorting protocol communication cost of: $2mn(\frac{7}{3} + \frac{32}{9}\kappa) + 3m\kappa$.

The heavy hitters protocol requires the following for each item out of the total m items:

- Compute two secure comparisons over n bits. Assuming the state-of-the-art secure comparison protocol of Rabbit [MRVW21, Fig. 6], we get $\geq 4mn \log n$ from `LTBits` and `BitAdder` as well as mn to open the values.
- One secure multiplication over two secret shared n -bit variables: For m values it would be at least mn bits.
- Secure shuffling over and n -bit secret shared value, where the semi-honest shuffling takes $2m$ field element communication.

Asharov et al. [AHI⁺22] considers the compiler of Chida *et al.* [CGH⁺18] that converts a semi-honest protocol to a malicious protocol. However, this results in increased communication cost (i.e., $2\times$ the semi-honest cost): $2(4mn \log n + mn + 2mn) = 8mn \cdot \log n + 6mn$. The per-server communication cost for their maliciously secure heavy-hitters protocol is at least:

$$2mn\left(\frac{7}{3} + \frac{32}{9} \cdot \kappa\right) + 3m\kappa + 8mn \log n + 6mn \text{ bits.}$$

Setting the security parameter κ to 60 bits, the number of items m to 10^6 , and the number of bits of each item n to 256 bits we get that the communication cost should be at least:

$$2 \cdot 10^6 \cdot 256\left(\frac{7}{3} + \frac{32}{9}60\right) + 3 \cdot 10^6 \cdot 60 + (8 \cdot 10^6 \cdot 256 \cdot \log 256 + 6 \cdot 10^6 \cdot 256) = 14.96 \text{ giga bytes}$$

Therefore, the total server-server communication cost is at least $14.96 \cdot 3 \approx 45$ gigabytes for computing the heavy hitters over 256-bit keys between three servers for 10^6 clients.

H Private Histogram Protocol

We present our histogram protocol π_{HIST} in Fig. 22, 23 for the sake of completeness. The histogram protocol is a building block for our heavy-hitters protocol and *is not* our final protocol. It suffers from the limitation that the client’s input should lie in the subset \mathbf{X} that the servers evaluate, i.e., $\alpha_i \in \mathbf{X}$ for $i \in [\ell]$. This leaks whether the client’s input lies in \mathbf{X} or not based on whether the evaluated DPF output in the consistency check is 0 or not. This issue can be addressed by using techniques from Section 3.4, mainly replacing the VDPF with a VIDPF, and using the four consistency checks discussed in Section 3.4.

Private Histogram π_{HIST}

We denote a vector $\mathbf{Y} \in \mathbb{F}^m$ component-wise as $\mathbf{Y} := \{y_1, y_2, \dots, y_m\}$, where $y_j \in \mathbb{F}$ for $j \in [m]$.

- **Input:** Each client \mathcal{C}_i has an input point $\alpha_i \in \mathbf{X}$ for $i \in [\ell]$ and $m := |\mathbf{X}|$.
- **Output:** $\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2$ output a histogram of the ℓ clients' data. If the servers abort then it denotes a malicious server involvement.
- **Primitive:** $\text{VDPF} := (\text{Gen}, \text{BatchEval})$ is a verifiable distributed point function. $\text{H} : \{0, 1\}^* \rightarrow \{0, 1\}^\kappa$ is a random oracle.

Client \mathcal{C} Computation. (Repeated for ℓ clients, each of which has their own private input α)

1. Client \mathcal{C} with input α prepares three pairs of DPF keys with independent randomness $u, v, w \xleftarrow{r} \{0, 1\}^\kappa$, as follows:

$$(\text{key}_{(0,1)}, \text{key}_{(1,0)}) := \text{Gen}(1^\kappa, \alpha, 1, \mathbb{G}), \quad (\text{key}_{(1,2)}, \text{key}_{(2,1)}) := \text{Gen}(1^\kappa, \alpha, 1, \mathbb{G}),$$

$$(\text{key}_{(2,0)}, \text{key}_{(0,2)}) := \text{Gen}(1^\kappa, \alpha, 1, \mathbb{G})$$
2. The client sends $(\text{key}_{(0,1)}, \text{key}_{(0,2)}, \text{key}_{(2,1)})$ to \mathcal{S}_0 , $(\text{key}_{(1,0)}, \text{key}_{(1,2)}, \text{key}_{(2,0)})$ to \mathcal{S}_1 and $(\text{key}_{(2,1)}, \text{key}_{(2,0)})$ to \mathcal{S}_2 .

Server Computation.

If this is the first client, each server \mathcal{S}_b initializes $\text{HIST}_{(b,b+1)}$ and $\text{HIST}_{(b+1,b)}$ for $b \in \{0, 1, 2\}$ as follows:

\mathcal{S}_0 initializes $\text{HIST}_{(0,1)} := 0^m, \text{HIST}_{(0,2)} := 0^m$, and $\text{HIST}_{(2,1)} := 0^m$

\mathcal{S}_1 initializes $\text{HIST}_{(1,2)} := 0^m, \text{HIST}_{(1,0)} := 0^m$, and $\text{HIST}_{(2,0)} := 0^m$,

\mathcal{S}_2 initializes $\text{HIST}_{(2,0)} := 0^m$ and $\text{HIST}_{(2,1)} := 0^m$

1. **VDPF Evaluation:** Each server \mathcal{S}_b computes $\mathbf{Y}_{(b,b+1)}$ and $\mathbf{Y}_{(b,b+2)}$ for $b \in \{0, 1, 2\}$ as follows: (Repeated for ℓ clients)

\mathcal{S}_0 computes $\mathbf{Y}_{(0,1)}, \pi_{(0,1)} := \text{VDPF.BatchEval}(0, \text{key}_{(0,1)}, \mathbf{X})$

and $\mathbf{Y}_{(0,2)}, \pi_{(0,2)} := \text{VDPF.BatchEval}(1, \text{key}_{(0,2)}, \mathbf{X})$

\mathcal{S}_1 computes $\mathbf{Y}_{(1,2)}, \pi_{(1,2)} := \text{VDPF.BatchEval}(0, \text{key}_{(1,2)}, \mathbf{X})$

and $\mathbf{Y}_{(1,0)}, \pi_{(1,0)} := \text{VDPF.BatchEval}(1, \text{key}_{(1,0)}, \mathbf{X})$

\mathcal{S}_0 and \mathcal{S}_2 compute $\mathbf{Y}_{(2,1)}, \pi_{(2,1)} := \text{VDPF.BatchEval}(1, \text{key}_{(2,1)}, \mathbf{X})$

\mathcal{S}_1 and \mathcal{S}_2 compute $\mathbf{Y}_{(2,0)}, \pi_{(2,0)} := \text{VDPF.BatchEval}(0, \text{key}_{(2,0)}, \mathbf{X})$

Each server \mathcal{S}_b computes $\tau_{(b,b+1)}$ and $\tau_{(b,b+2)}$ for $b \in \{0, 1, 2\}$ as follows:

\mathcal{S}_0 parses $\mathbf{Y}_{(0,1)} = \{y_{(0,1),1}, y_{(0,1),2}, \dots, y_{(0,1),m}\}$ and computes $\tau_{(0,1)} := \sum_{j=1}^m y_{(0,1),j}$

\mathcal{S}_0 parses $\mathbf{Y}_{(0,2)} = \{y_{(0,2),1}, y_{(0,2),2}, \dots, y_{(0,2),m}\}$ and computes $\tau_{(0,2)} := \sum_{j=1}^m y_{(0,2),j}$

\mathcal{S}_1 parses $\mathbf{Y}_{(1,2)} = \{y_{(1,2),1}, y_{(1,2),2}, \dots, y_{(1,2),m}\}$ and computes $\tau_{(1,2)} := \sum_{j=1}^m y_{(1,2),j}$

\mathcal{S}_1 parses $\mathbf{Y}_{(1,0)} = \{y_{(1,0),1}, y_{(1,0),2}, \dots, y_{(1,0),m}\}$ and computes $\tau_{(1,0)} := \sum_{j=1}^m y_{(1,0),j}$

\mathcal{S}_1 and \mathcal{S}_2 parse $\mathbf{Y}_{(2,0)} = \{y_{(2,0),1}, y_{(2,0),2}, \dots, y_{(2,0),m}\}$ and compute $\tau_{(2,0)} := \sum_{j=1}^m y_{(2,0),j}$

\mathcal{S}_0 and \mathcal{S}_2 parse $\mathbf{Y}_{(2,1)} = \{y_{(2,1),1}, y_{(2,1),2}, \dots, y_{(2,1),m}\}$ and compute $\tau_{(2,1)} := \sum_{j=1}^m y_{(2,1),j}$

\mathcal{S}_0 computes $h_0 := \text{H}(\mathbf{Y}_{(0,1)} - \mathbf{Y}_{(0,2)} \parallel \mathbf{Y}_{(0,2)} - \mathbf{Y}_{(2,1)})$ and \mathcal{S}_1 computes $h_1 := \text{H}(\mathbf{Y}_{(2,0)} - \mathbf{Y}_{(1,0)} \parallel \mathbf{Y}_{(1,2)} - \mathbf{Y}_{(2,0)})$.

Fig. 22: Private Histogram Protocol π_{HIST} (continues in Fig. 23).

Private Histogram π_{HIST}

We denote a vector $\mathbf{Y} \in \mathbb{F}^m$ component-wise as $\mathbf{Y} := \{y_1, y_2, \dots, y_m\}$, where $y_j \in \mathbb{F}$ for $j \in [m]$.

Server Computation (Cont.).

2. **Batch-Verification.** The servers batch-verify the client inputs for all three sessions and across the three sessions by invoking π_{check} (Fig. 9):

- (a) \mathcal{S}_0 sets $u_i := \{(\pi_{(0,1)}, \pi_{(0,2)}, \pi_{(2,1)}, \tau_{(0,1)}, \tau_{(0,2)}, \tau_{(2,1)}, h_0)\}$ values for client $i \in [\ell]$. \mathcal{S}_1 sets $v_i := \{(\pi_{(1,0)}, \pi_{(2,0)}, \pi_{(1,2)}, 1 - \tau_{(1,0)}, 1 - \tau_{(2,0)}, 1 - \tau_{(1,2)}, h_1)\}$ values for client $i \in [\ell]$. \mathcal{S}_0 sets $\mathbf{u} := \{u_i\}_{i \in [\ell]}$ and \mathcal{S}_1 sets $\mathbf{v} := \{v_i\}_{i \in [\ell]}$. \mathcal{S}_0 and \mathcal{S}_1 batch-verify all the client inputs by computing the bit ver and list \mathbf{L} (comprising of invalid client inputs) by running π_{check} with inputs \mathbf{u} and \mathbf{v} respectively: $(\text{ver}, \mathbf{L}) := \pi_{\text{check}}(\mathbf{u}, \mathbf{v})$:

$$\text{ver} := 0 \text{ if there exists a client such that : } (\pi_{(0,1)} \neq \pi_{(1,0)}) \vee (\pi_{(0,2)} \neq \pi_{(2,0)}) \vee (\pi_{(2,1)} \neq \pi_{(1,2)}) \vee$$

$$(\tau_{(0,1)} + \tau_{(1,0)} \neq 1) \vee (\tau_{(0,2)} + \tau_{(2,0)} \neq 1) \vee (\tau_{(2,1)} + \tau_{(1,2)} \neq 1) \vee (h_0 \neq h_1)$$

and $\mathbf{L} := \{\text{list of invalid clients' that failed to pass the above checks}\}$. If $\text{ver} = 1$, then all the clients' inputs are valid.

- (b) \mathcal{S}_2 possesses $\pi_{(2,0)}, \pi_{(2,1)}, \tau_{(2,0)}, \tau_{(2,1)}$ values for each client. \mathcal{S}_2 verifies that \mathcal{S}_2 's version of $\pi_{(2,1)}, \tau_{(2,1)}$ matches with \mathcal{S}_0 's version of $\pi_{(2,1)}, \tau_{(2,1)}$. \mathcal{S}_2 also attests that \mathcal{S}_2 's version of $R_{(2,0)}^k$ matches with \mathcal{S}_0 's version of $\pi_{(0,2)}, \tau_{(0,2)}$ by computing $(\text{ver}', \mathbf{L}')$ as follows:

$$(\text{ver}', \mathbf{L}') := \pi_{\text{check}}(\{\pi_{(2,1)}, \tau_{(2,1)}, \pi_{(2,0)}, \tau_{(2,0)}\}_{\ell \text{ clients of } \mathcal{S}_2}, \{\pi_{(2,1)}, \tau_{(2,1)}, \pi_{(0,2)}, \tau_{(0,2)}\}_{\ell \text{ clients of } \mathcal{S}_0}).$$

- (c) \mathcal{S}_2 verifies that \mathcal{S}_2 's version of $\pi_{(2,0)}, \tau_{(2,0)}$ matches with \mathcal{S}_1 's version of $\pi_{(2,0)}, \tau_{(2,0)}$. \mathcal{S}_2 also attests that \mathcal{S}_2 's version of $\pi_{(2,1)}, \tau_{(2,1)}$ matches with \mathcal{S}_1 's version of $\pi_{(1,2)}, \tau_{(1,2)}$ by computing $(\text{ver}'', \mathbf{L}'')$ as follows:

$$(\text{ver}'', \mathbf{L}'') := \pi_{\text{check}}(\{\pi_{(2,0)}, \tau_{(2,0)}, \pi_{(2,1)}, \tau_{(2,1)}\}_{\ell \text{ clients of } \mathcal{S}_2}, \{\pi_{(2,0)}, \tau_{(2,0)}, \pi_{(1,2)}, \tau_{(1,2)}\}_{\ell \text{ clients of } \mathcal{S}_0}).$$

After batch verification, the servers identify the list of bad clients as $\mathbf{L} := \mathbf{L} \cup \mathbf{L}' \cup \mathbf{L}''$. The servers ignore the inputs of all clients in \mathbf{L} .

The servers locally perform the following computation:

The servers aggregate all correct client inputs into the histogram as follows: **(Repeated for all validated clients in $[\ell] \setminus \mathbf{L}$)**

$$\mathcal{S}_0 \text{ updates } \text{HIST}_{(0,1)} := \text{HIST}_{(0,1)} + \mathbf{Y}_{(0,1)}, \text{HIST}_{(0,2)} := \text{HIST}_{(0,2)} + \mathbf{Y}_{(0,2)} \text{ and } \text{HIST}_{(2,1)} := \text{HIST}_{(2,1)} + \mathbf{Y}_{(2,1)}$$

$$\mathcal{S}_1 \text{ updates } \text{HIST}_{(1,2)} := \text{HIST}_{(1,2)} + \mathbf{Y}_{(1,2)}, \text{HIST}_{(1,0)} := \text{HIST}_{(1,0)} + \mathbf{Y}_{(1,0)} \text{ and } \text{HIST}_{(2,0)} := \text{HIST}_{(2,0)} + \mathbf{Y}_{(2,0)}$$

$$\mathcal{S}_2 \text{ updates } \text{HIST}_{(2,0)} := \text{HIST}_{(2,0)} + \mathbf{Y}_{(2,0)} \text{ and } \text{HIST}_{(2,1)} := \text{HIST}_{(2,1)} + \mathbf{Y}_{(2,1)}$$

Output Phase.

1. Each two servers \mathcal{S}_b and \mathcal{S}_{b+1} exchange $H(\text{HIST}_{(b,b+1)}, r_{(b,b+1)})$ and $H(\text{HIST}_{(b+1,b)}, r_{(b+1,b)})$ for random $r_{(b,b+1)}, r_{(b+1,b)} \xleftarrow{r} \{0, 1\}^\kappa$.
2. \mathcal{S}_0 sends $(\text{HIST}_{(0,1)}, \text{HIST}_{(0,2)}, \text{HIST}_{(2,1)}, r_{(0,1)}, r_{(0,2)})$ to \mathcal{S}_1 . \mathcal{S}_1 sends $(\text{HIST}_{(1,2)}, \text{HIST}_{(1,0)}, \text{HIST}_{(2,0)}, r_{(1,2)}, r_{(1,0)})$ to \mathcal{S}_0 . \mathcal{S}_2 broadcasts $(r_{(2,0)}, r_{(2,1)})$.
3. \mathcal{S}_0 and \mathcal{S}_1 verify the above hashes. If any of the hashes fail then the servers abort. Else, they perform the following:

$$\mathcal{S}_0 \text{ and } \mathcal{S}_1 \text{ compute } \text{HIST}_0 := \text{HIST}_{(0,1)} + \text{HIST}_{(1,0)}, \text{HIST}_1 := \text{HIST}_{(1,2)} + \text{HIST}_{(2,1)}, \text{ and}$$

$$\text{HIST}_2 := \text{HIST}_{(2,0)} + \text{HIST}_{(0,2)}$$

4. \mathcal{S}_0 and \mathcal{S}_1 abort if $\text{HIST}_0 \neq \text{HIST}_1$ or $\text{HIST}_1 \neq \text{HIST}_2$. Else, they output HIST where $\text{HIST} = \text{HIST}_0 = \text{HIST}_1 = \text{HIST}_2$.

Fig. 23: Private Histogram Protocol π_{HIST} (continuing from Fig. 22).