Mastic: Private Weighted Heavy-Hitters and Attribute-Based Metrics

Dimitris Mouris*
Nillion & University of Delaware
dimitris@nillion.com

Christopher Patton
Cloudflare, Inc.
cpatton@cloudflare.com

Hannah Davis
Seagate Technology
hannah.e.davis@seagate.com

Pratik Sarkar
Supra Research
iampratiksarkar@gmail.com

Nektarios Georgios Tsoutsos
University of Delaware
tsoutsos@udel.edu

Abstract

Insight into user experience and behavior is critical to the success of large software systems and web services. Gaining such insights, while preserving user privacy, is a significant challenge. Recent advancements in multi-party computation have made it practical to securely compute aggregates over secret shared data. Two such protocols have emerged as candidates for standardization at the IETF: Prio (NSDI 2017) for general-purpose statistics; and Poplar (IEEE S&P 2021) for heavy hitters, where the goal is to compute the most popular inputs held by users without learning the inputs themselves. While each of these protocols is well-suited to certain applications, there remain a number of use cases identified by IETF for which neither Prio nor Poplar is practical.

We introduce Mastic, a protocol for the following functionality: each of a large number of clients holds an *input* (e.g., a URL) and its corresponding *weight* (e.g., page load time); for a given candidate input (or prefix), a small number of non-colluding servers wish to securely aggregate the weights of clients that hold that input (or some input with that prefix), without learning the weights or which client holds which input. This functionality makes two new classes of applications possible. The first is a natural generalization of heavy hitters we call *weighted heavy-hitters*. The second is an enhancement of Prio-style metrics we call *attribute-based metrics* in which aggregates are grouped by hierarchical user attributes (e.g., their geographic location or software version). We demonstrate Mastic's practicality for these applications with a real-world example of each. We also compare our protocol with Prio and Poplar on a wide area network. Overall, we report over one order of magnitude performance improvement over Poplar for plain heavy-hitters and $1.5 - 2\times$ improvement over Prio for attribute-based metrics.

Keywords

Distributed point functions, zero-knowledge proofs, privacy enhancing technologies, secure multiparty computation.

1 Introduction

Critically important to the success of today's digital technology is the ability to gain insight into user behavior and experience. Web browsers, operating systems, and web services collect telemetry to detect performance issues, bugs, and security vulnerabilities [28]; advertisers track ad impressions to make sense of which ads drive revenue [24, 34, 38, 43]; and AI models are trained on user data for specific purposes, like detecting landmarks in photos [3] or classifying malicious behavior [16].

Enabling such use cases requires collecting privacy-sensitive measurements from users. However, often these measurements are not consumed directly, but only in some aggregated form, such as a summary statistic (e.g., mean, median, or standard deviation) [19], a probabilistic data structure (e.g., Bloom filter [8] or count-min sketch [18]), or machine learning model (such as linear regression or gradient descent) [20]. In these situations, to preserve user privacy it is desirable to collect only what the application needs – i.e., the *aggregate* – and avoid gathering raw measurements and computing over the plaintexts directly.

Many in the tech industry are investing in *multi-party computation (MPC)* to help address this problem [3, 17, 28, 30, 36, 42] as it allows multiple parties to jointly compute a function on private inputs. Much of this work revolves around a special class of lightweight and highly parallelizable MPC protocols undergoing standardization at IETF [5, 27]. Called *verifiable, distributed aggregation functions (VDAFs)* [22], these protocols delegate the computation of the aggregate to a small number of servers – typically two. Measurements are secret shared (i.e., cryptographically split) and uploaded to the servers such that no one server sees any measurement in the clear. In addition, the computation is *verifiable* in the sense that the servers are guaranteed to compute an aggregate of only valid measurements.[1] The validity of the measurement depends on the application: each measurement might be a bit (0 or 1) and the aggregate result would be the frequency of these binary outcomes; each measurement might fall within a predetermined integer range and the aggregate would be the sum of the integers; or each measurement might be a one-hot vector (i.e., everywhere zero except a single one) and the aggregate would be a histogram computed from the sum of the vectors.

---

[1] In verifiable distributed aggregation functions (VDAFs), verifiability refers to the ability of the servers to assert that the result was computed from valid inputs – not to be confused with the notion of publicly verifiable MPC [6].

VDAFs are being integrated into a variety of real-world applications as we speak. One of the keys to the success of these schemes is their flexibility. The seminal example is Prio, first proposed by Corrigan-Gibbs and Boneh [19] and now a candidate for standardization [5]. Prio can compute any aggregation function $F(m_1, \ldots, m_N)$ that can be represented as the sum of (some encoding of) the measurements $m_1, \ldots, m_N$. Validity is then defined by an arithmetic circuit evaluated over each encoded measurement and proven using a *zero-knowledge proof on secret shared data* [10]. This enables servers to privately validate that their shares of each m_i sum up to a valid measurement.

While many private aggregation problems can be solved with Prio, others do not fit cleanly into this rubric. The most important example for IETF currently is privately computing *heavy hitters*. Here each client C_i holds a private input $\alpha \in \{0,1\}^n$ and the servers' goal is to compute the subset of inputs held by at least T clients for some target *threshold* T. For example, if each input α represents a URL, then the T-*heavy-hitters* would represent the set of URLs visited by at least T users. Poplar [11] was the first solution for heavy hitters with enough scalability for use cases of practical interest. Poplar is based on the *function secret sharing* [13] paradigm, where each client produces secret shares of some function f such that the servers can compute secret shares of $f(x)$ for a chosen input x with minimal ($\mathcal{O}(|x|)$) communication overhead. In particular, the authors construct a secret sharing of a function $f_{\alpha,\beta}$ for which $f_{\alpha,\beta}(x) = \beta$ for each *prefix x* of bitstring α and $f_{\alpha,\beta}(x) = 0$ otherwise. Given such a scheme, called an *incremental distributed point function (IDPF)*, the servers can count how many of the inputs begin with a given candidate prefix.

Given the practical importance of the heavy hitters, Poplar is also being considered for standardization at IETF [5]. However, both Prio and Poplar have some functional and operational limitations observed by participants in the IETF working group.

First, many applications of Prio require breaking down metrics by client attributes, such as the HTTP user agent, software version, geographic location, etc. [26]. For example, a website developer may wish to track metrics across different web browsers in order to gain insight into performance differences between them. Moreover, it is desirable to break down metrics this way without reducing the anonymity set by revealing attributes to the server in plaintext. We refer to this feature as *attribute-based metrics*. Prio can be extended to support this feature, but the communication overhead would be too high to be practical in most cases. (See Section 4.2 for details.)

Second, a natural generalization of heavy hitters is to associate with each input α a *weight β* such that the servers compute the subset of inputs for which (some function of) the sum of the weights exceeds the threshold T. For example, if the input is a URL as before, and the weight is the page load time measured by the browser, the heavy hitters might represent the URLs with the highest average or median load times. We refer to this problem as *weighted heavy hitters*. Poplar solves only a special case of this problem, namely when $\beta \in \{0,1\}$.

Third, a feature of Poplar that limits its deployability is the procedure by which the servers verify that their shares of $f_{\alpha,\beta}(x)$ for each candidate prefix x add up to a valid value. The main drawback of their approach is that it requires two rounds of communication, whereas many VDAFs, like Prio, require just one round to complete

verification. Multiple rounds inhibits performance (especially over wide-area networks) and also requires the servers to keep state, which adds complexity to the protocol [27]. Ideally, our solution for (weighted) heavy hitters or (attribute-based) metrics would require just one round trip over the network.

An important goal of this paper is to advance the state of the art by devising a protocol that overcomes these limitations, fits well within the VDAF framework, and is also aligned with IETF's work on privacy-preserving measurement.

1.1 Our Contribution

We present *Mastic*, a new VDAF for the following functionality: Each client holds a private pair (α, β), where $\alpha \in \{0,1\}^n$ is called the *input* and β is the corresponding *weight*. For a given *candidate prefix* $x \in \{0,1\}^{\le n}$, Mastic allows the servers to aggregate the weights β for all clients for which x is a prefix of their input α. This functionality immediately enables the two classes of applications described above, namely attribute-based metrics and weighted heavy-hitters. Mastic also achieves our goal of completing verification in one round of communication, while also having zero bandwidth overhead compared to Poplar (for plain heavy-hitters).

Our contributions are summarized as follows:

(1) We introduce Mastic (Section 3), a flexible protocol oriented toward real-world use cases and operational considerations, which elegantly adopts two primitives from prior work: The first is the *Verifiable IDPF (VIDPF)* of Mouris et al. (PETS 2024 [35]), a lightweight extension of IDPF that leads to reduced round complexity compared to Poplar. The second is the zero-knowledge proof system on secret-shared data of Boneh et al. (CRYPTO 2019 [10]), giving us the flexibility needed to support our use cases. As we will see, the novel composition of these primitives enables us to achieve tasks that were infeasible in prior works. Notably, Mastic's simplicity makes it relatively easy to implement and offers confidence in its security.

(2) We prove Mastic is secure in the same threat model as Poplar. Specifically, Mastic ensures (i) *privacy of the input measurements* in the presence of malicious clients and one malicious server (i.e., in each round the servers learn only the intended output: a weighted prefix histogram) and (ii) *robustness of the aggregate result* in the presence of malicious clients (i.e., servers correctly compute aggregate of valid measurements only). See Section 5.

(3) We implement Mastic and demonstrate its suitability for two real-life applications of interest to IETF [29], namely *Network Error Logging (NEL)* [44] and *Browser Telemetry* [28]. In the first, we propose a privacy-preserving version of NEL, a tool that provides telemetry crucial for detecting and diagnosing connectivity issues between clients and servers on the internet. In the second, we introduce attribute-based browser telemetry with improved usability compared to Prio without sacrificing privacy. We describe these applications in Section 4 and report our implementation and experimental evaluation in Section 6.

Table 1: Comparison of protocols from distributed ZKP and distributed point functions (DPFs).

Protocol[*]	Number of Servers	Rounds (per level) for Input Validation	Heavy Hitters	Attribute-Based Metrics	Robustness Against Malicious Clients & a Malicious Server
Prio [19]	≥ 2	1	○[†]	◐[‡]	○
Poplar [11]	2	2	●	○	○
Doplar [22]	2	1	●	○	○
PLASMA [35]	3	1	●	○	●
Mastic (this work)	2	1	●	●	○

[*] All protocols protect the honest clients' privacy against malicious clients and a malicious server. Additionally, all are robust against malicious clients.
[†] Prio could theoretically do heavy hitters but it would incur multiple rounds and leakages, rendering it impractical.
[‡] Prio is not able to perform attribute-based metrics but can do a simplified version which we refer to as "plain metrics" (see Section 4.2).

1.2 Related Works

We discuss related works for private heavy hitters and statistics along with their models and the applications they enable. In Table 1, we provide a theoretical comparison with works relying on distributed ZKPs and DPFs, which are most closely related to Mastic.

Private Statistics from Distributed Point Functions. Poplar [11] proposed IDPFs to address the problem of heavy hitters by computing the total number of private client inputs that begin with a given prefix. This incremental property allowed efficient evaluation of strings based on prefixes which was not possible with DPFs [13]. Poplar's servers ensure that the client measurements are valid based on a two-round MPC. PLASMA (PETS 2024 [35]) improved on IDPFs by proposing the Verifiable IDPF (VIDPF) primitive that combined IDPFs with the verifiability property of [23]. As a result, PLASMA relies solely on hashing rather than MPC operations to assert client measurement validity. PLASMA ensures robustness when one of the three servers colludes with the malicious clients; Mastic targets the two-server setting instead. Notably, Mastic can also be instantiated with three servers, thereby achieving security in the same threat model as PLASMA. The concurrent work of Doplar (PETS 2023 [22]) introduced a construction they also call a "VIDPF", but the construction is less efficient than that of [35]. Like Poplar and PLASMA, Doplar only solves the plain heavy hitters problem. Boneh et al. recently [12] generalized the MPC used in Poplar to verify richer sets of outputs, thereby making it suitable for weighted heavy hitters. However, their protocol requires more rounds of communication than ours.

Private Statistics from Differential Privacy. Differential Privacy (DP)-based techniques [39, 46] have been used to solve the problem of heavy hitters. Thus, they offer a trade-off between utility and privacy: to obtain higher data utility, they leak some information about the client measurements to the server. Other works that focus on stronger guarantees require at least two non-colluding servers.

The work of Anderson et al. [2] relies on oblivious shuffling to unlink reports from the clients that generated them, combines them with artificial data to add DP noise, and finally computes approximate statistics over the shuffled data.

A recent line of work beginning with Bell et al. [7] studies the related problem of aggregating *sparse histograms*. Like DPF-based protocols, their protocols also rely on two servers to achieve privacy. However, their protocols rely on shuffling and homomorphic encryption rather than secret sharing, leading to a different set of operational considerations for deployments. Notably, their protocols rely heavily on mechanisms for differential privacy and cannot be used for exact statistics. On the other hand, DPF-based protocols like Poplar, PLASMA, and Mastic can be composed with such mechanisms to enhance their privacy properties, albeit with lower utility than the protocol of Bell et al.

Statistics from (Non-DPF) MPC. The work of [9] utilizes general-purpose MPC [15, 32] for computing heavy hitters but results in prohibitively expensive solutions. Asharov et al. [4] propose an MPC sorting-based protocol for heavy hitters that improves on Poplar and PLASMA in terms of performance but requires significantly higher communication between the servers, which in many scenarios (especially over WAN) is not practical. Similarly, Vogue [31] relies on MPC sorting but inherently has high server-to-server communication. Both of these protocols specifically solve the plain heavy hitters problem and have not been extended to weighted heavy hitters. More specifically, the client votes for their input α by sharing α with the servers. The servers sort the shared inputs, and then in the aggregation phase, they consider a point to be heavy-hitting if there are more than a threshold number of votes for it. In the weighted version, the clients also have to share their weights for each point α, and the servers need to consider that during aggregation. The weights have to be sorted based on the input values, and then the weights corresponding to each unique input have to be obliviously added (using techniques from [37]) and then compared against the threshold. This would significantly affect the performance of the final protocol. We would also like to note that the weighted heavy hitter problem is not related to the recent notion of weighted MPC introduced in [25].

Heavy Hitters via Threshold Secret Sharing. A recent line of work beginning with Davidson et al. (CCS 2022 [21]) shows how to compute heavy hitters using a threshold secret sharing scheme along with other primitives. The high-level idea is that clients who generate the same measurement will arrange to upload shares of a t-of-n secret sharing of the measurement: once the aggregation server has received the desired threshold of shares t, it can immediately recover the underlying measurement. Clients submit their shares through an anonymizing proxy in order to ensure any recovered measurements cannot be linked to them. This scheme was later improved by Li et al. (USENIX 2024 [33]) by reducing the amount of information leaked to the attacker beyond the t-heavy-hitters. Their protocol leaks about the same amount of information as IDPF-based protocols do.
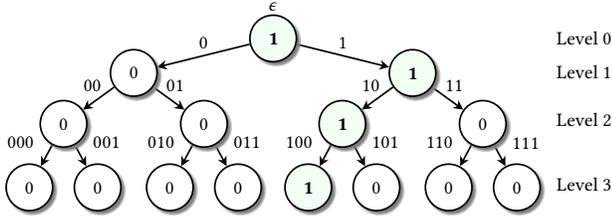
Figure 1: IDPF tree for input $\alpha = 100$ and weight $\beta = 1$.

2 Preliminaries

2.1 Notation

Fix a prime order field $\mathbb{F}$. We use $[\![\cdot]\!]$ to denote additive secret shares: if $x \in \mathbb{F}^m$ for some m, then $[\![x]\!]_0$ and $[\![x]\!]_1$ denote shares of x for which $x = [\![x]\!]_0 + [\![x]\!]_1$. We use $[n]$ to denote the set of integers $\{1, \ldots, n\}$. We use "$:=$" for assignment, "$\xleftarrow{\$} D$" for sampling uniformly from a finite set D, and "$\|$" for concatenation of strings. For bitstring s, we denote by $s|^n$ the n-bit prefix of s. We use square brackets to denote vectors, e.g., $v = [v_1, \ldots, v_n]$ is the length-n vector with elements $v_1, \ldots, v_n$. For any two sets D and R, we denote by $\mathcal{AF}(D, R)$ the set of all functions with domain D and range R. Our protocol involves a large number of *clients* and a small number of *servers*, which we also call *aggregators*. Let $\mathcal{S}_0$ and $\mathcal{S}_1$ denote the two aggregator servers.

2.2 Distributed Point Functions

Mastic is based on function secret sharing [13]. Secret sharing a function f allows the share holders to locally compute $[\![f(x)]\!]_0$, $[\![f(x)]\!]_1$ for a given input x without revealing anything about f. A *distributed point function (DPF)* [13] is a special case in which f is defined as follows. Fix $\alpha, \beta \in \{0, 1\}^n \times \mathbb{F}^m$ for some n, m: we call f the *point function* for α, β if $x = \alpha$ implies $f(x) = \beta$ and $x \neq \alpha$ implies $f(x) = 0^m$. Mastic requires two additional properties.

Incrementality. The *incremental DPF (IDPF)* of [11] secret shares an *incremental* point function, where $f(p) = \beta$ for any non-empty *prefix p of α*.[2] This allows for richer statistics, including (weighted) heavy hitters. Each share holder has a share of the *prefix tree* for α, β. As illustrated in Fig. 1, a prefix tree is a complete binary tree whose nodes are labeled with output values, which we call *weights*. Each path in this tree corresponds to a unique prefix; when we evaluate the incremental point function f at prefix p, the output gives the weight of the node that we reach when we traverse path p. For example, in Fig. 1: $f(1) = \beta$; $f(10) = \beta$; and $f(101) = 0^m$.

Verifiability. In many applications (including Mastic), the DPF shares are generated by an untrusted client, e.g., in a web browser. Since the shares reveal nothing about the underlying function, it is trivial for a malicious client to craft malformed shares that evaluate to something other than an (incremental) point function, thus breaking the correctness of the application. To detect such attacks, Boneh et al. [11] use a two-round MPC protocol between the servers to verify that for a given level of the prefix tree, at most one of the outputs has a non-zero value. de Castro and Polychroniadou [23]'s

verifiable DPF ensures a similar property for the DPF's output. However, their approach is based on hashing, is lighter weight, and requires just one round of communication rather than two.

2.3 Verifiable IDPFs (VIDPFs)

Our starting point for Mastic is the *verifiable IDPF (VIDPF)* of Mouris et al. [35]. This scheme adapts the techniques of de Castro and Polychroniadou to the IDPF of Boneh et al. to allow the servers to verify the correctness of the computation without revealing the input or its weight.

For each node it traverses in its share of the prefix tree, each server computes a short bit string π^p, where p is the path to the node from the root. We refer to this string as the *VIDPF proof* (or simply *proof*) for prefix p. Each proof is computed from the proof of its parent node in such a way that, if each server computes the same proof for p, then there is at most one prefix of length $|p|$ that evaluates to a non-zero value.

A VIDPF is comprised of the following algorithms:[3]

- $\mathcal{V}.\mathsf{Gen}(\alpha, \beta) \to (\mathsf{pub}, \mathsf{key}_0, \mathsf{key}_1)$: The *key generation* takes an *input* $\alpha \in \{0, 1\}^n$ and its *weight* $\beta \in \mathbb{F}^m$ and outputs the *keys* key_0 and key_1 and a *public share* pub.
- $\mathcal{V}.\mathsf{Eval}(\mathsf{key}_b, \mathsf{pub}, p, \mathsf{st}_b^{p'}, \pi_b^{p'}) \to ([\![y^p]\!]_b, \mathsf{st}_b^p, \pi_b^p)$: The *prefix evaluation algorithm* takes one of the keys and the public share and returns a share of a node of the prefix tree. It operates on a prefix $p \in \{0, 1\}^{\leq n}$ and the *state* $\mathsf{st}_b^{p'}$ and *proof* $\pi_b^{p'}$ for a prefix p' of p. (Usually p' is the parent of p, i.e., $p = p'\|z$ for some bit z.) Its outputs are the state and proof for p and a share of the weight y^p associated with p.[4] The proof and state for the root ($p' = \epsilon$) are defined to be ϵ.
- $\mathcal{V}.\mathsf{EvalRoot}(\mathsf{key}_b, \mathsf{pub}) \to [\![\beta]\!]_b$: The *root evaluation algorithm* takes in one of the keys and the public share and outputs the aggregator's share of β.
- $\mathcal{V}.\mathsf{Valid}(\pi_0^p, \pi_1^p) \to \{\mathsf{true}, \mathsf{false}\}$: The *validity algorithm* takes in both servers' proofs and determines whether they hold keys that represent a valid (i.e., one-hot) prefix tree. For the specific construction used by Mastic and [35], Valid simply checks equality of π_0^p and π_1^p. Therefore, it can be efficiently batch-evaluated over many prefixes using a collision-resistant hash function.

Since the specific VIDPF scheme we use proves its security in the random oracle model, we require it to define two sets Dom and Rng. Every security game using the scheme must begin by sampling a uniformly random function from the set of all functions with domain Dom and range Rng. It must then provide the adversary and Gen, Eval, and EvalRoot with oracle access to this function.

In Appendix A we define the properties of VIDPFs we need for Mastic. Briefly, $\mathcal{V}$ should be:

- *Correct*: when the client and aggregators are honest, the aggregators correctly evaluate shares of the prefix tree;

---

[2]Boneh et al. [11] consider a slightly richer function parameterized by $\beta_1, \ldots, \beta_n$ where β_i is the output for the length-i prefix of α. For our purposes, it is sufficient for each output to be the same.

[3]There are a few differences between our syntax and [35]. The main one is the addition of the root evaluation algorithm, which Mastic uses (and is implicit in [35]). We also have a public share, which contains parts of the keys that are sent to both servers (namely, correction words [35, Fig. 14]). Finally, we explicitly capture in syntax that validity only holds when both servers derive identical proofs.

[4]If p is a prefix of α, then $y^p = \beta$; otherwise $y^p = 0^m$.

- *Verifiable:* when the aggregators are honest, malicious clients cannot construct a public share and keys $(\mathsf{pub}, \mathsf{key}_0, \mathsf{key}_1)$ for which the evaluation tree contains more than one non-zero node at any level $k \in [n]$; and
- *Private:* the information revealed to each aggregator reveals nothing about an honest client's input.

PLASMA [35] proved the security of their VIDPF construction based on a pseudorandom generator and XOR-collision resistant hash function. The XOR-collision resistant hash is instantiated using a random oracle. Mastic uses the same construction based on a pseudorandom generator in the random oracle model.

Limitations. Observe that VIDPF already provides most of what we need for Mastic: using the same tree traversal and pruning strategy for PLASMA [35], modified slightly to use sums of weights rather than prefix counts (see Section 3 for details), we can compute the subset of inputs whose total weight exceeds the desired threshold. However, the degree of verifiability it provides is not enough for most applications, including ours (Section 4). In Section 3.2.2 we describe in detail how Mastic overcomes this gap.

2.4 Shared ZK Proofs

Zero-knowledge (ZK) proof systems enable a prover to demonstrate to a verifier that a value has some specific property in zero-knowledge, i.e., without revealing anything about the value to the verifier beyond its validity. The specific variant we use, from [10], operates on secret-shared inputs, and verification is distributed amongst the share holders. This proof system has been incorporated into the candidate standard for Prio [5]; here, we devise a syntax that is suitable for both Prio and Mastic.

A *shared ZK proof* system defines three algorithms: Prove, Query, and Decide. In practice, a client splits its input into linear shares before providing them to Prove, which generates proofs for each aggregator. We require the existence of an algorithm Extract that extracts a single input from the shares; this eliminates ambiguity about the witness. The aggregators then run Query, exchange their verification strings, and call Decide to compute a final verdict.

Like VIDPFs, we define the security of shared ZK proofs in the random oracle model, so we equip the system with two sets Dom and Rng. Again, at the start of each security game involving shared ZK, we will sample a uniformly random function H from the set of all functions with domain Dom and range Rng. Then all three algorithms and the adversary are given oracle access to H.

- $\mathcal{Z}.\mathsf{Prove}^{\mathsf{H}}(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1) \rightarrow (\pi_0^{\mathsf{szk}}, \pi_1^{\mathsf{szk}}, \mathsf{nonce})$: Let x denote the private measurement, presumably in the language $\mathcal{L} \subseteq \mathbb{F}^m$ recognized by $\mathcal{Z}$. The *proof generation* algorithm takes linear shares $\mathsf{Extract}(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1) = x$ of x and outputs a *partial proof* for each aggregator along with a nonce. This algorithm is run by each client.
- $\mathcal{Z}.\mathsf{Query}^{\mathsf{H}}(vk, \mathsf{nonce}, \llbracket x \rrbracket_b, \pi_b^{\mathsf{szk}}) \rightarrow (\mathsf{st}_b, \sigma_b)$: The *query generation algorithm* takes in the *verification key* ($vk \in \{0,1\}^{\mathsf{vkl}}$) shared by the aggregators, the nonce, one of the secret shares ($\llbracket x \rrbracket_b$), and the corresponding partial proof (π_b^{szk}). It returns the aggregator's state (st_b) and a *partial verifier* (σ_b). This algorithm is run by each aggregator.
- $\mathcal{Z}.\mathsf{Decide}^{\mathsf{H}}(\sigma_0, \sigma_1, \mathsf{st}) \rightarrow \mathsf{Accept/Reject}$: Finally, the *decision algorithm* takes in the verifier shares and the state (st)

of an aggregator. It outputs Accept if x was recognized as a member of $\mathcal{L}$ and Reject otherwise. This algorithm is run by each aggregator.

We define the security properties required for $\mathcal{Z}$ in Appendix B. Briefly, we require $\mathcal{Z}$ to be: *Complete*: when the client and aggregators are honest, the aggregators accept the measurement; *Sound*: when the aggregators are honest, invalid measurements are detected with high probability; and *Zero-knowledge*: when the client and at least one aggregator is honest, execution of the proof system reveals nothing about the measurement beyond its validity.

The shared ZK system $\mathcal{Z}$ can be instantiated from a *fully linear proof* system [22] similar to the candidate standard for Prio [5]. See Fig. 9 of Appendix B for details.

3 Weighted Heavy-Hitters and Attribute-Based Metrics

This section describes the Mastic protocol, built from a VIDPF and shared ZK proof system. We focus the presentation on weighted heavy hitters and capture attribute-based metrics as a mode of operation for Mastic in Section 3.2.3. We begin by describing our threat and communication model.

3.1 Threat and Communication Model

There are many clients (thousands or millions), each of which holds a private *measurement* consisting of an *input* $\alpha \in \{0,1\}^n$ and its *weight* β. There are two aggregators who are responsible for gathering the private client measurements and aggregating them. One of these aggregators plays a special role: we call it the *leader* $\mathcal{S}_0$ and the other the *helper* $\mathcal{S}_1$. For the most part, the leader and helper perform the same computation except that the leader picks a shared *verification key* that is used for verifying the weights. No other cryptographic asset is required to execute the protocol (except those needed to establish secure channels). Fundamentally, we trust the leader and helper not to collude against client privacy, and we trust them to honestly enforce integrity of the results, which is a goal in their own interest.

Each client sends a message to each aggregator that encodes a secret share of its measurement. Thereafter, the aggregators interact with one another to verify each pair of shares is a valid measurement and to compute the (weighted) heavy hitters. The clients do not participate in the protocol beyond sending their initial messages.

We make the following assumptions about the adversary. First, we assume that all malicious parties are active: they may deviate arbitrarily from the protocol. Second, the set of clients and aggregators is fixed for the adversary's attack. Third, multiple instances of the protocol may be executed simultaneously. Fourth, corruptions are static: the adversary chooses the set of parties it controls before beginning its attack. Finally, the adversary may eavesdrop on any communication channel except for those between each honest client and each honest aggregator (in practice, a secure channel will be established between them). We consider two security goals:

Privacy in the presence of malicious clients and a malicious aggregator. Our primary goal is that the attacker learns nothing more than the aggregates of the honest clients' measurements over the

course of a single protocol execution.[5] Particularly, the individual measurements themselves are not revealed to the adversary. We are interested in protecting the measurements from a malicious aggregator, but will further allow the attacker to corrupt some fraction of the clients in order to target the remaining honest participants.

Robustness in the presence of malicious clients. Malicious clients may attempt to disrupt the protocol by providing malformed inputs. Our goal is to prevent an attacker from forcing the aggregators to compute anything other than aggregate measurements submitted by honest clients and valid *partial* measurements submitted by malicious clients. (A valid partial measurement is of the form (α_i, β_i) where β_i is valid, i.e., $\beta_i \in \mathcal{L}$, and α_i is a string of length $\leq n$ bits instead of n bits; see Section 5 for details.)

We formalize these goals in Section 5. Note that privacy is not achievable if *both* aggregators are corrupted by the adversary. Likewise, we achieve robustness only when both aggregators behave honestly. Robustness can be added in the presence of a corrupt aggregator at a significant cost: either by adding a third aggregator (where one out of three can be corrupted) as in PLASMA [35], or using verifiable computation in the two aggregator setting where the aggregators prove the correctness of their computation using a zero-knowledge proof [45]. General-purpose zero-knowledge proofs are prohibitively expensive for this application [45] since the ZK circuit (which each aggregator has to prove) would scale linearly with the number of clients.

3.2 The Mastic Protocol

Mastic is specified in Fig. 2 in terms of a VIDPF $\mathcal{V}$, a shared ZK proof system $\mathcal{Z}$, and a hash function H with domain Dom and range Rng. The input to the computation is the set $\{(\alpha_i, \beta_i)\}_{i \in [N]}$ of client input/weight pairs; the output is the subset of α_i's that have the highest total weight. Its execution is associated with a *threshold* $\mathsf{T} \geq 0$ and a function order $: \mathbb{F}^m \mapsto \mathbb{R}$ defining a total ordering of sums of weights.

3.2.1 Client Report Generation. Each client C_i holds a private measurement that comprises two parts: a bit-string α_i of length n and an associated weight β_i in some language $\mathcal{L} \subseteq \mathbb{F}^m$ of valid weights. The language $\mathcal{L}$ depends on the application. For example, we might want β to be a bit (i.e., $\mathcal{L} = \{0, 1\}$), an integer in a public range ($\mathcal{L} = [R]$ for some R), or a one-hot vector ($\mathcal{L} = \{[1, 0, \ldots, 0], [0, 1, 0, \ldots, 0], \ldots, [0, \ldots, 0, 1]\}$). The language must be agreed upon by the clients and aggregators before report generation and processing begins.

The client report generation phase is shown in Fig. 2 under "Client Computation". During this phase, each client first encodes its α_i, β_i using the VIDPF key generation algorithm that produces a pair of keys $\text{key}_{(i,0)}, \text{key}_{(i,1)}$ and a public share pub_i. Next, it generates a shared ZK proof that asserts to the aggregators that its private weight β_i is indeed a member of the language $\mathcal{L}$. To do this, the client passes the exact secret shares of β_i to the proof generation algorithm that the aggregators will receive when evaluating the

VIDPF at the root (i.e., the empty bit-string ϵ). The client computes these shares by invoking $\mathcal{V}$.EvalRoot twice, each time with the inputs of each aggregator. The $\mathcal{Z}$.Prove algorithm returns the partial proofs $\pi^{\text{szk}}_{(i,0)}, \pi^{\text{szk}}_{(i,1)}$ as well as a nonce to the client. Finally, the client sends the nonce nonce_i, the public share pub_i, the VIDPF key $\text{key}_{(i,b)}$, and the partial proof $\pi^{\text{szk}}_{(i,b)}$ to each aggregator $\mathcal{S}_b$. After each client sends its messages, it goes offline.

3.2.2 Aggregator Computation. Next, we delve into the main phase of our protocol in Fig. 2, the "Aggregator Computation".

Initialization. The aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ must first agree on a verification key $vk \in \{0, 1\}^{\text{vkl}}$. For the sake of simplicity, we assume the leader $\mathcal{S}_0$ chooses this value unilaterally and sends it to the helper $\mathcal{S}_1$. This is sufficient for security, as long as the aggregators commit to this value before the protocol begins (see [22, Section 3.2]). [6] They then initialize a set of sets $\text{HH}^{\leq n} \coloneqq \{\text{HH}^0, \text{HH}^1, \ldots \text{HH}^n\}$ as $\{\epsilon, \emptyset, \ldots, \emptyset\}$, where n is the size of the clients' bit-strings. As the aggregators evaluate all n levels, they will start populating $\text{HH}^{\leq n}$ with the weighted heavy hitter bit-strings. Finally, the aggregators initialize a list Reports that contains the *candidate reports* from the clients. A report is removed from this set if a validity check fails.

Verifying Client Inputs. The aggregators can now start generating secret shares of the client reports by evaluating the VIDPF keys of each client. It is crucial for the aggregators to verify that the client's input is well-formed in a privacy-preserving manner. There are three things to check: *Weight Check:* $[\![\beta_i]\!]_0 + [\![\beta_i]\!]_1 \in \mathcal{L}$, where $[\![\beta_i]\!]_0, [\![\beta_i]\!]_1$ are the aggregators' shares of β_i; *One-Hotness Check:* evaluating the keys of C_i on different prefixes *at the same level* should yield secret shares of a one-hot vector (i.e., only one pair of secret shares should correspond to β_i and the rest to zero); and *Path Check:* Each node along the α_i-path of the prefix tree has weight β_i. We emphasize that all the aforementioned checks need to be performed by the aggregators while keeping the measurement (α_i, β_i) private. If any of these checks are omitted, then a malicious client will be able to trick the aggregators and affect the robustness of the protocol (i.e., data poisoning attack).

We split these checks into two different categories; the checks we perform at the root level of the VIDPF tree and the checks we perform at any other level (including the leaves).

Root-Level Checks. The aggregators need to verify that the private β_i is indeed a valid weight (i.e., $\beta_i \in \mathcal{L}$, where $\mathcal{L}$ depends on the type of statistic). This is a two-fold process: First, the aggregators evaluate each client's VIDPF keys at the root level of the tree (i.e., the empty bit string ϵ) to get secret shares of β_i (i.e., $[\![\beta_i]\!]_0$ and $[\![\beta_i]\!]_1$). Each $\mathcal{S}_b$ acquires $[\![\beta_i]\!]_b$ by invoking $\mathcal{V}$.EvalRoot using the client's key $\text{key}_{(i,b)}$ and the public share information pub_i. Next, each aggregator $\mathcal{S}_b$ performs a shared ZK query using their retrieved $[\![\beta_i]\!]_b$ share, the verification key vk shared by the aggregators, the nonce nonce_i associated with the client's report, and the partial proof $\pi^{\text{szk}}_{(i,b)}$ generated by the client. Each $\mathcal{S}_b$ receives a state $\text{st}_{i,b}$ and partial verifier $\sigma_{i,b}$. The aggregators exchange their partial verifiers and run the shared ZK decision algorithm to verify if the client's weight

---

Π_{Mastic} for weighted heavy hitters

CLIENT COMPUTATION:

Input: Each client C_i for $i \in [N]$ holds measurement $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$ composed of an input α_i and its weight β_i.

1. C_i runs $(\text{pub}_i, \text{key}_{(i,0)}, \text{key}_{(i,1)}) := \mathcal{V}.\text{Gen}(\alpha_i, \beta_i)$.

2. C_i runs $(\pi^{\text{szk}}_{(i,0)}, \pi^{\text{szk}}_{(i,1)}, \text{nonce}_i) := \mathcal{Z}.\text{Prove}(\llbracket\beta_i\rrbracket_0, \llbracket\beta_i\rrbracket_1)$ where $\llbracket\beta_i\rrbracket_b := \mathcal{V}.\text{EvalRoot}(\text{key}_{(i,b)}, \text{pub}_i)$ for $b \in \{0,1\}$.

3. C_i sends report share $(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi^{\text{szk}}_{(i,b)})$ to $\mathcal{S}_b$ for each $b \in \{0,1\}$.

AGGREGATOR COMPUTATION:

Input: The aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ start with a verification key $vk \in \{0,1\}^{\text{vkl}}$ established out-of-band. Each sets $\text{HH}^{\leq n} := \{\text{HH}^0, \text{HH}^1, \dots \text{HH}^n\}$ as $\{\epsilon, \emptyset, \dots, \emptyset\}$, the initial set of *candidate prefixes* for each level and sets $\text{Reports} := [N]$, the initial set of *candidate reports*. Finally each initializes $(\llbracket y_i^\epsilon\rrbracket_b, \text{st}^\rho_{(i,b)}, \pi^\epsilon_{(i,b)}) = (\bot, \epsilon, \epsilon)$ for each $i \in [N]$.

1. **For each** client $i \in \text{Reports}$:   ▷ Weight check using $\mathcal{Z}$ at the root.

   a. Remove i from Reports if $(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi^{\text{szk}}_{(i,b)})$ does not follow the correct formatting. ▷ Input-Formatting check.

   b. $\mathcal{S}_b$ runs $(\text{st}_b, \sigma_b) := \mathcal{Z}.\text{Query}(vk, \text{nonce}_i, \llbracket\beta_i\rrbracket_b, \pi^{\text{szk}}_{(i,b)})$, where $\llbracket\beta_i\rrbracket_b := \mathcal{V}.\text{EvalRoot}(\text{key}_{(i,b)}, \text{pub}_i)$.

   c. $\mathcal{S}_b$ sends σ_b to $\mathcal{S}_{1-b}$. If $\mathcal{Z}.\text{Decide}(\sigma_0, \sigma_1, \text{st}_b) \neq \text{Accept}$, then $\mathcal{S}_b$ removes i from Reports.

2. **For each** level $k \in [0, \dots, n-1]$ and prefix $p \in \text{HH}^k$:   ▷ Path & One-hot Verifiability checks.

   a. **For each** candidate report $i \in \text{Reports}$:

      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p\rrbracket_b, \text{st}^\rho_{(i,b)}, \pi^\rho_{(i,b)})$ from memory corresponding to prefix p and client C_i.

      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma\rrbracket_b, \text{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)}) := \mathcal{V}.\text{Eval}(\text{key}_{(i,b)}, \text{pub}_i, \gamma, \text{st}^\rho_{(i,b)}, \pi^\rho_{(i,b)})$ for each prefix $\gamma \in \{p \parallel 0, p \parallel 1\}$ and stores the results in memory.

      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \parallel 0$ and $p \parallel 1$. To do so, each $\mathcal{S}_b$ computes $h^p_{(i,b)} := (-1)^b \cdot (\llbracket y_i^p\rrbracket_b - \llbracket y_i^{p\parallel 0}\rrbracket_b - \llbracket y_i^{p\parallel 1}\rrbracket_b)$.   ▷ Observe that $h^p_{(i,0)} = h^p_{(i,1)}$.

      iv. $\mathcal{S}_b$ accumulates its local state as $R^k_{(i,b)} := \text{H}(\|_{p \in \text{HH}^k}(p, h^p_{(i,b)}, \pi^{p\parallel 0}_{(i,b)}, \pi^{p\parallel 1}_{(i,b)}))$.   ▷ This is for all heavy-hitter prefixes.

      v. $\mathcal{S}_b$ sends $R^k_{(i,b)}$ to $\mathcal{S}_{1-b}$. If $R^k_{(i,0)} \neq R^k_{(i,1)}$, then $\mathcal{S}_b$ removes i from Reports.   ▷ One hash for each client.

   b. **For each** k-bit heavy-hitting prefix $p \in \text{HH}^k$ the aggregators prune on $\gamma \in \{p \parallel 0, p \parallel 1\}$ as: ▷ Aggregation & Pruning.

      i. Each $\mathcal{S}_b$ accumulates $\llbracket\text{weight}^\gamma\rrbracket_b := \llbracket\text{weight}^\gamma\rrbracket_b + \llbracket y_i^\gamma\rrbracket_b$.   ▷ Each $\llbracket y_i^\gamma\rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.

      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\text{weight}^\gamma := \llbracket\text{weight}^\gamma\rrbracket_0 + \llbracket\text{weight}^\gamma\rrbracket_1$. If $\text{order}(\text{weight}^\gamma) < \text{T}$, then prune γ from the candidate prefix set. Otherwise, accumulate $\text{HH}^{k+1} := \text{HH}^{k+1} \cup \{\gamma\}$.   ▷ order$(\cdot)$ is decided by the aggregators.

3. Finally, the servers output HH^n as the set of weighted T-heavy-hitters.

Figure 2: Protocol Π_{Mastic} for (T, order)-weighted-heavy-hitters built from VIDPF $\mathcal{V}$, shared ZK $\mathcal{Z}$, and hash function H.

is valid. Finally, the aggregators remove from the candidate reports Reports each client C_i whose proof was not verified successfully.

This check is described in "Step 1." in "Aggregator Computation" in Fig. 2 and is performed for every client. Essentially, the shared ZK proof allows the aggregators to check that they have valid shares of a weight $\beta_i \in \mathcal{L}$ without reconstructing it. However, so far the aggregators have only verified that the client has submitted a report with a valid weight for the root of the VIDPF tree (i.e., the empty bit string). The next step is to verify that this weight is correctly propagated down the tree and compute the weighted heavy hitters.

Intermediate Levels and Leaves Checks. The protocol continues iteratively by processing one level at a time, starting from the two children of the root node (ϵ). The goal of the aggregators is to evaluate all the client reports on both children of the root node (namely $\epsilon \parallel 0$ and $\epsilon \parallel 1$), verify for each client that they have valid shares of these evaluations, aggregate them all together, and finally, only keep the bit strings whose aggregate weight is above the threshold T. For instance, if 0, 1 are both heavy hitters at level one, then the aggregators will evaluate the children of both (i.e., 00, 01, 10, and 11) at level two. For each evaluation of C_i's report at each of these prefixes, the aggregators acquire secret shares of weight β_i if p is a prefix of α_i and secret shares of zero (0^m) otherwise. So, if C_i's measurement is $\alpha_i = 111 \dots$, then all the evaluations of C_i's VIDPF

keys at level two yield shares of zero, except for the evaluation on $p = 11$ which returns shares of β_i. At the next level, the aggregators need to verify the evaluated path, i.e., the evaluation on $p = 111$ returns shares of the same weight β_i as in the previous level, while all other evaluations at level three return secret shares of zero.

The way to do this is surprisingly *simple* and *efficient*: check that the output for p is equal to the sum of the outputs of its two children (i.e., the prefixes $p \parallel 0$ and $p \parallel 1$):

$$\mathcal{S}_0 \text{ computes } h^p_{(i,0)} := \llbracket y_i^p\rrbracket_0 - \llbracket y_i^{p\parallel 0}\rrbracket_0 - \llbracket y_i^{p\parallel 1}\rrbracket_0; \text{ and}$$

$$\mathcal{S}_1 \text{ computes } h^p_{(i,1)} := -\llbracket y_i^p\rrbracket_1 + \llbracket y_i^{p\parallel 0}\rrbracket_1 + \llbracket y_i^{p\parallel 1}\rrbracket_1.$$

Observe that checking that $h^p_{(i,0)}$ and $h^p_{(i,1)}$ are equal can be done by hashing both and comparing the hashes. This equality guarantees us that the evaluation of the parent node is the sum of the evaluations of the children since $h^p_{(i,0)} = h^p_{(i,1)}$ then:

$$\llbracket y_i^p\rrbracket_0 - \llbracket y_i^{p\parallel 0}\rrbracket_0 - \llbracket y_i^{p\parallel 1}\rrbracket_0 = -\llbracket y_i^p\rrbracket_1 + \llbracket y_i^{p\parallel 0}\rrbracket_1 + \llbracket y_i^{p\parallel 1}\rrbracket_1$$

$$\Leftrightarrow \llbracket y_i^p\rrbracket_0 + \llbracket y_i^p\rrbracket_1 = \llbracket y_i^{p\parallel 0}\rrbracket_0 + \llbracket y_i^{p\parallel 1}\rrbracket_0 + \llbracket y_i^{p\parallel 0}\rrbracket_1 + \llbracket y_i^{p\parallel 1}\rrbracket_1$$

$$\Leftrightarrow y_i^p = y_i^{p\parallel 0} + y_i^{p\parallel 1}.$$

However, this path check alone is not sufficient as the weight of the parent can be split between the children. This check will still pass, but this is not a valid report since it is not one-hot.

One-hotness is assured by the verifiability property of VIDPF (Section 2.3). For each prefix evaluation, the two aggregators also receive a proof. By combining multiple proofs at the same level, the aggregators can verify that the evaluations for a specific level are one-hot. This one-hot check, in conjunction with the path and weight check, is sufficient to ensure the validity of each client's report at a given level.

The aforementioned checks are described in "Step 2.a." in "Aggregator Computation" in Fig. 2 and are performed for every client. In more detail, the aggregators evaluate the VIDPF in the two children $\gamma \in \{p \| 0, p \| 1\}$ of the current path p, each receives a secret share $[\![y_i^\gamma]\!]_b$ and VIDPF proof $\pi_{(i,b)}^\gamma$. Then, they use the secret shares for the path verifiability check and generate $h_{(i,b)}^p$ for each prefix p. As there are usually multiple candidate prefixes at each, we batch all the checks for all the prefixes per client: the aggregators hash all $h_{(i,b)}^p, \pi_{(i,b)}^{p\|0}, \pi_{(i,b)}^{p\|1}$ for all $p \in \mathsf{HH}^k$ into $R_{(i,b)}^k$. Finally, the aggregators exchange $R_{(i,b)}^k$ for each client C_i and check that $R_{(i,0)}^k = R_{(i,1)}^k$. If they are not equal, then the aggregators remove i from the set of candidate reports Reports as either the path or the one-hot check has failed. The aggregators exchange as many hashes as the total clients, but this can be improved as described in Section 3.2.4.

Aggregation and Pruning. After the "Verifying Client Inputs" step is done, the two aggregators have removed all malformed reports for the root level as well as for the current level (e.g., level k). The next step is to aggregate all the client reports together and compute the heavy hitter prefixes HH^k of length k. This step is shown in "Step 2.b." in "Aggregator Computation" in Fig. 2. Each $\mathcal{S}_b$ accumulates all secret shares $[\![y_i^\gamma]\!]_b$ for prefixes $\gamma \in \{p \| 0, p \| 1\}$ into $[\![\mathsf{weight}^\gamma]\!]_b$. Then, both aggregators recover weight^γ and decide to keep (or prune) the prefix γ based on whether or not the output of a function $\mathsf{order}(\cdot)$ is greater than a threshold T. Note that each weight β_i is a vector of field elements ($\beta_i \in \mathbb{F}^m$ as defined in Section 2.3). Finally, the aggregators prune all prefixes with order less than T and continue to the next level.

3.2.3 Modes of Operation.
Mastic is designed to support a variety of secure aggregation tasks. These can be categorized into three modes of operation.

Weighted Heavy Hitters. Mastic is designed primarily to solve the weighted heavy hitters problem. In this problem, each measurement is a pair $(\alpha, \beta) \in \{0,1\}^n \times \mathcal{L}$ and the output is the subset of inputs for which $\mathsf{order}(\mathsf{weight}) \geq \mathsf{T}$, where weight is the sum of the weights and $\mathcal{L} \subseteq \mathbb{F}^m$, $\mathsf{order} : \mathbb{F}^m \mapsto \mathbb{R}$, and $\mathsf{T} \in \mathbb{R}$ are determined by the application. We provide an example of such an application in Section 4.1.

Plain Heavy Hitters. The plain heavy hitters problem is a special case of weighted heavy hitters where each $m = 1$, $\mathcal{L} = \{[0], [1]\}$, and $\mathsf{order}(\cdot)$ is the identity function. In this case, Mastic reduces to a two-aggregator variant of PLASMA.

Attribute-Based Metrics. Mastic admits an enhanced variant of general-purpose metrics (à la Prio [5]) in which aggregates are broken down by client attributes. Just as with weighted heavy hitters, each measurement consists of a pair $(\alpha, \beta) \in \{0,1\}^n \times \mathcal{L}$, where we call α the client's *attribute* and β the client's *value*. The aggregators evaluate the reports on a set of attributes of interest

$[x_1, \ldots, x_A]$. The result is the aggregate (i.e., total of the values) of reports that share the same attribute. We give an example in Section 4.2.

Mastic in this mode of operation is essentially a subset of the tree traversal algorithm in Fig. 2. The main difference is that, instead of using the weights to decide which paths to traverse, the aggregators decide in advance which leaves they want to evaluate (i.e., the attributes) and traverse the path to each of these leaves. The shared ZK proof is checked at the root. The VIDPF proof, state, and path checks are all computed as usual, except they only exchange the $R_{(i,b)}^k$-values at the last level of the tree. (The shared ZK partial proofs are exchanged in the same flow.) Note that, in order to compute the path checks, it is also necessary to evaluate the sibling of each node traversed along a path to a leaf.

3.2.4 Optimizations.
A few techniques can be used to reduce the concrete communication cost of Mastic.

Batched Path and One-Hot Checks. In the "Client Input Verification" of Section 3.2.2, during the intermediate levels (say k) and leaves verification, the aggregators need to exchange (up to)[7] N hashes (i.e., $R_{(i,b)}^k$ for $i \in [\mathsf{N}]$) to verify the inputs of N clients. We adopt a batching optimization based on Merkle trees proposed by PLASMA [35] to reduce this number and allow the aggregators to verify N clients in a batch. Each $\mathcal{S}_b$ creates a Merkle tree from the N hashes and sends the Merkle tree root to the other aggregator. If they match, then we know that the one-hot and path checks passed for all clients. If the roots are different, then the aggregators repeat the same process for the next level until they traverse down the tree and remove the malicious clients.

Batched Weight Checks. In concurrent work (IEEE S&P 2024 [40]), it was shown how to transform the ZK proof system on secret-shared data of Boneh et al. [10]. The same transformation could be applied to shared ZK during the first round of aggregation. We remark that, as this work is concurrent with our own, we did not evaluate this optimization in our implementation (Section 6).

Minimizing Communication for Weight Checks. In plain heavy hitters mode, the shared ZK evaluation can be foregone completely in favor of the simpler, $\beta = 1$ check of PLASMA. This decreases the size of the reports sent from clients to aggregators (shared ZK partial proofs are omitted) as well as the size of the messages sent between the aggregators (shared ZK partial verifiers are omitted from the root checks).

The size of the shared ZK partial proofs used in our construction of $\mathcal{Z}$ is $\mathcal{O}(m)$ in general. With standard PRG-based secret sharing techniques, we can ensure the concrete communication cost of the $\mathcal{Z}$ partial proofs is small. (The helper's partial proof can be represented by a PRG seed; see [5, Section 7] for details.)

4 Applications

This section describes two key applications that motivate Mastic. Both were described in discussions in and around the working group at the IETF developing standards for secure aggregation [29]. In fact, Mastic is designed to overcome limitations of existing protocols already being considered [5]: Poplar [11] and Prio [19].

---

[7]Note that the number of hashes here may be less than N as some reports may have been removed at a previous level of the tree.

Poplar (and PLASMA [35]) are designed specifically to solve the plain (i.e., non-weighted) heavy hitters problem. There is a plethora of applications for which this limited functionality falls short, one of which we describe in detail below (Section 4.1). Our primary goal for Mastic is to provide a replacement for Poplar that is more flexible and therefore useful in a wider variety of applications.

Another goal is to extend Prio with a feature often requested at the IETF [26, 41]: the ability to group Prio-style metrics by client attributes. For example, when aggregating a histogram, we may want to split the results into a separate histogram for clients grouped by some property, like their geographical location. We describe a motivating application in Section 4.2.

Prio can be extended to support such functionality: instead of aggregating a length-m histogram, we would aggregate a length-$(m \cdot A)$ histogram comprised of A length m histograms, one for each of the A attributes of interest. However, this would result in $\mathcal{O}(A)$ overhead in communication, which is impractical in most situations. Mastic, on the other hand, provides the same functionality, but with only $\mathcal{O}(n)$ overhead, where n is the length in bits of each attribute.

4.1 Network Error Logging

Network Error Logging (NEL) [44] is a mechanism used by web browsers to report errors that occur while attempting to establish a connection to a server. Some of these errors are visible to the server, but not all: failures in DNS, TCP, TLS, and HTTP can occur without the server having any visibility into the issue. A small number of connection errors is expected, even under normal operating conditions; but a sudden increase in errors may be an indication of an outage, or a configuration issue impacting millions of users. Without a reporting mechanism like NEL, these events would manifest in the server's telemetry merely as a drop in overall traffic.

NEL is particularly important for content delivery networks, such as Akamai, AWS, or Cloudflare, that handle HTTP traffic for a large number of websites (typically millions). A content delivery network acts as a reverse proxy between clients and origin servers that provides a layer of caching and security services.

Reports are comprised of the URL the client attempted to navigate to (e.g., `https://example.com`), the type of error that occurred, and metadata related to the attempt, such as the time that elapsed between when the connection attempt began and when the error was observed (e.g., [44, Section 7]). Clients may also report *successful* connection attempts to give the server a sense of the *error rate*. The exact client behavior is determined by the reporting policy specified by the server (see [44, Section 5.1]).

NEL data is privacy-sensitive for two reasons. First, it exposes information that the server would not otherwise have access to, which can be abused to probe the client's network configuration as described in [44, Section 9]. Second, for operational reasons, the reporting endpoint may be organizationally separated from the server (i.e., run on different cloud infrastructures), leading to an increased risk of the client's browsing history being exposed (e.g., in a data breach).

Initial experiments with Poplar for this use case were reported on the IETF working group mailing list [1]. They found that NEL was too latency-sensitive for Poplar to be useful, as there is a need for real-time detection of errors and impacted origins. In this section, we show how to apply Mastic to help overcome this challenge.

Private NEL with Mastic. We consider here a simplified version of NEL (similar to [1]) where each client reports a tuple $(\mathsf{dom}, \mathsf{err})$ consisting of a domain name dom (e.g., `example.com`) and a value err that represents an error (for instance, `dns.unreachable`) or an indication that no error occurred (e.g., `ok`). Notably, this can be easily extended in Mastic to represent more elaborate metrics. e.g., where each weight includes the time it took each browser to report the error (and the aggregate is the average error reporting time), user agent (browser type and version), etc. However, our main goal is to understand 1) the distribution of errors and 2) which domains are impacted.

We expect there to be a large number of distinct domain names (millions in the case of content delivery networks) and only a small number of error variants (the NEL spec [44] defines 30 variants). The following Mastic parameters are suitable for this application.

Inputs: Each α encodes the domain dom truncated to $n = 256$ bits, which is sufficient to represent most of the domains on the internet [11, 35]. Shorter domains are padded with 0s.

Weights: Each weight β represents the error variant dom. To compute the distribution of errors, we encode each error variant as a distinct bucket of a histogram so that $[1, 0, 0, \ldots]$ represents `ok`, $[0, 1, 0, \ldots]$ represents `dns.unreachable`, $[0, 0, 1, \ldots]$ represents `dns.name_not_resolved`, and so on. There are 30 such variants (see [44, Section 6]), so the language $\mathcal{L}$ of weights is exactly the set of length-30 vectors over $\mathbb{F}$ containing all 0s except for a single 1.

Ordering: Our order($\cdot$) function computes the ratio of reports with err $\neq$ ok to reports with err = ok. The latter is simply the first bucket of the aggregated histogram; the former is the sum of the remaining 29 buckets. Note that our ordering of aggregated weights considers the error rate rather than the raw error count. This ensures that the signal for less popular domains is not swamped by the noise generated by popular sites (network issues may impact some domains but not others). Another benefit is that, under normal operating conditions, there will be a small number of heavy-hitting domains, which means Mastic will run very efficiently. During an incident, there will be more heavy hitters, which means it will take longer to compute the set of impacted domains. However, *we get the errors immediately* at the root of the prefix tree, which is the most important information needed to begin remediation. As more levels are evaluated, we get more detailed errors.

4.2 Attribute-Based Browser Telemetry

Web browsers, like Chrome, Firefox, or Safari, collect telemetry generated by users as they surf the web to gain insights into trends that guide product decisions. In many cases, Prio can be used to privately aggregate this telemetry. However, this comes at the cost of flexibility.

For example, Mozilla is using Prio [28] to collect page load metrics from Firefox for a list of known popular sites (e.g., `google.com`). The purpose of these metrics is to detect if changes to these sites cause regressions that might be correlated with an increased average load time or error rate. A subtle, but important requirement for this system is the ability to break down the metrics by client

attributes. The most crucial attributes are 1) the software *version*, and 2) the information about the client's *location*.

Meeting this requirement by increasing the size of the histogram leads to intolerable communication overhead. An alternative is to have each client upload this information in the clear alongside its Prio report so that the reports can be grouped by version and location. The downside of this approach is that it significantly reduces the anonymity set of each user since they are only mixed with their attribute group rather than the entire population.

Private Browser Telemetry with Mastic. Mastic provides a simple solution to this problem. For the sake of presentation, we consider a simplified version of Mozilla's use case (the same approach can be applied to any aggregation task for which Prio is suitable). Each client reports a tuple $(\mathsf{ver}, \mathsf{loc}, \mathsf{site}, \mathsf{time})$ where: ver is a string representing the client's software version (e.g., `Firefox/122.0`); loc is a string encoding its country code (e.g., `GR`, `US`, `IN`, etc.); site is one of a fixed set of sites (e.g., `google.com`, `wikipedia.org`, etc.); and time is the load time of the site in seconds. The version and location are included in the Mastic input; the site and load time are encoded by the corresponding weight. Notably, this is just one example of what Mastic can do; the same idea can be applied to other types of metrics.

Compared to the private NEL application in Section 4.1, the number of possible inputs here is relatively small: there are fewer than 200 country codes and a handful of browser versions in wide use at any given time. This means the aggregators can enumerate a set of inputs of interest and evaluate them immediately. Consider the following parameters for Mastic, in its attribute-based metrics mode of operation (Section 3.2.3):

Attributes: Two-letter country codes can easily be encoded in 2 bytes. Likewise, the number of distinct browser versions is easily less than 2^{16}, so 2 bytes are sufficient. Therefore, each α can be encoded with just $n = 32$ bits.

Values: Similar to private NEL, each weight β is a 0-vector except for a single 1 representing a bucket in a histogram. We represent $(\mathsf{site}, \mathsf{time})$ as a histogram bucket as follows. First, we quantize time (in seconds) into one of four buckets: $[0, 0.1)$, $[0.1, 1)$, $[1, 5)$, and $[5, \infty)$. Let $t \in [4]$ denote the time bucket for time. Next, suppose we wish to track metrics for 25 sites. Let $s \in [25]$ denote the index of site in this list. Then the index of 1 in β is simply $t \cdot s$ such that $|\beta| = m = 4 \cdot 25 = 100$.

5 Security Analysis of Mastic

In this section, we present our security analysis of Π_{Mastic} (Fig. 2). Following [19, 22], we consider privacy and robustness separately.

5.1 Privacy: Malicious Clients and Aggregator

To define privacy, we first consider what information is protected and what is leaked. The Mastic protocol is designed to reveal the sum of the weights for every valid report whose point α is prefixed by the query p of interest to the aggregators. We require that even if one of the two aggregators is malicious, neither aggregator learns more than this sum, even if it knows something else about the individual client measurements.

We capture this property in a simulation-based model, presented in Figs. 10 and 11 found in Appendix C. Essentially, we ask for

the existence of a stateful algorithm Sim that can interact with a malicious aggregator exactly as an honest aggregator and clients would. This simulator should be indistinguishable from the real protocol operations, even if the adversary knows all honest clients' measurements, and the simulator knows only the aggregate results. When one of the aggregators is fully malicious, Mastic can be used in a variety of ways that do not conform to the weighted heavy-hitters application, so we consider privacy over all modes of operation in Section 3.2.3. This means that a malicious aggregator can ask for aggregate results across any set of client reports it desires without violating our privacy notion. Consequently, we do not protect against inference or Sybil attacks and leave the according defenses up to higher-level systems.

We define two games, $\mathcal{G}_{\mathsf{Mastic}}^{\mathsf{priv-real}}$ (c.f. Fig 10), and $\mathcal{G}_{\mathsf{Sim}}^{\mathsf{priv-ideal}}$ (c.f. Fig 11). Each of these games initially requests a verification key vk and the corrupt aggregator index b from the adversary $\mathcal{A}$, then presents an interface of six oracles capturing all the interactions of an honest aggregator: with honest clients, corrupt clients, and the malicious aggregator. The adversary may interact with these oracles at will; when it halts, it must output a bit denoting whether it believes the interface is real or simulated.

Informally, the game $\mathcal{G}_{\mathsf{Mastic}}^{\mathsf{priv-real}}$ presents the adversary with a view of the real Mastic protocol, and $\mathcal{G}_{\mathsf{Sim}}^{\mathsf{priv-ideal}}$ presents a simulated view that doesn't depend on the value of individual client measurements. Our model is actually slightly stronger, in that we give the adversary a little *more* information than it would learn from a real Mastic interaction. This is because in Mastic, aggregators process many prefixes simultaneously for each report, and they verify one-hot and path verifiability proofs in batches that are hashed together. However, since Mastic doesn't place any restrictions on the size of these batches or whether they overlap, we let the model capture only the worst-case scenario and have aggregators process prefixes individually and return all proofs without hashing. It should be clear that any information that is leaked by the batched proofs will also be leaked by the inputs to the hash function, so this strictly strengthens the security definition.

We define the advantage $\mathbf{Adv}_{\mathsf{Mastic},\mathsf{Sim}}^{\mathsf{priv}}(\mathcal{A})$ of $\mathcal{A}$ against the privacy of Mastic with respect to simulator Sim as:

$$\left| \Pr[\mathcal{G}_{\mathsf{Mastic}}^{\mathsf{priv-real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\mathcal{G}_{\mathsf{Sim}}^{\mathsf{priv-ideal}}(\mathcal{A}) \Rightarrow 1] \right|.$$

Informally, we say that Mastic is private if there is a PPT simulator for which all PPT adversaries get negligible advantage in the security parameter.

Let $\mathbf{Adv}_{\mathcal{Z},\mathsf{Sim}_{\mathsf{SZK}}}^{\mathsf{priv}}(\cdot)$ denote the advantage of an adversary in attacking the privacy of $\mathcal{Z}$ with respect to simulator $\mathsf{Sim}_{\mathsf{SZK}}$ (we define this function precisely in Appendix B). Likewise, let $\mathbf{Adv}_{\mathcal{V}}^{\mathsf{priv}}(\cdot)$ denote the advantage of an adversary in attacking the privacy of $\mathcal{V}$ (defined in Appendix A). We claim the following theorem and prove it via a full pseudocode specification of the simulator and a series of game hops in Appendix C.

THEOREM 1. *For any simulator* $\mathsf{Sim}_{\mathsf{SZK}}$, *there exists a simulator* Sim *such that for any adversary* $\mathcal{A}$, *there exist* $\mathcal{Z}$-*attacker* $\mathcal{B}$ *and* $\mathcal{V}$-*attacker* $\mathcal{D}$ *such that*

$$\mathbf{Adv}_{\mathsf{Mastic},\mathsf{Sim}}^{\mathsf{priv}}(\mathcal{A}) \le \mathbf{Adv}_{\mathcal{Z},\mathsf{Sim}_{\mathsf{SZK}}}^{\mathsf{priv}}(\mathcal{B}) + q \cdot \mathbf{Adv}_{\mathcal{V}}^{\mathsf{priv}}(\mathcal{D}),$$

where q is the number of queries made to the "Honest Client Computation" oracle, and the runtime of $\mathcal{D}$ is about that of an honest aggregator in the Mastic protocol, and the runtime of $\mathcal{B}$ is about that of an honest aggregator plus the time to run Sim_{SZK} *once per interaction with the honest aggregator.*

Proof sketch. For each honest client, our simulator generates one false input share to pass on to the adversary. These false input shares consist of a VIDPF key for a random measurement and a seemingly valid proof generated by the zero-knowledge simulator for the SZK proof system without knowledge of the corresponding input. It deletes the honest aggregator's key immediately, so VIDPF privacy guarantees that the adversary cannot detect that the single remaining key corresponds to an incorrect measurement.

When the adversary asks the simulator in its role as honest aggregator to produce an SZK verifier share or decision for an honest client, the simulated aggregator has no input share to query. Instead, we refer the request to the SZK simulator's Query and Decide interfaces, which provide the required values. The final component that the simulator must produce is the honest aggregator's share of the path verifier checks. For malicious reports, the simulator simply runs the honest query algorithm on the share it is provided. For honest reports, recall that the simulator generates the adversary's input share, and can thus perform the path verifier checks on behalf of the adversary. Since all honest reports are well-formed, the malicious and honest path verifier check shares will sum to zero, so the simulator negates the adversary's checks for a perfect simulation. The reduction, therefore, has two components: one reduction to VIDPF privacy for each false input share, and a reduction to SZK privacy capturing the substitution of the SZK simulator for a real prover. The simulation of the path-verifier checks is always correct and indistinguishable, so it does not factor into the theorem bound.

5.2 Robustness: Malicious Clients

Next, we focus on the robustness guarantees provided by Mastic against malicious clients. To argue robustness we assume the aggregators follow the Mastic protocol steps correctly.

We capture this property in a simulation-based model [14]. The adversary $\mathcal{A}$ initially corrupts a set Reports' of clients. In the real-world game $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ (Fig. 18 in Appendix D), the parties run the Mastic protocol using their input measurements. Both the honest and corrupt clients provide their report shares to the aggregators, who compute the output (set of heavy-hitter strings and their children, and also the weights of the heavy-hitting strings and their children) and return it to $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$. The game forwards this to $\mathcal{A}$.

We also define a corresponding ideal-world game, $\mathcal{G}_{\text{Sim}, \mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$ (Fig. 18). In this game, the corrupt clients (controlled by $\mathcal{A}$) participate in the same way as in the real world, but with a pair of simulated aggregators. This simulation is performed by a PPT algorithm called Sim (Fig. 19). The simulator also has access to an ideal functionality $\mathcal{F}_{\text{wHH}}$ (Fig. 16) for weighted heavy-hitters. The simulator obtains the report shares of the corrupt clients, extracts the input measurements (α_i', β_i') from those shares, and invokes $\mathcal{F}_{\text{wHH}}$ (on behalf of the corrupt clients) with the extracted measurements. Upon obtaining the client measurements from both honest and corrupt clients, $\mathcal{F}_{\text{wHH}}$ checks the measurements and then computes

the output. For the consistency check, $\mathcal{F}_{\text{wHH}}$ discards the measurements if they are not correctly formatted, i.e., $(\alpha_i, \beta_i) = (\bot, \bot)$, or if the weight is not valid, i.e., $\beta_i \notin \mathcal{L}$. The functionality also allows corrupt clients to submit prefix strings, i.e., $\alpha_i \in \{0, 1\}^{\leq n}$. Once the measurements are validated, the functionality aggregates the validated measurements by computing the weights and heavy-hitting set similar to the aggregation phase in Mastic. $\mathcal{F}_{\text{wHH}}$ computes the output and this is returned to $\mathcal{G}_{\text{Sim}, \mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$. Finally, the game forwards this to $\mathcal{A}$.

Define the advantage $\mathbf{Adv}_{\text{Mastic,Sim}}^{\text{rob}}(\mathcal{A})$ of $\mathcal{A}$ in breaking the robustness of Π_{Mastic} with respect to simulator Sim as

$$\left| \Pr[\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\mathcal{G}_{\text{Sim}, \mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}(\mathcal{A}) \Rightarrow 1] \right|.$$

Informally, we say Π_{Mastic} is robust if there exists a PPT simulator Sim such that $\mathbf{Adv}_{\text{Mastic,Sim}}^{\text{rob}}(\mathcal{A})$ is negligible in the security parameter for all PPT adversaries $\mathcal{A}$.

We write $\mathbf{Adv}_{\text{H}}^{\text{coll}}(\cdot)$ to denote the probability of an attacker finding a collision against hash function H. We write $\mathbf{Adv}_{\mathcal{V},k}^{\text{verif}}(\cdot)$ to denote the probability of an attacker breaking verifiability of $\mathcal{V}$ at level $k \in [n]$ (we define this function in Appendix A). We write $\mathbf{Adv}_{\mathcal{Z}}^{\text{sound}}(\cdot)$ for the probability that an attacker breaks the soundness of $\mathcal{Z}$ (Appendix B).

THEOREM 2. *There exists a simulator* Sim *such that for all* $k \in [n]$ *and all* Π_{Mastic}*-robustness attackers* $\mathcal{A}$*, there exist a* $\mathcal{Z}$*-soundness attacker* $\mathcal{B}$*, an algorithm* C *for finding* H*-collisions, and a* $\mathcal{V}$*-verifiability attacker* $\mathcal{D}$ *such that*

$$\mathbf{Adv}_{\text{Mastic,Sim}}^{\text{rob}}(\mathcal{A}) \leq \mathsf{N}' \cdot \left(\mathbf{Adv}_{\mathcal{Z}}^{\text{sound}}(\mathcal{B}) + \right.$$
$$\left. n \cdot (\mathbf{Adv}_{H}^{\text{coll}}(C) + \mathbf{Adv}_{\mathcal{V},k}^{\text{verif}}(\mathcal{D})) \right),$$

where at most N' *clients are corrupted by* $\mathcal{A}$ *and the run times of* Sim*,* $\mathcal{B}$*,* C*, and* $\mathcal{D}$ *are upper bounded by the combined run time of honest aggregators in* Π_{Mastic} *and the run time of* $\mathcal{A}$*.*

Proof sketch. To argue robustness, the aggregators need to ensure that the report shares provided by a corrupt client encode a valid measurement (α_i, β_i). This involves checking three things.

- The weight β_i encoded inside the report shares is valid, i.e., $\beta_i \in \mathcal{L}$. This is ensured by the aggregators by evaluating the $\mathcal{V}$ keys at the root layer to obtain β_i, and then running the $\mathcal{Z}$ to validate β_i (without reconstructing β_i). An adversarial client who provides an invalid β_i, i.e., $\beta_i \notin \mathcal{L}$, but passes the checks of $\mathcal{Z}$ breaks soundness of $\mathcal{Z}$.
- Next, the aggregators must ensure that the same β_i value is propagated across a single path in the evaluation tree (encoded inside the $\mathcal{V}$ keys). This reduces to verifying that each level $k(\in [n])$ in the evaluation tree contains only a single non-zero node. This is ensured by verifying the $\mathcal{V}$ proofs for every node considered (as part of an evaluation path) during the computation of the heavy-hitter set. Suppose an adversarial client's report shares encode more than one non-zero node (as part of two evaluation paths), and those nodes are also encountered during the heavy-hitter evaluation. If such a client passes the checks, then we show how to use such an adversarial report to break the verifiability of $\mathcal{V}$.
- Lastly, the aggregator nodes must ensure that the non-zero nodes at each level are along the same path, say p. This is

performed by checking that the output for prefix p is equal to the sum of the output of its children, namely $p \parallel 0$ and $p \parallel 1$. This check is information-theoretic. By combining the one-hot verifiability and path-verifiability properties, we provide stronger guarantees where the adversarial input $\alpha_i \in \{0, 1\}^{\leq n}$ can be extracted from the report shares.

Finally, we optimize the communication by allowing the aggregator nodes to hash the results of the consistency checks for each client and then match the hash values. So, here we also need to rely on the collision-resistance of the hash function to ensure that if the hashes match then the underlying preimages are also equal. We defer the simulation-based proof to Appendix D.

6 Experimental Evaluation

Our goal in this section is to assess whether Mastic is efficient enough for the applications described in Section 4.

Setup. We perform experiments on two `c5.18xlarge` AWS instances, each with 72 vCPUs at 3.60 GHz, 144 GB memory, and 25 Gbps of network bandwidth. All our experiments are over a wide area network (WAN), with one server in Ohio (`us-east-2`) and the other in Oregon (`us-west-2`). We only focus on WAN as this is the most common way an MPC protocol will be deployed in the real world. We measure the runtime from the moment the aggregators receive all client report shares and start running the protocol. Mastic is implemented in Rust 1.74 and uses `tarpc` for asynchronous Remote Procedure Calls (RPC) and `rayon` for multi-threading. We conduct client-side experiments using a laptop with an Intel i7-8650U CPU (1.90 GHz). For heavy hitters (plain and weighted) we set the T to be 1% of the clients' bit strings.[8]

Target experiments. The goals of our experiments are to: 1) evaluate Mastic for weighted heavy-hitters and attribute-based metrics; 2) compare Mastic with related works; and 3) demonstrate the feasibility of Mastic for the applications described in Section 4 – i.e., NEL and attribute-based telemetry.

Weighted Heavy Hitters & Attribute-Based Metrics. Mastic is the only protocol for weighted heavy hitters and we use various weight sizes m (namely 5, 10, and 30) for a fixed number of bits n=256. Note that there are multiple ways to implement pruning based on the weights. For consistency with our plain heavy-hitter examples, we increase m by one and use the identity function for order($\cdot$); i.e., the last index of the aggregated weight counts the number of reports. This way, we can use the same threshold as in plain heavy-hitters, T = 1% of N. Additionally, we explore how the presence of malicious clients affects the protocol latency. For plain and attribute-based metrics, we fix the m to 100 and vary the number of attributes A between 1 (for plain metrics) to 128 and 1024 for attribute-based metrics. In all these cases, our field size is 128 bits.

Related Works. First, we compare Mastic to Poplar [11], the state of the art for *plain* heavy hitters in the two-aggregator setting. For plain heavy hitters with Mastic, we use $m = 1$ and a field size of 64 bits. Poplar's implementation[9] uses 62-bit fields for intermediate levels and 256 bits for the leaves.
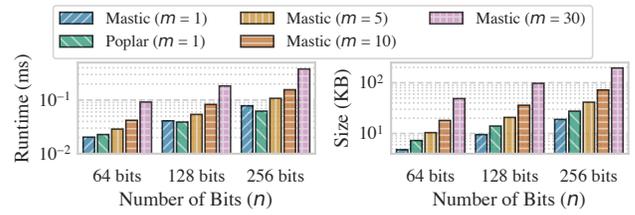
Figure 3: Report generation times (left) and report size (right) for different values of n and m.

Second, we wish to benchmark Mastic's attribute-based metrics mode with Prio [19]. Prio can provide a (less flexible) form of this functionality, as described in Section 4. Concretely, our goal is to aggregate a histogram of length $m = 100$ (as in Section 4.2) for each of A (fixed, in the case of Prio) attributes. We emulate this in Prio by aggregating a length $m \cdot A$ histogram, where the first m buckets correspond to the first attribute, the next m buckets correspond to the next attribute, and so on. We consider various numbers of attributes, $A = 1$, 128, and 1024. Correspondingly: for Mastic we set the bit length to $n = 1$, 7, and 10, respectively; and for Prio we set the input length to 100, 12800, and 102400, respectively. For Prio we use `libprio-rs`,[10] which implements the candidate standard [5]. We apply the same parallelization techniques across all protocols (Mastic, Poplar, and Prio) and run them in the same WAN setup.

Applications. Our last goal is to demonstrate the feasibility of Mastic for NEL and attribute-based browser telemetry. In this case, we use the exact parameters described in Sections 4.1 and 4.2, respectively.

Client Cost. First, we benchmark the client costs: the time it takes for a client to generate a report as well as the size of each report share (the message sent to each aggregator). Both of these costs vary based on the length of α and the size m of the weight β. We report the combination of three different sizes of α ($n = 64$, 128, and 256 bits) with $m = 1$ (for plain heavy hitters), $m = 5$, 10, and 30 (for weighted heavy hitters).

The client costs are shown in Fig. 3. On the left-hand side, we show the report generation time; as expected, the larger values of n and m take more time than the smaller variants, but in all cases, the report generation time is minimal. On the right-hand side of Fig. 3, we show the report sizes for the different α and m configurations. The smallest report size (64 bits, $m = 1$) is less than 5 KB, while the biggest key size (256 bits, $m = 30$) is less than 200 KB.

Prio is not suitable for weighted heavy hitters and therefore does not appear in Fig. 3. Instead, we compare the client cost for our application from Section 4.2: the total size of Mastic report shares for $n = 32$ and $m = 100$ uploaded by each client is about 53KB, while Prio for the same purpose with $A = 1024$ attributes needs about 1.6MB (30× overhead). Prio is also less flexible since the set of attributes would need to be known to the clients in advance.

Aggregator Cost. Next, we focus on aggregator costs for plain and weighted heavy hitters. In Fig. 4, we show the aggregator runtime over WAN for an increasing number of clients, 1k, 10k, 100k, and 1 million, and for $n = 256$ bit strings. For plain heavy hitters ($m = 1$), Mastic outperforms Poplar by almost an order of magnitude. In

---

[8]Mastic is open-source at https://github.com/TrustworthyComputing/mastic.
[9]https://github.com/henrycg/heavyhitters

[10]https://github.com/divviup/libprio-rs
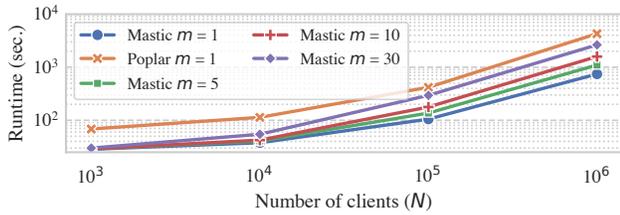
Figure 4: *Plain* ($m = 1$) **and** *Weighted* ($m > 1$) *heavy hitters* **for** $n = 256$**. Aggregator runtime (WAN) for an increasing** N.
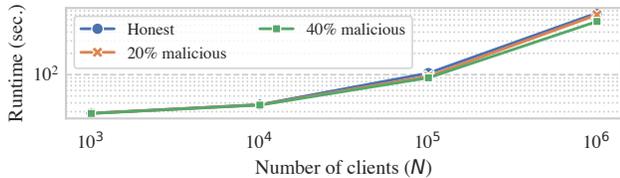
Figure 5: Mastic plain heavy-hitters with malicious clients.

the same figure, we show Mastic's aggregator runtime for three different weight sizes (i.e., m = 5, 10, and 30). As expected, the bigger m values have an impact on the runtime of the protocol; due both to increased communication and increased computation (i.e., path verifiability now has to consider m values). For 10^6 clients, going from $m = 5$ (18 minutes) to $m = 30$ (44 minutes) more than doubles the latency (but still outperforms Poplar).

In Fig. 5, we vary the number of malicious clients for plain heavy hitters between 20% and 40% of the total number of clients. We observe that a higher percentage of malicious clients results in lower latency for Mastic, which is primarily because malicious clients fail to pass the shared ZK verification so they are eliminated from the protocol. If the weight check succeeds, but the path or one-hot check fails, Mastic performs similarly to PLASMA [35]. Note that Mastic exhibits the same scaling for weighted heavy hitters as with plain, as the only difference between the two is the size of the weight. The Poplar paper did not evaluate how it scales in the presence of malicious clients but we expect a similar trend.

In Fig. 6, we show the aggregator cost for attribute-based metrics with Mastic versus Prio. With this experiment, we aim to answer the question: *"at what cost can Mastic replace Prio?"*. As expected, for a single attribute ($A = 1$) Prio is faster than Mastic after 10^5 clients as it uses $n = 0$. However, as we increase the value of A, we observe that Mastic quickly outperforms Prio (about $1.5 - 2\times$ as fast). Another benefit of Mastic is that the attributes do not need to be fixed *a priori*; Mastic can support any number.

Suitability of Mastic for NEL. The NEL application (Section 4.1) is time-sensitive: the sooner the results are available, the sooner they can be acted upon to diagnose and resolve the issue that precipitated the errors. As shown in Fig. 4, for realistic parameters ($n = 256$; $m = 30$; N $= 10^6$ clients), Mastic takes 44 minutes to compute the high-error rate domains (i.e., inputs α) and their error distributions. The end-to-end latency could be further improved by issuing parallel RPCs and utilizing all the available network bandwidth. Still, waiting even 5 minutes for results might be too long, depending on the conditions. Fortunately, Mastic produces
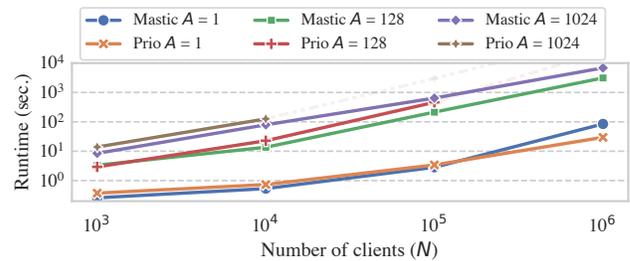
Figure 6: *Plain* ($A = 1$) **and** *Attribute-based* ($A > 1$) *metrics* **with** $m = 100$**. For Mastic, we use** $n = \lfloor \log_2(A + 1) \rfloor$**. For Prio, the faded-out lines represent extrapolated values as the evaluation did not finish after exceeding the available server memory. Aggregator runtime over WAN for an increasing** N.

the errors themselves much faster, immediately after evaluating the first level of the prefix tree for each report (in just 10.4 seconds in our experiment). To summarize, Mastic is the first work that enables such an application; this can be used in a real NEL instantiation as long as some latency for learning the domains can be tolerated.

Suitability of Mastic for Browser Telemetry. Lastly, we consider the parameters described in Section 4.2, i.e., $n = 32$, $m = 100$ and 200 with a 128-bit field. For a single attribute, the latency for Mastic is 103 and 168 seconds for 1 million clients for $m = 100$ and $m = 200$, respectively. Going up to 10 attributes, Mastic takes approximately 9 minutes for $m = 100$. Based on this result, we believe it is clear that Mastic is concretely efficient enough for this application.

7 Concluding Remarks

This work presents Mastic, the first two-server MPC protocol for general-purpose metrics that supports both weighted heavy-hitters and aggregation grouped by user attributes. Mastic offers notable benefits over the previous state-of-the-art Prio and Poplar frameworks that focused on either plain aggregations or plain heavy hitters. At its core, our protocol leverages verifiable incremental distributed point functions along with shared zero-knowledge proofs to enable privacy and robustness in the presence of malicious clients and privacy in the presence of a malicious aggregator. Mastic is efficient for real-life applications, such as network error logging and attribute-based telemetry, and outperforms Poplar and Prio on heavy-hitters and attribute-based scenarios, respectively.

References

[1] Syed Suleman Ahmad. 2023. Part 2: Exploring heavy hitter applications (continued). PPM mailing list archive: https://mailarchive.ietf.org/arch/msg/ppm/6DEXt2PS3IGOGWcg5hHzgsIGYyg/.
[2] Erik Anderson, Melissa Chase, F. Betul Durak, Esha Ghosh, Kim Laine, and Chenkai Weng. 2021. Aggregate Measurement via Oblivious Shuffling. Cryptology ePrint Archive, Report 2021/1490. https://eprint.iacr.org/2021/1490.

[3] Apple. 2023. Learning Iconic Scenes with Differential Privacy. https://machinelearning.apple.com/research/scenes-differential-privacy.

[4] Gilad Asharov, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Ariel Nof, Benny Pinkas, Katsumi Takahashi, and Junichi Tomida. 2022. Efficient Secure Three-Party Sorting with Applications to Data Analysis and Heavy Hitters. In *ACM CCS 2022: 29th Conference on Computer and Communications Security*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM Press, Los Angeles, CA, USA, 125–138. https://doi.org/10.1145/3548606.3560691

[5] Richard Barnes, David Cook, Christopher Patton, and Phillipp Schoppmann. 2022. Verifiable Distributed Aggregation Functions. Internet Engineering Task Force (IETF), Internet-Draft. https://datatracker.ietf.org/doc/draft-irtf-cfrg-vdaf.

[6] Carsten Baum, Ivan Damgård, and Claudio Orlandi. 2014. Publicly Auditable Secure Multi-Party Computation. In *SCN 14: 9th International Conference on Security in Communication Networks (Lecture Notes in Computer Science, Vol. 8642)*, Michel Abdalla and Roberto De Prisco (Eds.). Springer, Heidelberg, Germany, Amalfi, Italy, 175–196. https://doi.org/10.1007/978-3-319-10879-7_11

[7] James Bell, Adrià Gascón, Badih Ghazi, Ravi Kumar, Pasin Manurangsi, Mariana Raykova, and Phillipp Schoppmann. 2022. Distributed, Private, Sparse Histograms in the Two-Server Model. In *ACM CCS 2022: 29th Conference on Computer and Communications Security*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM Press, Los Angeles, CA, USA, 307–321. https://doi.org/10.1145/3548606.3559383

[8] Burton H Bloom. 1970. Space/time trade-offs in hash coding with allowable errors. *Commun. ACM* 13, 7 (1970), 422–426.

[9] Jonas Böhler and Florian Kerschbaum. 2021. Secure Multi-party Computation of Differentially Private Heavy Hitters. In *ACM CCS 2021: 28th Conference on Computer and Communications Security*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, Virtual Event, Republic of Korea, 2361–2377. https://doi.org/10.1145/3460120.3484557

[10] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2019. Zero-Knowledge Proofs on Secret-Shared Data via Fully Linear PCPs. In *Advances in Cryptology – CRYPTO 2019, Part III (Lecture Notes in Computer Science, Vol. 11694)*, Alexandra Boldyreva and Daniele Micciancio (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 67–97. https://doi.org/10.1007/978-3-030-26954-8_3

[11] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2021. Lightweight Techniques for Private Heavy Hitters. In *2021 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, San Francisco, CA, USA, 762–776. https://doi.org/10.1109/SP40001.2021.00048

[12] Dan Boneh, Elette Boyle, Henry Corrigan-Gibbs, Niv Gilboa, and Yuval Ishai. 2023. Arithmetic Sketching. In *Advances in Cryptology – CRYPTO 2023, Part I (Lecture Notes in Computer Science, Vol. 14081)*, Helena Handschuh and Anna Lysyanskaya (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 171–202. https://doi.org/10.1007/978-3-031-38557-5_6

[13] Elette Boyle, Niv Gilboa, and Yuval Ishai. 2015. Function Secret Sharing. In *Advances in Cryptology – EUROCRYPT 2015, Part II (Lecture Notes in Computer Science, Vol. 9057)*, Elisabeth Oswald and Marc Fischlin (Eds.). Springer, Heidelberg, Germany, Sofia, Bulgaria, 337–367. https://doi.org/10.1007/978-3-662-46803-6_12

[14] Ran Canetti. 2000. Security and Composition of Multiparty Cryptographic Protocols. *Journal of Cryptology* 13, 1 (Jan. 2000), 143–202. https://doi.org/10.1007/s001459910006

[15] Koji Chida, Daniel Genkin, Koki Hamada, Dai Ikarashi, Ryo Kikuchi, Yehuda Lindell, and Ariel Nof. 2018. Fast Large-Scale Honest-Majority MPC for Malicious Adversaries. In *Advances in Cryptology – CRYPTO 2018, Part III (Lecture Notes in Computer Science, Vol. 10993)*, Hovav Shacham and Alexandra Boldyreva (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 34–64. https://doi.org/10.1007/978-3-319-96878-0_2

[16] Cloudflare. 2022. Stop attacks before they are known: making the Cloudflare WAF smarter. https://blog.cloudflare.com/stop-attacks-before-they-are-known-making-the-cloudflare-waf-smarter.

[17] Cloudflare. 2023. Privacy-preserving measurement and machine learning. https://blog.cloudflare.com/deep-dive-privacy-preserving-measurement.

[18] Graham Cormode and S. Muthukrishnan. 2004. An Improved Data Stream Summary: The Count-Min Sketch and Its Applications. In *LATIN 2004: Theoretical Informatics, 6th Latin American Symposium (Lecture Notes in Computer Science, Vol. 2976)*, Martin Farach-Colton (Ed.). Springer, Heidelberg, Germany, Buenos Aires, Argentina, 29–38.

[19] Henry Corrigan-Gibbs and Dan Boneh. 2017. Prio: Private, Robust, and Scalable Computation of Aggregate Statistics. In *Proceedings of the 14th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) (*NSDI'17*). USENIX Association, USA, 259–282.

[20] Ivan Damgård, Daniel Escudero, Tore Kasper Frederiksen, Marcel Keller, Peter Scholl, and Nikolaj Volgushev. 2019. New Primitives for Actively-Secure MPC over Rings with Applications to Private Machine Learning. In *2019 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, San Francisco, CA, USA, 1102–1120. https://doi.org/10.1109/SP.2019.00078

[21] Alex Davidson, Peter Snyder, E. B. Quirk, Joseph Genereux, Benjamin Livshits, and Hamed Haddadi. 2022. STAR: Secret Sharing for Private Threshold Aggregation Reporting. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* (Los Angeles, CA, USA) (*CCS '22*). Association for Computing Machinery, New York, NY, USA, 697–710. https://doi.org/10.1145/3548606.3560631

[22] Hannah Davis, Christopher Patton, Mike Rosulek, and Phillipp Schoppmann. 2023. Verifiable Distributed Aggregation Functions. *Proceedings on Privacy Enhancing Technologies* 2023, 4 (July 2023), 578–592. https://doi.org/10.56553/popets-2023-0126

[23] Leo de Castro and Antigoni Polychroniadou. 2022. Lightweight, Maliciously Secure Verifiable Function Secret Sharing. In *Advances in Cryptology – EUROCRYPT 2022, Part I (Lecture Notes in Computer Science, Vol. 13275)*, Orr Dunkelman and Stefan Dziembowski (Eds.). Springer, Heidelberg, Germany, Trondheim, Norway, 150–179. https://doi.org/10.1007/978-3-031-06944-4_6

[24] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response. In *ACM CCS 2014: 21st Conference on Computer and Communications Security*, Gail-Joon Ahn, Moti Yung, and Ninghui Li (Eds.). ACM Press, Scottsdale, AZ, USA, 1054–1067. https://doi.org/10.1145/2660267.2660348

[25] Sanjam Garg, Abhishek Jain, Pratyay Mukherjee, Rohit Sinha, Mingyuan Wang, and Yinuo Zhang. 2023. Cryptography with Weights: MPC, Encryption and Signatures. In *Advances in Cryptology – CRYPTO 2023, Part I (Lecture Notes in Computer Science, Vol. 14081)*, Helena Handschuh and Anna Lysyanskaya (Eds.). Springer, Heidelberg, Germany, Santa Barbara, CA, USA, 295–327. https://doi.org/10.1007/978-3-031-38557-5_10

[26] Tim Geoghegan. 2023. DAP draft repository issue #489. https://github.com/ietf-wg-ppm/draft-ietf-ppm-dap/issues/489.

[27] Tim Geoghegan, Christopher Patton, Brandon Pitman, Eric Rescorla, and Christopher Wood. 2021. Distributed Aggregation Protocol for Privacy Preserving Measurement. Internet Engineering Task Force (IETF), Internet-Draft. https://datatracker.ietf.org/doc/draft-ietf-ppm-dap.

[28] Bobby Holley. 2023. Built for privacy: Partnering to deploy Oblivious HTTP and Prio in Firefox. https://blog.mozilla.org/en/products/firefox/partnership-ohttp-prio.

[29] IETF. 2022. Charter for Privacy Preserving Measurement Working Group. Internet Engineering Task Force (IETF). https://datatracker.ietf.org/doc/charter-ietf-ppm/.

[30] Internet Security Research Group (ISRG). 2024. Divvi Up: A privacy-respecting system for aggregate statistics. https://divviup.org.

[31] Pranav Jangir, Nishat Koti, Varsha Bhat Kukkala, Arpita Patra, Bhavish Raj Gopal, and Somya Sangal. 2022. Poster: Vogue: Faster Computation of Private Heavy Hitters. In *ACM CCS 2022: 29th Conference on Computer and Communications Security*, Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi (Eds.). ACM Press, Los Angeles, CA, USA, 3371–3373. https://doi.org/10.1145/3548606.3563544

[32] Marcel Keller. 2020. MP-SPDZ: A Versatile Framework for Multi-Party Computation. In *ACM CCS 2020: 27th Conference on Computer and Communications Security*, Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna (Eds.). ACM Press, Virtual Event, USA, 1575–1590. https://doi.org/10.1145/3372297.3417872

[33] Hanjun Li, Sela Navot, and Stefano Tessaro. 2024. POPSTAR: Lightweight Threshold Reporting with Reduced Leakage. Cryptology ePrint Archive, Paper 2024/320. https://eprint.iacr.org/2024/320 https://eprint.iacr.org/2024/320.

[34] Dimitris Mouris, Daniel Masny, Ni Trieu, Shubho Sengupta, Prasad Buddhavarapu, and Benjamin M Case. 2024. Delegated Private Matching for Compute. *Proceedings on Privacy Enhancing Technologies* 2024, 2 (July 2024), 1–24.

[35] Dimitris Mouris, Pratik Sarkar, and Nektarios Georgios Tsoutsos. 2024. PLASMA: Private, Lightweight Aggregated Statistics against Malicious Adversaries. *Proceedings on Privacy Enhancing Technologies* 2024, 3 (July 2024), 1–19.

[36] MPC Deployments. 2023. The Deployment Dilemma: Merits & Challenges of Deploying MPC. https://mpc.cs.berkeley.edu/blog/deployment-dilemma.html.

[37] Kartik Nayak, Xiao Shaun Wang, Stratis Ioannidis, Udi Weinsberg, Nina Taft, and Elaine Shi. 2015. GraphSC: Parallel Secure Computation Made Easy. In *2015 IEEE Symposium on Security and Privacy*. IEEE Computer Society Press, San Jose, CA, USA, 377–394. https://doi.org/10.1109/SP.2015.30

[38] Private Advertising Technology Community Group (PATCG). 2023. Private Advertising Technology Community Group Charter. https://patcg.github.io/charter.html.

[39] Zhan Qin, Yin Yang, Ting Yu, Issa Khalil, Xiaokui Xiao, and Kui Ren. 2016. Heavy Hitter Estimation over Set-Valued Data with Local Differential Privacy. In *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi (Eds.). ACM Press, Vienna, Austria, 192–203. https://doi.org/10.1145/2976749.2978409

[40] M. Rathee, Y. Zhang, H. Corrigan-Gibbs, and R. Ada Popa. 2024. Private Analytics via Streaming, Sketching, and Silently Verifiable Proofs. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE Computer Society, Los Alamitos, CA, USA, 194–194. https://doi.org/10.1109/SP54263.2024.00245

[41] Eric Rescorla. 2022. DAP draft repository issue #183. https://github.com/ietf-wg-ppm/draft-ietf-ppm-dap/issues/183.

[42] Trey Titone. 2022. Interoperable Private Attribution (IPA) Explained. https://adtechexplained.com/interoperable-private-attribution-ipa-explained.

[43] Vincent Toubiana, Arvind Narayanan, Dan Boneh, Helen Nissenbaum, and Solon Barocas. 2010. Adnostic: Privacy Preserving Targeted Advertising. In *ISOC Network and Distributed System Security Symposium – NDSS 2010*. The Internet Society, San Diego, CA, USA, 1–16.

[44] W3C Working Group. 2023. Network Error Logging. https://www.w3.org/TR/network-error-logging/.

[45] Kang Yang, Pratik Sarkar, Chenkai Weng, and Xiao Wang. 2021. QuickSilver: Efficient and Affordable Zero-Knowledge Proofs for Circuits and Polynomials over Any Field. In *ACM CCS 2021: 28th Conference on Computer and Communications Security*, Giovanni Vigna and Elaine Shi (Eds.). ACM Press, Virtual Event, Republic of Korea, 2986–3001. https://doi.org/10.1145/3460120.3484556

[46] Wennan Zhu, Peter Kairouz, Brendan McMahan, Haicheng Sun, and Wei Li. 2020. Federated Heavy Hitters Discovery with Differential Privacy. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 108)*, Silvia Chiappa and Roberto Calandra (Eds.). PMLR, Online, 3837–3847. https://proceedings.mlr.press/v108/zhu20a.html

Appendix

A   Security Definitions for VIDPF

Let $\mathcal{V} = (\text{Gen}, \text{Eval}, \text{EvalRoot}, \text{Valid})$ be a VIDPF as defined in Section 2.3. We recall the security definitions for $\mathcal{V}$ of [35], adapted to our refined syntax.

Correctness. For all inputs $\alpha \in \{0, 1\}^n$, all weights $\beta \in \mathbb{F}^m$, all $k \in [n]$, and all $p \in \{0, 1\}^k$, when we let $(\text{pub}, \text{key}_0, \text{key}_1) \coloneqq \text{Gen}(\alpha, \beta)$ and $([\![y^{p|^i}]\!]_b, \text{st}_b^{p|^i}, \pi_b^{p|^i}) \coloneqq \text{Eval}(\text{key}_b, \text{pub}, p|^i, \text{st}_b^{p|^{i-1}}, \pi_b^{p|^{i-1}})$ for all $i \le k$ and $b \in \{0, 1\}$, it holds that $\text{Valid}(\pi_0^p, \pi_1^p)$ and

$$[\![y^p]\!]_0 + [\![y^p]\!]_1 = \begin{cases} \beta & \text{if } p = \alpha|^k \\ 0^m & \text{otherwise.} \end{cases}$$

Privacy. The information revealed to each aggregator leaks nothing about the underlying measurement. We formalize this via an indistinguishability game. First, for adversary $\mathcal{A}$ and challenge bit c, define $\mathcal{G}_{\mathcal{V},c}^{\text{priv}}(\mathcal{A})$ as the following experiment:

(1) Run $\mathcal{A}$ to get $(\alpha^0, \beta^0) \in \{0, 1\}^n \times \mathbb{F}^m$ and corrupt aggregator index $b \in \{0, 1\}$.
(2) Sample (α^1, β^1) uniformly from $\{0, 1\}^n \times \mathcal{L}$.
(3) Run $(\text{pub}, \text{key}_0, \text{key}_1) \coloneqq \text{Gen}(\alpha^c, \beta^c)$ and give the public share and corrupt aggregator's key $(\text{pub}, \text{key}_b)$ to $\mathcal{A}$.
(4) Run $\mathcal{A}$ to get its guess of the challenge bit and output it.

The advantage of $\mathcal{A}$ in attacking the privacy of $\mathcal{V}$ is

$$\mathbf{Adv}_{\mathcal{V}}^{\text{priv}}(\mathcal{A}) \coloneqq \left| \Pr[\mathcal{G}_{\mathcal{V},1}^{\text{priv}}(\mathcal{A}) \Rightarrow 1] - \Pr[\mathcal{G}_{\mathcal{V},0}^{\text{priv}}(\mathcal{A}) \Rightarrow 1] \right|.$$

Informally, we say that $\mathcal{V}$ is private if no PPT adversary gets a non-negligible advantage in breaking its privacy.

Note that, if the $\mathcal{V}$ is correct, then an aggregator learns nothing from its interaction with its peer. This is because both parties compute the same proof string for each prefix.

Verifiability. No corrupt client can construct $(\text{pub}, \text{key}_0, \text{key}_1)$ for which the portion of the prefix tree traversed by the aggregators contains more than one non-zero node at any level $k \in [n]$. Formally, we define the advantage $\mathbf{Adv}_{\mathcal{V},k}^{\text{verif}}(\mathcal{A})$ of adversary $\mathcal{A}$ in breaking the verifiability of $\mathcal{V}$ at level k as the probability that the following experiment outputs true:

(1) Run $\mathcal{A}(1^k)$ to get $(\text{pub}, \text{key}_0, \text{key}_1)$ and $u, v \in \{0, 1\}^k$.

(2) Evaluate the prefix tree for $(\text{pub}, \text{key}_0, \text{key}_1)$ and u, v. That is, for all $(i, b) \in [k] \times \{0, 1\}$ let

$$([\![y^{u|^i}]\!]_b, \text{st}_b^{u|^i}, \pi_b^{u|^i}) \coloneqq \text{Eval}(\text{key}_b, \text{pub}, u|^i, \text{st}_b^{u|^{i-1}}, \pi_b^{u|^{i-1}})$$

$$([\![y^{v|^i}]\!]_b, \text{st}_b^{v|^i}, \pi_b^{v|^i}) \coloneqq \text{Eval}(\text{key}_b, \text{pub}, v|^i, \text{st}_b^{v|^{i-1}}, \pi_b^{v|^{i-1}})$$

(3) Output $([\![y^u]\!]_0 + [\![y^u]\!]_1 \ne 0^m) \wedge ([\![y^v]\!]_0 + [\![y^v]\!]_1 \ne 0^m) \wedge (\pi_0^u = \pi_1^u) \wedge (\pi_0^v = \pi_1^v) \wedge (u \ne v)$.

The adversary wins if the output is true, i.e., it finds two distinct, equal length paths in the prefix tree for which 1) both nodes are non-zero and 2) the proofs for path u and v are valid. We say that $\mathcal{V}$ is ϵ-*verifiable* if for all PPT adversaries $\mathcal{A}$ and $k \in [n]$ it holds that $\mathbf{Adv}_{\mathcal{V},k}^{\text{verif}}(\mathcal{A}) \le \epsilon$.

PLASMA [35] proved that their $\mathcal{V}$ construction satisfied this property based on the XOR-collision resistance of the hash function used in the construction. [23] proved the verifiability property of their verifiable distributed point function based on the XOR-collision resistance property of a hash function. The hash function is modeled as a random oracle to prove XOR-collision resistance. Mastic requires the same assumption as it uses the same VIDPF construction.

B   Security Definitions for shared ZK

Let $\mathcal{Z} = (\text{Prove}, \text{Query}, \text{Decide}, \text{Dom}, \text{Rng})$ be a shared ZK scheme as defined in Section 2.4. In this section, we define the security properties we require for $\mathcal{Z}$. These are comparable to definitions from [10, Section 6] but apply to our syntax.

Completeness. We say that $\mathcal{Z}$ is *complete* if for all $b \in \{0, 1\}$, $vk \in \{0, 1\}^{\text{vkl}}$, $\text{H} : \text{Dom} \to \text{Rng}$, and all $[\![x]\!]_0, [\![x]\!]_1$ such that $\text{Extract}([\![x]\!]_0, [\![x]\!]_1) \in \mathcal{L}$, it holds that

$$\Pr[\text{Decide}^{\text{H}}(\sigma_0, \sigma_1, \text{st}_b) = 1] = 1,$$

where the state and verifier shares were generated by running $(\pi_0^{\text{szk}}, \pi_1^{\text{szk}}, \text{nonce}) \coloneqq \text{Prove}^{\text{H}}([\![x]\!]_0, [\![x]\!]_1)$ then $(\text{st}_b, \sigma_b) \coloneqq \text{Query}^{\text{H}}(vk, \text{nonce}, [\![x]\!]_b, \pi_b^{\text{szk}})$ for $b \in \{0, 1\}$.

Zero-Knowledge. The aggregators should learn nothing about any honest client's input while validating it as long as one aggregator is honest. To formalize this, let $\mathcal{A}$ be an adversary (the corrupt aggregator) and let Sim be a simulator with interfaces Init, Prove, Query, and Decide. We define games $\mathcal{G}_{\mathcal{Z}}^{\text{shared ZK-real}}(\mathcal{A})$ and $\mathcal{G}_{\text{Sim}}^{\text{shared ZK-ideal}}(\mathcal{A})$ as shown in Figs. 7 and 8, respectively. In these games, the adversary may request the honest client to prove the validity of a shared input x, and it may ask the honest aggregator to run the SZK query or decision algorithms on the input and proof shares it has received from any honest client. Define the advantage of $\mathcal{A}$ in distinguishing execution of $\mathcal{Z}$ from simulator Sim as

$$\mathbf{Adv}_{\mathcal{Z},\text{Sim}}^{\text{priv}}(\mathcal{A}) \coloneqq \Big| \Pr[\mathcal{G}_{\mathcal{Z}}^{\text{shared ZK-real}}(\mathcal{A}) \Rightarrow 1]$$
$$- \Pr[\mathcal{G}_{\text{Sim}}^{\text{shared ZK-ideal}}(\mathcal{A}) \Rightarrow 1] \Big|.$$

Informally, we say that $\mathcal{Z}$ is *zero-knowledge* if there exists a PPT simulator for which no PPT adversary gets a non-negligible distinguishing advantage.

Soundness. The proof system is sound if a malicious client cannot fool honest aggregators into accepting an invalid measurement. Formally, define the advantage $\mathbf{Adv}_{\mathcal{Z}}^{\text{sound}}(\mathcal{A})$ of $\mathcal{A}$ in breaking the

$\mathcal{G}_{\mathcal{Z}}^{\text{shared ZK-real}}$

Init(vk, b)

1. $H \xleftarrow{\$} \mathcal{AF}(\text{Dom}, \text{Rng})$
2. Store vk and b in memory.

Prove$(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1)$

We require that $\text{Extract}(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1) \in \mathcal{L}$

1. $(\text{nonce}, \pi^{\text{szk}}{}_0, \pi^{\text{szk}}{}_1) \xleftarrow{\$} \mathcal{Z}.\text{Prove}^H (\llbracket x \rrbracket_0, \llbracket x \rrbracket_1)$
2. $T_{\text{client}}[\text{nonce}] \leftarrow \llbracket x \rrbracket_{1-b}, \pi^{\text{szk}}{}_{1-b}$
3. return $(\text{nonce}, \pi^{\text{szk}}{}_b)$

Query(nonce)

1. $(\llbracket \sigma \rrbracket, \text{st}) \leftarrow \mathcal{Z}.\text{Query}^H(vk, 1-b, \text{nonce}, T_{\text{client}}[\text{nonce}])$
2. $T_{\text{szk}}[\text{nonce}] \leftarrow \llbracket \sigma \rrbracket, \text{st}$
3. return $\llbracket \sigma \rrbracket$

Decide$(\text{nonce}, \llbracket \sigma \rrbracket_b)$

1. $\llbracket \sigma \rrbracket_{1-b}, \text{st} \leftarrow T_{\text{szk}}[\text{nonce}]$
2. return $\mathcal{Z}.\text{Decide}^H(\llbracket \sigma \rrbracket_0, \llbracket \sigma \rrbracket_1, \text{st})$

RO(d)

1. return $H(d)$
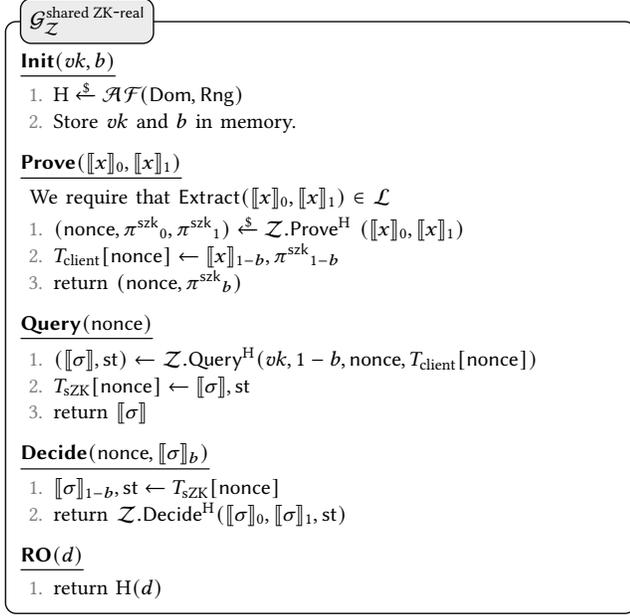
Figure 7: Real game for defining zero-knowledge of shared ZK scheme $\mathcal{Z}$.

$\mathcal{G}_{\text{Sim}_{\text{SZK}}}^{\text{shared ZK-ideal}}$

Init(vk, b)

1. Store vk and b in memory.

Prove$(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1)$

We require that $\text{Extract}(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1) \in \mathcal{L}$

1. $(\text{nonce}, \pi^{\text{szk}}{}_b, \text{st}) \xleftarrow{\$} \text{Sim}_{\text{SZK}}.\text{Prove}(\text{st}, b)$
2. return $(\text{nonce}, \pi^{\text{szk}}{}_b)$

Query(nonce)

1. $(\llbracket \sigma \rrbracket, \text{st}) \leftarrow \text{Sim}_{\text{SZK}}.\text{Query}(vk, 1-b, \text{nonce}, \text{st})$
2. return $\llbracket \sigma \rrbracket$

Decide$(\text{nonce}, \llbracket \sigma \rrbracket_b)$

1. return $\text{Sim}_{\text{SZK}}.\text{Decide}(\llbracket \sigma \rrbracket_b, 1-b, \text{nonce}, \text{st})$

RO(i, d)

1. $\text{st}, h \leftarrow \text{Sim}_{\text{SZK}}.\text{RO}(\text{st}, i, d)$
2. return h
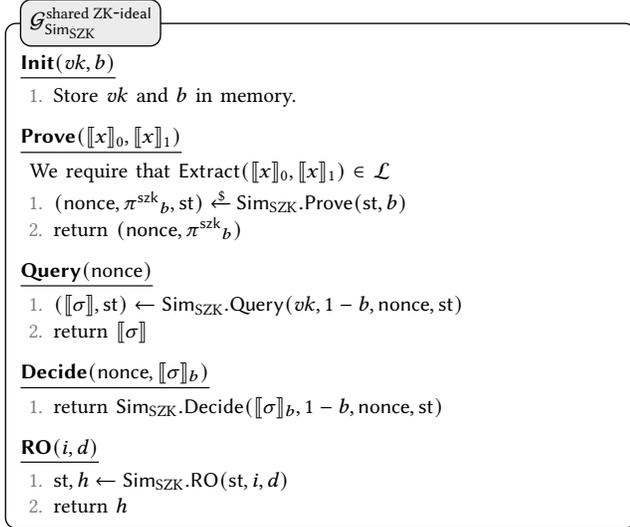
Figure 8: Ideal game for defining zero-knowledge for shared ZK scheme $\mathcal{Z}$.

soundness of $\mathcal{Z}$ as the probability that the following experiment outputs true:

(1) Run $\mathcal{A}$ to get $(\text{nonce}, \llbracket x \rrbracket_0, \llbracket x \rrbracket_1, \pi_0^{\text{szk}}, \pi_1^{\text{szk}})$.
(2) Sample $vk \xleftarrow{\$} \{0, 1\}^{\text{vkl}}$.
(3) Run $(\text{st}_b, \llbracket \sigma \rrbracket_b) := \text{Query}(vk, \text{nonce}, \llbracket x \rrbracket_b, \llbracket \pi^{\text{szk}} \rrbracket_b)$ for each $b \in \{0, 1\}$.
(4) Output $\big(\exists b\ \text{Decide}(\llbracket \sigma \rrbracket_0, \llbracket \sigma \rrbracket_1, \text{st}_b) = 1\big) \wedge \text{Extract}(\llbracket x \rrbracket_0, \llbracket x \rrbracket_1) \notin \mathcal{L}$.

The adversary wins if the output is true, i.e., the measurement is invalid but the proof verifies. Informally, we say that $\mathcal{Z}$ is ϵ-sound if for all PPT $\mathcal{A}$ it holds that $\mathbf{Adv}_{\mathcal{Z}}^{\text{sound}}(\mathcal{A}) \leq \epsilon$.

B.1 Constructing Shared ZK from Fully Linear Proofs

Similar to the candidate standard for Prio [5], the shared ZK system $\mathcal{Z}$ can be instantiated from a fully linear proof (FLP) system [22]. In the remainder, we give the high level idea of the construction.

First, FLPs are designed to provide the same properties as conventional zero-knowledge proof systems (for membership in a finite language $\mathcal{L} \subseteq \mathbb{F}^m$); in a addition, they are "fully linear" in the sense the verifier's computations over the input and proof is linear. This allows verification to be distributed amongst multiple verifiers, each of which holds only a share of the input and proof.

FLP is the core component of the shared ZK proof system; what remains for $\mathcal{Z}$ to specify is 1) secret sharing of the FLP itself and 2) generation of shared randomness used by the prover and verifiers. The FLP used in Prio [5, Section 7.3] involves "joint randomness" used for proof generation and evaluation and "query randomness" used by the verifiers to check the proof's correctness: in Prio, the former is derived from the input shares using an extension of the Fiat-Shamir heuristic to proofs on secret shared data; the query randomness is derived by applying a PRF to the nonce generating by client and using the key vk shared by the verifiers. The full construction is defined in Figure 9.

We refer to [22, Section 4] for concrete security bounds for the underlying FLP system. Note that their security model for privacy and robustness does not immediately yield bounds for our setting.

C Proof of Theorem 1 (Π_{Mastic} is Private)

In this section, we prove Theorem 1, which for convenience we restate here.

THEOREM. *For any simulator Sim_{SZK}, there exists a simulator Sim such that for any adversary $\mathcal{A}$, there exist $\mathcal{Z}$-attacker $\mathcal{B}$ and $\mathcal{V}$-attacker $\mathcal{D}$ such that*

$$\mathbf{Adv}_{\text{Mastic,Sim}}^{\text{priv}}(\mathcal{A}) \leq \mathbf{Adv}_{\mathcal{Z}, \text{Sim}_{\text{SZK}}}^{\text{priv}}(\mathcal{B}) + q \cdot \mathbf{Adv}_{\mathcal{V}}^{\text{priv}}(\mathcal{D}),$$

where q is the number of queries made to the "Honest Client Computation" oracle, and the runtime of $\mathcal{D}$ is about that of an honest aggregator in the Mastic protocol, and the runtime of $\mathcal{B}$ is about that of an honest aggregator plus the time to run Sim_{SZK} once per interaction with the honest aggregator.

The privacy advantage of $\mathcal{A}$ against the mastic protocol is its advantage in distinguishing between games $\mathcal{G}_{\text{Mastic}}^{\text{priv-real}}$ (c.f. Figure 10) and $\mathcal{G}_{\text{Sim}}^{\text{priv-ideal}}$ (c.f. Figure 11). We will establish the claim by constructing a simulator Sim and adversaries $\mathcal{B}_{\mathcal{Z}}$ against the zero-knowledge security of the shared ZK scheme and $\mathcal{B}_{\mathcal{V}}$ against the privacy of the VIDPF scheme. Then we transform game $\mathcal{G}_{\text{Sim}}^{\text{priv-ideal}}$ into game $\mathcal{G}_{\text{Mastic}}^{\text{priv-real}}$ via a series of hybrid transitions and bound the distinguishing advantage between each pair of hybrids.

First, we define simulator Sim in Figure 12. This simulator mimics the behavior of an honest aggregator in the Mastic protocol without knowledge of any client's underlying measurements. Broadly, the simulator handles two types of client reports: honest clients

---

Shared ZK Proof from Fully Linear Proof

$\mathcal{Z}.\text{Prove}[\mathcal{F}, \text{RO}]([\![x]\!]_0, [\![x]\!]_1)$

    (1) $\text{seed}_0, \text{seed}_1, \text{nonce} \xleftarrow{\$} \{0,1\}^{\text{kl}}$

    (2) $x \leftarrow [\![x]\!]_0 + [\![x]\!]_1$

    (3) for $i \in \{0, 1\}$:

      (a) $\text{jrpart}_i \leftarrow \text{RO}_1(\text{seed}_i \| \text{nonce} \| [\![x]\!]_i)$

    (4) $\text{jrseed} \leftarrow \text{RO}_2(\text{jrpart}_0 \| \text{jrpart}_1)$

    (5) $\text{jrand} \leftarrow \text{RO}_3(\text{jrseed})$

    (6) $\pi_{\text{FLP}} \xleftarrow{\$} \mathcal{F}.\text{Prove}(x; \text{jrand})$

    (7) $\pi_1 \leftarrow \text{RO}_4(\text{pf}, \text{seed}_1 \| \text{nonce} \| [\![x]\!]_1)$

    (8) $\pi_0 \leftarrow \pi_{\text{FLP}} - \pi_1$

    (9) return $(\pi_0, \text{seed}_0, \text{jrpart}_1), (\text{seed}_1, \text{jrpart}_0)$

$\mathcal{Z}.\text{Query}[\mathcal{F}, \text{RO}](vk, i, \text{nonce}, [\![x]\!]_i, \pi_i, \text{seed}_i, \text{jrpart}_{1-i})$

    (1) if $i = 1$ $\pi_1 \leftarrow \text{RO}(\text{seed}_1 \| \text{nonce} \| [\![x]\!]_1)$

    (2) $\text{jrpart}_i \leftarrow \text{RO}_1(\text{seed}_i \| \text{nonce} \| [\![x]\!]_i)$

    (3) $\text{jrseed} \leftarrow \text{RO}_2(\text{jrpart}_0 \| \text{jrpart}_1)$

    (4) $\text{jrand} \leftarrow \text{RO}_3(\text{jrseed})$

    (5) $\text{qrand} \leftarrow \text{RO}_5(vk, \text{nonce})$

    (6) $\sigma_i \xleftarrow{\$} \langle \mathcal{F}.\mathbf{Q}(\text{qrand}, \text{jrand}), [\![x]\!]_i \| \pi_i \rangle$

    (7) $\text{st}_i \xleftarrow{\$} \text{jrseed}$

    (8) return $(\sigma_i, \text{jrpart}_i), \text{st}_i$

$\mathcal{Z}.\text{Decide}[\mathcal{F}, \text{RO}](\sigma_0, \text{jrpart}_0, \sigma_1, \text{jrpart}_1, \text{st}_i)$

    (1) $\sigma \leftarrow \sigma_0 + \sigma_1$

    (2) $b \leftarrow \mathcal{F}.\text{Decide}(\sigma)$

    (3) if $(b \ \& \ [[\text{RO}(\text{jrseed}, \text{jrpart}_0 \| \text{jrpart}_1) = \text{st}_i]])$

    (4)   Return Accept

    (5) else return Reject

Figure 9: Construction of a shared Zero Knowledge proof system from a Fully Linear Proof system $\mathcal{F}$ and random oracles $\text{RO}_1 \dots \text{RO}_5$.

for whom the simulator must generate the malicious aggregator's report share, and malicious clients for whom the simulator receives only its own (maliciously generated) report share. The simulator processes these two types of reports separately in every oracle.

Let us first discuss maliciously generated client reports. The simulator receives these reports through the AcceptReport interface, and it simply stores them with a tag "mal" denoting that the share is malicious. In all other oracles, whenever the malicious tag is detected, the simulator honestly runs the Mastic protocol using its stored report share. Clearly, the behavior of the simulator and of an honest aggregator are identical for all malicious reports.

For honest client reports, the simulator's behavior is slightly more complex. In the GenReport interface, the simulator receives only an index i indicating that a new client report share should be generated. This report share must contain a nonce, a $\mathcal{V}$ public share and key, and a shared ZK partial proof, all corresponding to a measurement that the simulator does not know. Instead, the simulator picks its own measurement (α_i, β_i) at random from the set of all valid measurements, and creates a $\mathcal{V}$ public share and two keys by running the honest key generation algorithm. To hide the fact that these keys correspond to the wrong measurement, it immediately discards its own key key_{1-b}. It then uses the zero-knowledge simulator Sim_{SZK} to generate a simulated partial proof

and nonce (nonce, π_b^{szk}). Before it returns all these values to the adversary, it stores them with a tag "honest" indicating that the client is honest (and, by extension, that the simulator only knows the malicious aggregator's report share).

We must also simulate the validation of honest client reports, despite the lack of an honest report share. For shared ZK validation, the simulator refers again to the shared ZK zero-knowledge simulator to produce a verifier share and a decision. The remaining two portions of report validation are equality checks: the honest aggregator must produce a one-hot verifiability proof π for every prefix and level, and a path-verifiability proof h. Although we cannot compute these proofs directly, we know that for an honestly generated client report, the malicious and honest aggregators should always derive equal-valued proofs for both one-hot and path verifiability by the correctness of the VIDPF scheme. The simulator therefore uses its stored malicious report share to compute the malicious aggregator's proofs instead of the honest aggregator's, knowing they will be equal.

Formal Proof. The first change we make is to replace, one by one, the randomly sampled α_i, β_i in line 2 of GenReport. Let q be the number of queries $\mathcal{A}$ makes to the Honest Client Computation oracle. We design q hybrids, indexed by the integers j from 0 to q. In hybrid HYB_j(c.f. Figure 13), the first j queries to the Honest Client Computation oracle will compute pub, key_b by running GenReport honestly. In the remaining queries, it will instead compute $(\text{pub}, \text{key}_0, \text{key}_1) \coloneqq \mathcal{V}.\text{Gen}(\alpha, \beta)$, then discard the key key_{1-b}. HYB_0 is clearly identical to $\mathcal{G}_{\text{Sim}}^{\text{priv-ideal}}$, while HYB_q uses the honest client measurement to generate keys for every query. The distinction between HYB_j and HYB_{j+1} is that the $(j+1)^{\text{th}}$ Honest Client Computation query generates $(\text{pub}, \text{key}_b)$ from either the real measurement (α, β) or a randomly sampled measurement. Consequently, for each j we can build a reduction $\mathcal{D}$ that breaks the privacy of $\mathcal{V}$ with exactly the same probability that $\mathcal{A}$ distinguishes between HYB_j and HYB_{j+1}.

Our reduction $\mathcal{D}$ runs the hybrid HYB_j for $\mathcal{A}$. When the simulated "Honest Client Computation" oracle receives the $(j+1)^{\text{th}}$ query, it submits the client's measurement (α, β) as its own challenge along with the corrupt aggregator's index b. It uses the response pub, key_b in the corrupt aggregator's report share. If the challenge bit b in the $\mathcal{V}$ privacy game equals 0, the simulation of the j^{th} hybrid is perfect. Otherwise, $\mathcal{D}$ perfectly simulates the $(j+1)^{\text{th}}$ hybrid. We define $\mathcal{D}$ to return 1 when $\mathcal{A}$ returns 1 and return 0 otherwise. Then

$$\mathbf{Adv}_{\mathcal{V}}^{\text{priv}}(\mathcal{D}) \coloneqq \left| \Pr[\mathcal{G}_{\mathcal{V},1}^{\text{priv}}(\mathcal{D}) \Rightarrow 1] - \Pr[\mathcal{G}_{\mathcal{V},0}^{\text{priv}}(\mathcal{D}) \Rightarrow 1] \right| \leq$$

$$\left| \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{j+1}] - \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_j] \right|.$$

By a union bound over all $j \in [q]$, the probability of $\mathcal{A}$ distinguishing between hybrids HYB_0 and HYB_j is at most

$$\left| \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_q] - \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_0] \right| \leq q \cdot \mathbf{Adv}_{\mathcal{V}}^{\text{priv}}(\mathcal{D}).$$

In our next hybrid, HYB_{q+1}(c.f. Figure 14), we stop relying on the shared ZK simulator Sim_{SZK}. We replace all calls to Sim_{SZK} with the corresponding $\mathcal{Z}$ scheme operations. We first stop referring random oracle queries with $j = 2$ to the simulator and instead, sample a random function H_2. Then in GenReport we compute

$\mathcal{G}^{\text{priv-real}}_{\text{Mastic}}$

At the start of the game, the adversary outputs a key vk and the index b of the corrupt aggregator and the game samples $H_1 \xleftarrow{\$} \mathcal{AF}(\mathcal{V}.\text{Dom}, \mathcal{V}.\text{Rng})$, and $H_2 \xleftarrow{\$} \mathcal{AF}(\mathcal{Z}.\text{Dom}, \mathcal{Z}.\text{Rng})$, initializes lists T_{client} and T_{agg}, and sets $i := 0$.

Honest Client Computation:

Input: measurement $(\alpha, \beta) \in (\{0,1\}^n, \mathcal{L})$.

1. $i := i + 1$
2. Generate $(\text{nonce}, \text{pub}, \{\text{key}_c, \pi_c^{\text{szk}}\}_{c \in \{0,1\}})$ as in Fig. 2.
3. $T_{\text{client}}[i] := (\varepsilon, \text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$
4. $\text{st}_{(i,1-b)}^{\varepsilon} := \pi_{(i,1-b)}^{\varepsilon} := \varepsilon$
5. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Corrupt Client Computation:

Input: report share $(\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

1. $i := i + 1$
2. If $T_{\text{client}}[i] \neq \perp$ return $\perp$
3. $\text{st}_{(i,1-b)}^{\varepsilon} := \pi_{(i,1-b)}^{\varepsilon} := \varepsilon$
4. $T_{\text{client}}[i] := (\varepsilon, \text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

Aggregation

Input: set Reports of report indices and a prefix p

1. if $T_{\text{agg}}[p] \neq \perp$ then return $\perp$
2. else $T_{\text{agg}}[p] := \text{Reports}$; $[\![\text{weight}]\!]_{1-b} := 0$
3. for i in Reports
4. $[\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} + [\![y_i^p]\!]_{1-b}$
5. return $[\![\text{weight}]\!]_{1-b}$

Random Oracle:

Input: Index j and payload d. Outputs an independent random string (or vector)

(1) return $H_j(d)$

Honest Aggregator Evaluation:

Input: report index i and prefix p

1. $(\text{st}, \text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}) := T_{\text{client}}[i]$
2. if $\text{st} = \varepsilon$ then
3. $[\![y_i^\varepsilon]\!]_{1-b} := \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}_{1-b}, \text{pub})$
4. $(\text{st}^{\text{szk}}, \sigma) := \mathcal{Z}.\text{Query}^{H_2}(vk, \text{nonce}, [\![y_i^\varepsilon]\!]_{1-i}, \pi^{\text{szk}})_{1-b}$
5. $T_{\text{client}}[i] := ((\text{st}^{\text{szk}}, \sigma), \text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$
6. return σ
7. if $(\text{st}_{(i,1-b)}^p = \varepsilon$ or $[\![y_i^p]\!]_{1-b} = \varepsilon)$ and $p \neq \varepsilon$
8. return $\perp$
9. for $\gamma \in \{p \| 0, p \| 1\}$
10. $([\![y_i^\gamma]\!]_{1-b}, \text{st}_{(i,1-b)}^\gamma, \pi_{(i,1-b)}^\gamma) := \mathcal{V}.\text{Eval}^{H_1}(\text{key}_{1-b}, \text{pub}, \gamma, \text{st}_{(i,1-b)}^p, \pi_{(i,1-b)}^p)$
11. $h_{(i,1-b)}^p := (-1)^{1-b} \cdot ([\![y_i^p]\!]_{1-b} - [\![y_i^{p\|0}]\!]_{1-b} - [\![y_i^{p\|1}]\!]_{1-b})$
12. return $\pi_{(i,1-b)}^p, \pi_{(i,1-b)}^{p\|0}, \pi_{(i,1-b)}^{p\|1}, h_{(i,1-b)}^p$

Honest Aggregator Validation:

Input: report index i and partial shared ZK verifier σ_b

1. $(\text{st}, \text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := T_{\text{client}}[i]$
2. if $\text{st} \in \{\perp, \varepsilon\}$
3. return $\perp$
4. $(\text{st}^{\text{szk}}, \sigma_{1-i}) := \text{st}$
5. $T_{\text{client}}[i] \leftarrow (\perp, \text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}})$
6. return $\mathcal{Z}.\text{Decide}(\sigma_0, \sigma_1, \text{st}^{\text{szk}})$

Figure 10: Real game for defining privacy of Π_{Mastic}.

shares $[\![y_i^\epsilon]\!]_0$ and $[\![y_i^\epsilon]\!]_1$ of β. We can do this by using key_b to compute $[\![y_i^\epsilon]\!]_b$ via EvalRoot, then relying on the completeness of $\mathcal{V}$ and our knowledge of β to find $[\![y_i^\epsilon]\!]_{1-b}$. Then we call $\mathcal{Z}.\text{Prove}^{Hash_2}$ on these shares to generate the nonce and proofs. We store the shares and proofs, and in the Honest Aggregator Evaluation oracle, we compute the verifier share using $\mathcal{Z}.\text{Query}^{H_2}(vk, \text{nonce}, [\![y_i^\epsilon]\!]_{1-b}, \pi_{1-b}^{\text{szk}})$.

We then design a reduction $\mathcal{B}$ whose advantage in breaking the zero-knowledge security of $\mathcal{Z}$ is exactly the probability that $\mathcal{A}$ distinguishes hybrid HYB_{q+1} from HYB_q. Our reduction runs HYB_{q+1}, with a few changes to the highlighted lines. When HYB_{q+1} would query $\mathcal{Z}.\text{Prove}$ on the input shares in line 5 of the Honest Client Computation oracle, it instead calls its own **Prove** oracle to get back a nonce and proof. In line 9 of the Honest Aggregator Evaluation oracle, it queries **Query** on the honest client's nonce to get back a partial verifier string σ_{1-b}. Finally, in line 8 of the Honest Aggregator Validation oracle, the reduction queries its **Decide** oracle on the nonce and malicious verifier share instead of calling Decide. Additionally, everywhere HYB_{q+1} would call H_2, the reduction instead forwards the query to its own random oracle. Finally, $\mathcal{B}$ returns 1 if and only if $\mathcal{A}$ returns 1 in its simulated hybrid.

If we consider the behavior of the $\mathcal{B}$ when it plays $\mathcal{G}_{\mathcal{Z}}^{\text{shared ZK-real}}$, notice that like HYB_{q+1}, the random oracle implements a randomly sampled function from the correct set, and the Prove, Query,

and Decide oracles run $\mathcal{Z}$ on the proper inputs exactly as HYB_{q+1} would. Therefore the probability that $\mathcal{B}$ returns 1 is exactly $\Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{q+1}]$. Conversely, when $\mathcal{B}$ is playing $\mathcal{G}_{\text{Sim}_{\text{SZK}}}^{\text{shared ZK-ideal}}$, the random oracle and the Prove, Query, and Decide forward their inputs to Sim_{SZK} exactly as HYB_q would, and we have that $\Pr[\mathcal{B} \Rightarrow 1 | \mathcal{G}_{\text{Sim}_{\text{SZK}}}^{\text{shared ZK-ideal}}] = \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_q]$. Therefore,

$$\left| \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{q+1}] - \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_q] \right| \leq \mathbf{Adv}_{\mathcal{Z}, \text{Sim}_{\text{SZK}}}^{\text{priv}}(\mathcal{B}).$$

In our next hybrid, HYB_{q+2}(c.f. Figure 15), we are now ready to stop discarding the honest aggregator's VIDPF key key_{1-b}. Instead of computing the honest input share as $\beta - \mathcal{V}.\text{EvalRoot}^H(\text{key}_b, \text{pub})$, we start to compute it directly as $\mathcal{V}.\text{EvalRoot}^H(\text{key}_{1-b}, \text{pub})$. These values are equivalent due to the correctness of the VIDPF, so the adversary's view of the Honest Client Computation oracle does not change between HYB_{q+1} and HYB_{q+2}.

The next change is that we also use key_{1-b} in the Honest Aggregator Evaluation oracle to generate VIDPF proofs for honest reports. To enable this change, we now store key_{1-b} instead of key_b in $\text{st}[i]$ for all honest reports. This means that the output shares and proofs generated by $\mathcal{V}.\text{Eval}$ are now produced from the honest aggregator's key share instead of the malicious aggregator's key share. By $\mathcal{V}$ correctness and the definition of Valid as the equality function, however, the proofs π_i^p are equal for every prefix p

$\mathcal{G}_{\text{Sim}}^{\text{priv-ideal}}$

The simulator Sim is stateful and defines interfaces Init, GenReport, AcceptReport, Query, ValidateSZK, Aggregate, and RO. At the beginning of the game, the adversary outputs a key vk and the index b of the corrupt aggregator, and the game sets $i := 0$, initializes list T_{client}, and runs Sim.Init(vk, b).

Honest Client Computation:

Input: measurement $(\alpha, \beta) \in \{0, 1\}^n \times \mathcal{L}$

1. $i := i + 1$
2. $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}}) := \text{Sim.GenReport}(i)$.
3. $T_{\text{client}}[i] := (\alpha, \beta)$
4. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Corrupt Client Computation:

Input: report share $(\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

1. $i := i + 1$
2. Sim.AcceptReport$(i, \text{nonce}, \text{pub}, \pi^{\text{szk}}, \text{key})$

Random Oracle:

Input: Index j and payload d

1. return Sim.RO(j, d)

Honest Aggregator Evaluation:

Input: report index i and prefix p.

1. $(\text{pf}, \text{st}_{\text{Sim}_{\text{SZK}}}) := \text{Sim.Query}(i, p)$
   ▷ Depending on the state, pf may contain a verifier string σ or a tuple $(\pi_{(i,1-b)}^p, \pi_{(i,1-b)}^{p\|0}, \pi_{(i,1-b)}^{p\|1}, h_{(i,1-b)}^p)$
2. return pf

Honest Aggregator Validation:

Input: Report index i and partial sharedZK verifier string σ

1. return Sim.ValidateSZK(i, σ)

Aggregation

Input: set Reports of report indices, and a prefix p.

1. a := 0
2. for $i \in$ Reports
3. $(\alpha, \beta) := T_{\text{client}}[i]$
4.   if p is a prefix of α then a := a + β.
5. return Sim.Aggregate(p, N, a)

Figure 11: Ideal game for defining the privacy of Π_{Mastic}. Let Sim be a simulator.

regardless of which aggregator computes them, so the change in their derivation is undetectable.

What is left is to consider the path-verifiability check h output by Sim.Query. By $\mathcal{V}$ correctness, we have that for any $k \in [n]$ and any $p \in \{0, 1\}^k$, it holds that $[\![y_i^p]\!]_b = \beta - [\![y_i^p]\!]_{1-b}$ if p is a prefix of α and $[\![y_i^p]\!]_b = -[\![y_i^p]\!]_{1-b}$ otherwise. Consequently, we have that:

$$[\![y_i^p]\!]_b - [\![y_i^{p\|0}]\!]_b - [\![y_i^{p\|1}]\!]_b =$$
$$\begin{cases} \beta - [\![y_i^p]\!]_{1-b} - \beta + [\![y_i^{p\|0}]\!]_{1-b} + [\![y_i^{p\|1}]\!]_{1-b}, \\ \qquad\qquad\qquad \text{if } p \text{ is a prefix of } \alpha, \\ \\ -[\![y_i^p]\!]_{1-b} + [\![y_i^{p\|0}]\!]_{1-b} + [\![y_i^{p\|1}]\!]_{1-b}, \\ \qquad\qquad\qquad\qquad\qquad \text{otherwise.} \end{cases}$$

When β cancels out, we see that:

$$[\![y_i^p]\!]_b - [\![y_i^{p\|0}]\!]_b - [\![y_i^{p\|1}]\!]_b =$$
$$(-1)([\![y_i^p]\!]_{1-b} - [\![y_i^{p\|0}]\!]_{1-b} - [\![y_i^{p\|1}]\!]_{1-b}).$$

Accordingly, when we switch from key_b to key_{1-b} in the derivation of h, we must also multiply it by -1. With this tweak, the behavior of the Honest Aggregator Evaluation oracle in both hybrids is identical in the view of the adversary.

Now that we are storing key_{1-b} for all oracles, we must change the Aggregation oracle to maintain the consistency of the results. For each honest client reports $i \in$ Reports, in HYB_{q+1} we added the stored weight β to $[\![\text{weight}]\!]_{1-b}$ when p prefixed α. Then we subtracted the malicious aggregator's share $[\![y_i^p]\!]_b$ from $[\![\text{weight}]\!]_{1-b}$. Since we can no longer compute the malicious aggregator's share, in HYB_{q+2} we stop adding β and subtracting $[\![y_i^p]\!]_{1-b}$. Instead we

add the honest aggregator's share $[\![y_i^p]\!]_{1-b}$ to $[\![\text{weight}]\!]_{1-b}$. The correctness of $\mathcal{V}$ grants that $[\![y_i^p]\!]_{1-b} = \beta - [\![y_i^p]\!]_b$ when p prefixes α and $-[\![y_i^p]\!]_b$ otherwise, so the substitution we perform gives an identical value of $[\![\text{weight}]\!]_{1-b}$, and consequently

$$\Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{q+2}] = \Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{q+1}].$$

Notice that in HYB_{q+2}, we now treat both honest and malicious client reports identically regardless of the value of the tags ("honest" and "mal") in all oracles (except GenReport and AcceptReport, since we must generate the report ourselves in the former). Furthermore, the values computed and returned by each oracle in HYB_{q+2} are identical to those returned in $\mathcal{G}_{\text{Mastic}}^{\text{priv-real}}$, so

$$\Pr[\mathcal{A} \Rightarrow 1 | \text{HYB}_{q+2}] = \Pr[\mathcal{A} \Rightarrow 1 | \mathcal{G}_{\text{Mastic}}^{\text{priv-real}}].$$

Collecting bounds proves the claim.

D Proof of Theorem 2 (Π_{Mastic} is Robust)

In this section, we focus on the robustness guarantees provided by Mastic against malicious clients. To argue robustness against malicious clients, we assume the aggregators follow the Mastic protocol steps correctly. An adversary that maliciously corrupts multiple clients may attempt to disrupt the protocol by providing malformed report shares in Step. 3 of Fig. 2. A report share is considered to be malformed if 1) a client double-votes using the single report share, or 2) the report share contains an invalid measurement $(\alpha_i, \beta_i) \notin \{0, 1\}^n \times \mathcal{L}$. We argue that if the malformed report share passes the consistency checks and gets incorporated into the aggregation process by the honest aggregators then the malicious client breaks the verifiability of $\mathcal{V}$ in the first case, and soundness of $\mathcal{Z}$ in the second case.

$\text{Sim}_{\text{Sim}_{\text{SZK}}}(vk, b)$

The simulator Sim is stateful and defines interfaces Init, GenReport, AcceptReport, Query, ValidateSZK, Aggregate, and RO.

Sim.Init(vk, b)

1. $\text{st}[\mathcal{Z}] \leftarrow \epsilon$
2. $H_1 \xleftarrow{\$} \mathcal{AF}(\mathcal{V}.\text{Dom}, \mathcal{V}.\text{Rng})$
3. Store vk, b, and st as globally available state.

Sim.GenReport(i)

1. if $\text{st}[i] \neq \perp$ return $\perp$
2. $\alpha_i, \beta_i \xleftarrow{\$} \{0,1\}^n \times \mathcal{L}$
3. $(\text{pub}, \text{key}_0, \text{key}_1) \coloneqq \mathcal{V}.\text{Gen}(\alpha_i, \beta_i)$
4. $(\text{nonce}, \pi_b^{\text{szk}}, \text{st}[\mathcal{Z}]) \coloneqq \text{Sim}_{\text{SZK}}.\text{Prove}(\text{st}[\mathcal{Z}], b)$
5. $\text{st}[\mathcal{V}_i^\epsilon] \coloneqq \text{st}[\pi_i^\epsilon] \coloneqq \epsilon$
6. $\text{st}[i] \coloneqq (\text{``honest''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}}))$
7. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Sim.AcceptReport(i, nonce, pub, $\text{key}_{1-b}, \pi_{1-b}^{\text{szk}}$)

1. if $\text{st}[i] \neq \perp$ return $\perp$
2. $\text{st}[\mathcal{V}_i^\epsilon] \coloneqq \text{st}[\pi_i^\epsilon] \coloneqq \epsilon$
3. $\text{st}[i] \leftarrow (\text{``mal''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}))$

Sim.ValidateSZK(i, σ_b)

1. hon, status, params $\coloneqq \text{st}[i]$
2. if status $\in \{\epsilon, \perp\}$ return $\perp$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) \coloneqq \text{params}$
4. $\sigma_{1-b} \leftarrow \text{status}$
5. $\text{st}[i] \leftarrow (\text{hon}, \perp, \text{params})$
6. if hon = "mal"
7.   return $\mathcal{Z}.\text{Decide}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])$
8. return $\text{Sim}_{\text{SZK}}.\text{Decide}(\sigma_b, 1-b, \text{nonce}, \text{st}[\mathcal{Z}])$

Sim.RO(j, d)

1. if $j = 2$ then return $\text{Sim}_{\text{SZK}}.\text{RO}(d)$
2. else return $H_1(d)$

Sim.Query(i, p)

1. hon, status, params $\coloneqq \text{st}[i]$
2. if $(\text{st}[\mathcal{V}_i^p] = \epsilon$ or $\text{st}[\llbracket y_i^p \rrbracket] = \epsilon)$ and $p \neq \epsilon$ return $\perp$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) \coloneqq \text{params}$
4. if status $= \epsilon$
5. $\text{st}\llbracket y_i^\epsilon \rrbracket \coloneqq \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}, \text{pub})$
6.   if hon = "mal"
7. $(\text{st}[\mathcal{Z}, i], \sigma_{1-b}) \coloneqq \mathcal{Z}.\text{Query}(vk, \text{nonce}, \llbracket y_i^\epsilon \rrbracket, \pi^{\text{szk}})$
8.   else
9. $(\text{st}[\mathcal{Z}], \sigma_{1-b}) \coloneqq \text{Sim}_{\text{SZK}}.\text{Query}(vk, 1-b, \text{nonce}, \text{st}[\mathcal{Z}])$
10. $\text{st}[i] \leftarrow (\text{hon}, \sigma_{1-b}, \text{params})$
11.   return σ_{1-b}
12. for $\gamma \in \{p \parallel 0, p \parallel 1\}$
13. $(\text{st}[\llbracket y_i^\gamma \rrbracket], \text{st}[\mathcal{V}_i^\gamma], \text{st}[\pi_i^\gamma]) \coloneqq \mathcal{V}.\text{Eval}^{H_1}(\text{pub}, \text{key}, \gamma, \text{st}[\mathcal{V}_i^p], \text{st}[\pi_i^p])$
14. if hon = "mal"
15. $h \coloneqq (-1)^{1-b} \cdot (\text{st}[\llbracket y_i^p \rrbracket] - \text{st}[\llbracket y_i^{p\parallel 0} \rrbracket] - \text{st}[\llbracket y_i^{p\parallel 1} \rrbracket])$
16. else
17. $h \coloneqq (-1)^b \cdot (\text{st}[\llbracket y_i^p \rrbracket] - \text{st}[\llbracket y_i^{p\parallel 0} \rrbracket] - \text{st}[\llbracket y_i^{p\parallel 1} \rrbracket])$
18. return $\text{st}[\pi_i^p], \text{st}[\pi_i^{p\parallel 0}], \text{st}[\pi_i^{p\parallel 1}], h$

Sim.Aggregate($p, N, \mathbf{a}$)

1. $\llbracket \text{weight} \rrbracket_{1-b} \coloneqq \text{a}; k \coloneqq |p|$
2. for $i \in \text{Reports}$
3.   hon, status, params $\coloneqq \text{st}[i]$
4. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) \coloneqq \text{params}$
5.   if $(p|^{k-1} = \epsilon$ or $\text{st}[\mathcal{V}_i^{p|^{k-1}}] \neq \epsilon)$
6. $(\llbracket y_i^p \rrbracket, \_, \_) \coloneqq \mathcal{V}.\text{Eval}^{H_1}(\text{key}, \text{pub}, p, \text{st}[\mathcal{V}_i^{p|^{k-1}}], \text{st}[\pi_i^{p|^{k-1}}])$
7.   if hon = "mal"
8. $\llbracket \text{weight} \rrbracket_{1-b} \coloneqq \llbracket \text{weight} \rrbracket_{1-b} + \llbracket y_i^p \rrbracket$
9.   else
10. $\llbracket \text{weight} \rrbracket_{1-b} \coloneqq \llbracket \text{weight} \rrbracket_{1-b} - \llbracket y_i^p \rrbracket$
11. return $\llbracket \text{weight} \rrbracket_{1-b}$

Figure 12: Simulator Sim for the proof of Theorem 1. It takes as a parameter a simulator Sim_{SZK} for the privacy of $\mathcal{Z}$.

To argue robustness formally, for each pair of report shares submitted by a malicious client, the protocol needs to either "*extract*" a valid measurement from or detect that it is invalid. Once this distinction is performed the protocol needs to "*compute*" the aggregation function over the honest client *inputs* $(\alpha_i, \beta_i) \in \{0,1\}^n \times \mathcal{L}$ (represented via their report shares) and the valid measurements submitted by the malicious clients.

We capture this property in a simulation-based model [14], presented in Fig. 18 (Appendix D). The adversary $\mathcal{A}$ initially corrupts a set Reports′ of clients. In the real-world game $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ (Fig. 17), the parties run the Mastic protocol using their input measurements. Both the honest and corrupt clients provide their report shares to the aggregators, who compute the output (set of heavy-hitter strings and their children, and also the weights of the heavy-hitting strings and their children) and return it to $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$. The game forwards this to $\mathcal{A}$.

We also define a corresponding ideal-world game, $\mathcal{G}_{\text{Sim}, \mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$ (Fig. 18). In this game, the corrupt clients (controlled by $\mathcal{A}$) participate in the same way as in the real world, but with a pair of simulated aggregators. This simulation is performed by a PPT algorithm called Sim (Fig. 19). This simulator obtains the report shares of the corrupt clients, extracts the input measurements (α_i', β_i') from those shares, and invokes $\mathcal{F}_{\text{wHH}}$ (on behalf of the corrupt clients) with the extracted measurements.

Upon obtaining the client measurements from both honest and corrupt clients, $\mathcal{F}_{\text{wHH}}$ checks the measurements and then computes the output. For consistency check, $\mathcal{F}_{\text{wHH}}$ discards the measurements if they are not correctly formatted, i.e. $(\alpha_i, \beta_i) \neq (\perp, \perp)$, or if the weight is not valid, i.e. $(\alpha_i, \beta_i) \notin \mathcal{L}$. The functionality also allows corrupt clients to submit prefix strings, i.e. $\alpha_i \in \{0,1\}^{\leq n}$. Once the measurements are validated, the functionality aggregates the validated measurements by computing the weights and heavy-hitting set similar to the aggregation phase in Mastic. This is the Aggregation step in Fig. 16. $\mathcal{F}_{\text{wHH}}$ computes the output and this is returned to $\mathcal{G}_{\text{Sim}, \mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$. The game forwards this to $\mathcal{A}$.

Robustness guarantees that an adversary (who has corrupted an arbitrary number of clients) cannot distinguish between the real and ideal world outputs. We define the advantage of an adversary $\mathcal{A}$ against the robustness of Mastic with respect to simulator Sim

---

HYB$_j$

At the beginning of the game, the adversary outputs a key vk and the index b of the corrupt aggregator, and the game sets $i := 0$, initializes list T_{client}, and runs $\text{Sim.Init}(vk, b)$.

Honest Client Computation:

Input: measurement $(\alpha, \beta) \in \{0,1\}^n \times \mathcal{L}$

1. $i := i + 1$
2. $\alpha_i, \beta_i \xleftarrow{\$} \{0,1\}^n \times \mathcal{L}$
3. $\boxed{\text{if } i > j \text{ then } (\text{pub}, \text{key}_0, \text{key}_1) := \mathcal{V}.\text{Gen}(\alpha_i, \beta_i)}$
4. $\boxed{\text{else } (\text{pub}, \text{key}_0, \text{key}_1) := \mathcal{V}.\text{Gen}(\alpha, \beta)}$
5. $(\text{nonce}, \pi_b^{\text{szk}}, \text{st}[\mathcal{Z}]) := \text{Sim}_{\text{SZK}}.\text{Prove}(\text{st}[\mathcal{Z}, b])$
6. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
7. $\text{st}[i] := (\text{" honest"}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}}))$
8. $T_{\text{client}}[i] := (\alpha, \beta)$
9. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Corrupt Client Computation:

Input: report share $(\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

1. $i := i + 1$
2. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
3. $\text{st}[i] \leftarrow (\text{"mal"}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}))$

Aggregation

Input: set Reports of report indices, and a prefix p.

1. $\text{a} := 0$
2. for $i \in \text{Reports}$
3. $\quad (\alpha, \beta) := T_{\text{client}}[i]$
4. $\quad$ if p is a prefix of α then $\text{a} := \text{a} + \beta$.
5. $[\![\text{weight}]\!]_{1-b} := \text{a}; k := |p|$
6. for $i \in \text{Reports}$
7. $\quad \text{hon}, \text{status}, \text{params} := \text{st}[i]$
8. $\quad (\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
9. $\quad$ if $(p|^{k-1} = \epsilon$ or $\text{st}[\mathcal{V}_i^{p|^{k-1}}] \neq \epsilon)$
10. $\quad\quad ([\![y_i^p]\!], \_, \_) := \mathcal{V}.\text{Eval}^{H_1}(\text{key}, \text{pub}, p, \text{st}[\mathcal{V}_i^{p|^{k-1}}], \text{st}[\pi_i^{p|^{k-1}}])$
11. $\quad\quad$ if $\text{hon} = \text{"mal"}$
12. $\quad\quad\quad [\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} + [\![y_i^p]\!]$
13. $\quad\quad$ else
14. $\quad\quad\quad [\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} - [\![y_i^p]\!]$
15. return $[\![\text{weight}]\!]_{1-b}$

Honest Aggregator Evaluation:

Input: report index i and prefix p.

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $(\text{st}[\mathcal{V}_i^p] = \epsilon$ or $\text{st}[[\![y_i^p]\!]] = \epsilon)$ and $p \neq \epsilon$ return $\bot$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. if $\text{status} = \epsilon$
5. $\quad \text{st}[\![y_i^\epsilon]\!] := \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}, \text{pub})$
6. $\quad$ if $\text{hon} = \text{"mal"}$
7. $\quad\quad (\text{st}[\mathcal{Z}, i], \sigma_{1-b}) := \mathcal{Z}.\text{Query}^{\text{Sim}_{\text{SZK}}.\text{RO}}(vk, \text{nonce}, [\![y_i^\epsilon]\!], \pi^{\text{szk}})$
8. $\quad$ else
9. $\quad\quad (\text{st}[\mathcal{Z}], \sigma_{1-b}) := \text{Sim}_{\text{SZK}}.\text{Query}(vk, 1 - b, \text{nonce}, \text{st}[\mathcal{Z}])$
10. $\quad \text{st}[i] \leftarrow (\text{hon}, \sigma_{1-b}, \text{params})$
11. $\quad$ return σ_{1-b}
12. for $\gamma \in \{p \| 0, p \| 1\}$
13. $\quad (\text{st}[[\![y_i^\gamma]\!]], \text{st}[\mathcal{V}_i^\gamma], \text{st}[\pi_i^\gamma]) := \mathcal{V}.\text{Eval}^{H_1}(\text{pub}, \text{key}, \gamma, \text{st}[\mathcal{V}_i^p], \text{st}[\pi_i^p])$
14. if $\text{hon} = \text{"mal"}$
15. $\quad h := (-1)^{1-b} \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\|0}]\!]] - \text{st}[[\![y_i^{p\|1}]\!]])$
16. else
17. $\quad h := (-1)^b \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\|0}]\!]] - \text{st}[[\![y_i^{p\|1}]\!]])$
18. return $\text{st}[\pi_i^p], \text{st}[\pi_i^{p\|0}], \text{st}[\pi_i^{p\|1}], h$

Honest Aggregator Validation:

Input: Report index i and partial sharedZK verifier string σ

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $\text{status} \in \{\epsilon, \bot\}$ return $\bot$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. $\sigma_{1-b} \leftarrow \text{status}$
5. $\text{st}[i] \leftarrow (\text{hon}, \bot, \text{params})$
6. if $\text{hon} = \text{"mal"}$
7. $\quad$ return $\mathcal{Z}.\text{Decide}^{\text{Sim}_{\text{SZK}}.\text{RO}}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])$
8. return $\text{Sim}_{\text{SZK}}.\text{Decide}(\sigma_b, 1 - b, \text{nonce}, \text{st}[\mathcal{Z}])$

Random Oracle:

Input: Index j and payload d

1. if $j = 2$ then return $\text{Sim}_{\text{SZK}}.\text{RO}(d)$
2. else return $H_1(d)$

Figure 13: Hybrid HYB$_j$ in the proof of Theorem 1 . Changes from $\mathcal{G}_{\text{Sim}}^{\text{priv-ideal}}$ have been $\boxed{\text{highlighted}}$.

---

by $\textbf{Adv}_{\text{Mastic,Sim}}^{\text{rob}}(\mathcal{A})$ as:

$$\left| \Pr[\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}(\mathcal{A}) \Rightarrow 1] - \Pr[\mathcal{G}_{\text{Sim},\mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}(\mathcal{A}) \Rightarrow 1] \right|.$$

Mastic (also PLASMA and Poplar) permits a malicious client to submit report shares (in Mastic protocol) which are only valid until level $k \leq n$, after which robustness requires the report shares to be discarded.

Assume the adversary $\mathcal{A}$ corrupts $N' = |\text{Reports}'|$ clients. We formally introduce the games $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ and $\mathcal{G}_{\text{Sim},\mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$ for robustness in Figs. 17 and 18. We prove the robustness of Mastic in this section by proving that $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ is indistinguishable from $\mathcal{G}_{\text{Sim},\mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$ by providing the formal simulator algorithm Sim in Fig. 19.

Proof Sketch. To argue robustness, the aggregators need to ensure that the report shares provided by a corrupt client encode a valid measurement (α_i, β_i). As described in Section. 3.2.2, this involves checking three things.

The weight β_i encoded inside the report shares is valid, i.e. $\beta_i \in \mathcal{L}$. The aggregators ensure this by evaluating the $\mathcal{V}$ keys at the root layer to obtain β_i and then running the $\mathcal{Z}$ to validate β_i (without reconstructing β_i). An adversarial client who provides an invalid β_i, i.e. $\beta_i \notin \mathcal{L}$, but passes the checks of $\mathcal{Z}$ breaks soundness of $\mathcal{Z}$.

Next, the aggregators must ensure that the same β_i value is propagated across a single path in the evaluation tree (encoded inside the $\mathcal{V}$ keys). This reduces to verifying that each level $k(\in [n])$ in the evaluation tree contains only a single non-zero node. This is ensured by verifying the $\mathcal{V}$ proofs for every node considered (as

HYB_{q+1}

At the beginning of the game, the adversary outputs a key vk and the index b of the corrupt aggregator, and the game sets $i := 0$, initializes list T_{client}, and runs $\text{Sim.Init}(vk, b)$. We also sample $H_2 \overset{\$}{\leftarrow} \mathcal{AF}(\mathcal{Z}.\text{Dom}, \mathcal{Z}.\text{Rng})$.

Honest Client Computation:

Input: measurement $(\alpha, \beta) \in \{0,1\}^n \times \mathcal{L}$

1. $i := i + 1$
2. $(\text{pub}, \text{key}_0, \text{key}_1) := \mathcal{V}.\text{Gen}(\alpha, \beta)$
3. $[\![y_i^\epsilon]\!]_b := \text{EvalRoot}(\text{pub}, \text{key}_b)$
4. $[\![y_i^\epsilon]\!]_{1-b} := \beta - [\![y_i^\epsilon]\!]_b$
5. $(\text{nonce}, \pi_0^{\text{szk}}, \pi_1^{\text{szk}}) := \mathcal{Z}.\text{Prove}^{H_2}([\![y_i^\epsilon]\!]_0, [\![y_i^\epsilon]\!]_1)$
6. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
7. $\text{st}[i] := (\text{``honest''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}}))$
8. $T_{\text{client}}[i] := (\alpha, \beta)$
9. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Corrupt Client Computation:

Input: report share $(\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

1. $i := i + 1$
2. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
3. $\text{st}[i] \leftarrow (\text{``mal''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}))$

Aggregation

Input: set Reports of report indices, and a prefix p.

1. $\text{a} := 0$
2. for $i \in$ Reports
3. $\quad (\alpha, \beta) := T_{\text{client}}[i]$
4. $\quad$ if p is a prefix of α then $\text{a} := \text{a} + \beta$.
5. $[\![\text{weight}]\!]_{1-b} := \text{a}; \; k := |p|$
6. for $i \in$ Reports
7. $\quad \text{hon}, \text{status}, \text{params} := \text{st}[i]$
8. $\quad (\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
9. $\quad$ if $(p|^{k-1} = \epsilon$ or $\text{st}[\mathcal{V}_i^{p|^{k-1}}] \neq \epsilon)$
10. $\quad\quad ([\![y_i^p]\!], \_, \_) := \mathcal{V}.\text{Eval}^{H_1}(\text{key}, \text{pub}, p, \text{st}[\mathcal{V}_i^{p|^{k-1}}], \text{st}[\pi_i^{p|^{k-1}}])$
11. $\quad\quad$ if $\text{hon} = \text{``mal''}$
12. $\quad\quad\quad [\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} + [\![y_i^p]\!]$
13. $\quad\quad$ else
14. $\quad\quad\quad [\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} - [\![y_i^p]\!]$
15. return $[\![\text{weight}]\!]_{1-b}$

Honest Aggregator Evaluation:

Input: report index i and prefix p.

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $(\text{st}[\mathcal{V}_i^p] = \epsilon$ or $\text{st}[[\![y_i^p]\!]] = \epsilon)$ and $p \neq \epsilon$ return $\bot$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. if $\text{status} = \epsilon$
5. $\quad \text{st}[\![y_i^\epsilon]\!] := \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}, \text{pub})$
6. $\quad$ if $\text{hon} = \text{``mal''}$
7. $\quad\quad (\text{st}[\mathcal{Z}, i], \sigma_{1-b}) := \mathcal{Z}.\text{Query}^{H_2}(vk, \text{nonce}, [\![y_i^\epsilon]\!], \pi^{\text{szk}})$
8. $\quad$ else
9. $\quad\quad (\text{st}[\mathcal{Z}, i], \sigma_{1-b}) := \mathcal{Z}.\text{Query}(vk, \text{nonce}, [\![y_i^\epsilon]\!], \pi^{\text{szk}})$
10. $\quad \text{st}[i] \leftarrow (\text{hon}, \sigma_{1-b}, \text{params})$
11. $\quad$ return σ_{1-b}
12. for $\gamma \in \{p \| 0, p \| 1\}$
13. $\quad (\text{st}[[\![y_i^\gamma]\!]], \text{st}[\mathcal{V}_i^\gamma], \text{st}[\pi_i^\gamma]) := \mathcal{V}.\text{Eval}^{H_1}(\text{pub}, \text{key}, \gamma, \text{st}[\mathcal{V}_i^p], \text{st}[\pi_i^p])$
14. if $\text{hon} = \text{``mal''}$
15. $\quad h := (-1)^{1-b} \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\|0}]\!]] - \text{st}[[\![y_i^{p\|1}]\!]])$
16. else
17. $\quad h := (-1)^b \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\|0}]\!]] - \text{st}[[\![y_i^{p\|1}]\!]])$
18. return $\text{st}[\pi_i^p], \text{st}[\pi_i^{p\|0}], \text{st}[\pi_i^{p\|1}], h$

Honest Aggregator Validation:

Input: Report index i and partial sharedZK verifier string σ

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $\text{status} \in \{\epsilon, \bot\}$ return $\bot$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. $\sigma_{1-b} \leftarrow \text{status}$
5. $\text{st}[i] \leftarrow (\text{hon}, \bot, \text{params})$
6. if $\text{hon} = \text{``mal''}$
7. $\quad$ return $\mathcal{Z}.\text{Decide}^{H_2}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])$
8. return $\mathcal{Z}.\text{Decide}^{H_2}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])$

Random Oracle:

Input: Index j and payload d

1. if $j = 2$ then return $H_2(d)$
2. else return $H_1(d)$

Figure 14: Hybrid HYB_{q+1} in the proof of Theorem 1. Changes from HYB_q have been highlighted **.**

part of an evaluation path) during the computation of the heavy-hitter set. An adversarial client whose – report shares encode more than one non-zero node (as part of or two evaluation paths), and those nodes are also encountered during the heavy-hitter evaluation, and still the client passes the checks, can be used to break the verifiability of $\mathcal{V}$. Finally, the aggregator nodes must ensure that the non-zero nodes at each level are along the same path, say p. This is performed by checking that the output for prefix p is equal to the sum of the output of its children - ($p \| 0, p \| 1$). This check is information-theoretic. By combining the one-hot verifiability and path-verifiability guarantees, we provide stronger guarantees where the adversarial input $\alpha_i \in \{0,1\}^{\leq n}$ can be uniquely extracted from the report shares.

Finally, we optimized the communication by allowing the aggregator nodes to hash the results of the consistency checks for each client and then match the hash values. So, here we also need to rely on the collision-resistance of the hash function to ensure that if the hashes match then the underlying preimages are also equal.

Formal Proof. We prove Theorem 2 by showing that $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ and $\mathcal{G}_{\text{Sim},\mathcal{F}_{\text{wHH}}}^{\text{rob-ideal}}$ are computationally indistinguishable in the presence of our Sim against all PPT adversaries. We argue this via a sequence of hybrids.

- HYB_0: This is game $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$ the clients compute their report shares based on their inputs, provide these shares to the honest aggregators, the honest aggregators compute the

---

HYB$_{q+2}$

At the beginning of the game, the adversary outputs a key vk and the index b of the corrupt aggregator, and the game sets $i := 0$, initializes list T_{client}, and runs Sim.Init(vk, b). We also sample $H_2 \xleftarrow{\$} \mathcal{AF}(\mathcal{Z}.\text{Dom}, \mathcal{Z}.\text{Rng})$.

Honest Client Computation:

Input: measurement $(\alpha, \beta) \in \{0,1\}^n \times \mathcal{L}$

1. $i := i + 1$
2. $(\text{pub}, \text{key}_0, \text{key}_1) := \mathcal{V}.\text{Gen}(\alpha, \beta)$
3. $[\![y_i^\epsilon]\!]_b := \text{EvalRoot}(\text{pub}, \text{key}_b)$
4. $\boxed{[\![y_i^\epsilon]\!]_{1-b} := \text{EvalRoot}(\text{pub}, \text{key}_{1-b})}$
5. $(\text{nonce}, \pi_0^{\text{szk}}, \pi_1^{\text{szk}}) := \mathcal{Z}.\text{Prove}^{H_2}([\![y_i^\epsilon]\!]_0, [\![y_i^\epsilon]\!]_1)$
6. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
7. $\boxed{\text{st}[i] := (\text{``honest''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}))}$
8. $T_{\text{client}}[i] := (\alpha, \beta)$
9. Return $(\text{nonce}, \text{pub}, \text{key}_b, \pi_b^{\text{szk}})$

Corrupt Client Computation:

Input: report share $(\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}})$

1. $i := i + 1$
2. $\text{st}[\mathcal{V}_i^\epsilon] := \text{st}[\pi_i^\epsilon] := \epsilon$
3. $\text{st}[i] \leftarrow (\text{``mal''}, \epsilon, (\text{nonce}, \text{pub}, \text{key}_{1-b}, \pi_{1-b}^{\text{szk}}))$

Aggregation

Input: set Reports of report indices, and a prefix p.

1. $\boxed{[\![\text{weight}]\!]_{1-b} := 0}$; $k := |p|$
2. for $i \in$ Reports
3. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
4. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
5. if $(p^{|k-1} = \epsilon$ or $\text{st}[\mathcal{V}_i^{p^{|k-1}}] \neq \epsilon)$
6. $([\![y_i^p]\!], \_, \_) := \mathcal{V}.\text{Eval}^{H_1}(\text{key}, \text{pub}, p, \text{st}[\mathcal{V}_i^{p^{|k-1}}], \text{st}[\pi_i^{p^{|k-1}}])$
7.   if $\text{hon} = \text{``mal''}$
8. $[\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} + [\![y_i^p]\!]$
9.   else
10. $\boxed{[\![\text{weight}]\!]_{1-b} := [\![\text{weight}]\!]_{1-b} + [\![y_i^p]\!]}$
11. return $[\![\text{weight}]\!]_{1-b}$

Honest Aggregator Evaluation:

Input: report index i and prefix p.

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $(\text{st}[\mathcal{V}_i^p] = \epsilon$ or $\text{st}[[\![y_i^p]\!]] = \epsilon)$ and $p \neq \epsilon$ return $\perp$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. if $\text{status} = \epsilon$
5. $\text{st}[\![y_i^\epsilon]\!] := \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}, \text{pub})$
6.   if $\text{hon} = \text{``mal''}$
7. $(\text{st}[\mathcal{Z}, i], \sigma_{1-b}) := \mathcal{Z}.\text{Query}^{H_2}(vk, \text{nonce}, [\![y_i^\epsilon]\!], \pi^{\text{szk}})$
8.   else
9. $(\text{st}[\mathcal{Z}, i], \sigma_{1-b}) := \mathcal{Z}.\text{Query}(vk, \text{nonce}, [\![y_i^\epsilon]\!], \pi^{\text{szk}})$
10. $\text{st}[i] \leftarrow (\text{hon}, \sigma_{1-b}, \text{params})$
11.   return σ_{1-b}
12. for $\gamma \in \{p \parallel 0, p \parallel 1\}$
13. $(\text{st}[[\![y_i^\gamma]\!]], \text{st}[\mathcal{V}_i^\gamma], \text{st}[\pi_i^\gamma]) := \mathcal{V}.\text{Eval}^{H_1}(\text{pub}, \text{key}, \gamma, \text{st}[\mathcal{V}_i^p], \text{st}[\pi_i^p])$
14. if $\text{hon} = \text{``mal''}$
15. $h := (-1)^{1-b} \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\parallel 0}]\!]] - \text{st}[[\![y_i^{p\parallel 1}]\!]])$
16. else
17. $\boxed{h := (-1)^{b+1} \cdot (\text{st}[[\![y_i^p]\!]] - \text{st}[[\![y_i^{p\parallel 0}]\!]] - \text{st}[[\![y_i^{p\parallel 1}]\!]])}$
18. return $\text{st}[\pi_i^p], \text{st}[\pi_i^{p\parallel 0}], \text{st}[\pi_i^{p\parallel 1}], h$

Honest Aggregator Validation:

Input: Report index i and partial sharedZK verifier string σ

1. $\text{hon}, \text{status}, \text{params} := \text{st}[i]$
2. if $\text{status} \in \{\epsilon, \perp\}$ return $\perp$
3. $(\text{nonce}, \text{pub}, \text{key}, \pi^{\text{szk}}) := \text{params}$
4. $\sigma_{1-b} \leftarrow \text{status}$
5. $\text{st}[i] \leftarrow (\text{hon}, \perp, \text{params})$
6. if $\text{hon} = \text{``mal''}$
7.   return $\boxed{\mathcal{Z}.\text{Decide}^{H_2}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])}$
8. return $\boxed{\mathcal{Z}.\text{Decide}^{H_2}(\sigma_0, \sigma_1, \text{st}[\mathcal{Z}, i])}$

Random Oracle:

Input: Index j and payload d

1. if $j = 2$ then return $\boxed{H_2(d)}$
2. else return $H_1(d)$

Figure 15: Hybrid HYB$_{q+2}$ in the proof of Theorem 1 . Changes from HYB$_q$ have been $\boxed{\textbf{highlighted}}$.

Mastic protocol and then provide the output to $\mathcal{G}_{\text{Mastic}}^{\text{rob-real}}$. This output is returned to the adversary.

- HYB$_1$: This is same as HYB$_0$, except the simulated aggregators reconstruct β_i and removes i from Reports if $\beta_i \notin \mathcal{L}$. We present it in Fig. 20.

  A distinguisher distinguishes between the two hybrids if a malformed report containing $\beta_i \notin \mathcal{L}$ passes the consistency check for $\mathcal{Z}$ in HYB$_0$ and gets included in Reports, whereas in HYB$_1$ it gets removed from Reports. This alters the output distribution. Assuming a distinguisher $\mathcal{A}_{0,1}$ that distinguishes between the two hybrids then we build an adversary $\mathcal{B}$ that breaks the soundness of $\mathcal{Z}$ as follows.

When $\mathcal{A}_{0,1}$ (on behalf of a malicious client i) returns a malformed report share $\left(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi_{(i,b)}^{\text{szk}}\right)$ to $\mathcal{S}_b$ for each $b \in \{0,1\}$, our reduction $\mathcal{B}$ extracts

$$[\![\beta_i]\!]_b := \mathcal{V}.\text{EvalRoot}^{H_1}(\text{key}_{(i,b)}, \text{pub}_i)$$

for each $b \in \{0,1\}$. $\mathcal{B}$ then computes $\beta_i := [\![\beta_i]\!]_0 + [\![\beta_i]\!]_1$ and $(\text{st}_b, \sigma_b) := \mathcal{Z}.\text{Query}^{H_2}(vk, \text{nonce}_i, [\![\beta_i]\!]_b, \pi_{(i,b)}^{\text{szk}})$. If $\beta_i \notin \mathcal{L}$ and Accept $:= \mathcal{Z}.\text{Decide}^{H_2}(\sigma_0, \sigma_1, \text{st}_b)$, then $\mathcal{B}$ returns $(\text{nonce}_i, [\![\beta_i]\!]_0, [\![\beta_i]\!]_1, \pi_{(i,0)}^{\text{szk}}, \pi_{(i,1)}^{\text{szk}})$ to the challenger of $\mathcal{Z}$. It can be observed that $\mathcal{B}$ wins the game only when $\mathcal{A}_{0,1}$ constructs a malformed report such that the $\mathcal{Z}$ verifies but $\beta_i \notin \mathcal{L}$, and this directly translates into a win for $\mathcal{B}$. Assume the $\mathcal{A}_{0,1}$ has an advantage $\textbf{Adv}_{0,1}$ of distinguishing

$\mathcal{F}_{\mathsf{wHH}}$

PARAMETERS: Aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$. N clients C_i for $i \in [\mathsf{N}]$. $\mathcal{S}_0, \mathcal{S}_1$ agree on:

- A bound N on the number of client submissions.
- A bound T on the threshold for heavy hitters.
- Function order that defines a total ordering over sums of weights.

INPUTS:

Aggregators: $\mathcal{S}_0, \mathcal{S}_1$ do not have any input.

Clients: Each client C_i for $i \in [\mathsf{N}]$ holds "partial measurement" $(\alpha_i, \beta_i) \in (\{0,1\}^{\leq n} \cup \{\perp\}) \times (\mathbb{F}^m \cup \{\perp\})$ composed of an input α_i and its weight β_i. Each honest client holds a valid measurement, i.e. $(\alpha_i, \beta_i) \in \{0,1\}^n \times \mathcal{L}$. Let $\alpha_{i,k}$ represent the kth bit of α_i.

ALGORITHM:

1. *Init:* $\mathsf{HH}^{\leq n} := \{\mathsf{HH}^0, \mathsf{HH}^1, \ldots \mathsf{HH}^n\} := \{\{\epsilon\}, \emptyset, \ldots, \emptyset\}$. Set Reports $:= [\mathsf{N}]$.
2. *Check:* For each $i \in [\mathsf{N}]$: If $(\alpha_i, \beta_i) \notin \{0,1\}^{\leq n} \times \mathcal{L}$ then discard it from computation by updating Reports $:=$ Reports $\setminus \{i\}$.
3. *Aggregation:* For $k \in [0, \ldots, n-1]$ and for each prefix $p \in \mathsf{HH}^k$, consider $\gamma \in \{p \parallel 0, p \parallel 1\}$ and update $\mathsf{HH}^{k+1} := \mathsf{HH}^{k+1} \cup \gamma$ if $\mathsf{order}(\mathsf{weight}^\gamma) > \mathsf{T}$, where weight^γ is the sum of weights β_i of each input α_i with prefix γ. More formally:

$$\mathsf{weight}^\gamma := \sum_i \beta_i \text{ for } i \in \mathsf{Reports} \wedge (\alpha_i|^{k+1} = \gamma).$$

$\mathcal{F}_{\mathsf{wHH}}$ outputs the following:

- *Aggregators* $\mathcal{S}_0, \mathcal{S}_1$: Set of T-heavy hitters $\mathsf{HH}^{\leq n}$. For each heavy-hitting string p the aggregators also obtain weight^p, $\mathsf{weight}^{p\parallel 0}$ and $\mathsf{weight}^{p\parallel 1}$.
- *Clients* C_i for $i \in [\mathsf{N}]$: No output.

CORRUPTION: Adversary $\mathcal{A}_{\mathsf{HH}}$ corrupts multiple clients together. $\mathcal{A}_{\mathsf{HH}}$ sets the input of each corrupt client as a partial measurement (α_i, β_i).

NOTE: $\mathcal{F}_{\mathsf{wHH}}$ models the problem of "*weighted*" heavy-hitters. It can be weakened to capture "*plain*" heavy-hitters by letting *Check* enforce that $\beta_i = 1$ for every measurement.

Figure 16: Ideal Functionality for Weighted Heavy-Hitters

$\mathcal{G}_{\mathsf{Mastic}}^{\mathsf{rob-real}}$

Input: Each client C_i has input (α_i, β_i) for $i \in [\mathsf{N}]$.

1. Adversary $\mathcal{A}$ initially corrupts a set Reports' of clients and sends $\left(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi_{(i,b)}^{\mathsf{szk}}\right)$ to aggregator $\mathcal{S}_b$ for each $b \in \{0,1\}$ and $i \in \mathsf{Reports}'$.
2. Each aggregator $\mathcal{S}_b$ obtains $\left(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi_{(i,b)}^{\mathsf{szk}}\right)$ from ith honest client for $i \in \mathsf{Reports} \setminus \mathsf{Reports}'$, where $\mathsf{Reports} := [\mathsf{N}]$. Client i computes this by running the "Client Computation" protocol on $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$.
3. Each aggregator $\mathcal{S}_b$ runs the "Aggregation Computation" protocol (Fig. 2) on $\left(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi_{(i,b)}^{\mathsf{szk}}\right)$ for $b \in \{0,1\}$ to obtain $\mathsf{HH}^{\leq n}$, and for each heavy-hitting string $p \in \mathsf{HH}^{\leq n}$ obtain weight^p, $\mathsf{weight}^{p\parallel 0}$ and $\mathsf{weight}^{p\parallel 1}$.
4. Send the following to $\mathcal{A}$: Output $\mathsf{HH}^{\leq n}$, and weights weight^p, $\mathsf{weight}^{p\parallel 0}$ and $\mathsf{weight}^{p\parallel 1}$ for each $p \in \mathsf{HH}^{\leq n}$.

Figure 17: Real Game for defining robustness of Π_{Mastic}.

$\mathcal{G}_{\mathsf{Sim},\mathcal{F}_{\mathsf{wHH}}}^{\mathsf{rob-ideal}}$

Input: Each client C_i has input (α_i, β_i) for $i \in [\mathsf{N}]$.

1. Adversary $\mathcal{A}$ initially corrupts a set Reports' of clients and sends $\left(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi_{(i,b)}^{\mathsf{szk}}\right)$ to the simulated aggregators $\mathcal{S}_b$ (controlled by Sim) for each $b \in \{0,1\}$ and $i \in \mathsf{Reports}'$.
2. Sim extracts the ith malicious client inputs as follows for $i \in \mathsf{Reports}'$:

$$(i, \alpha_i', \beta_i') \leftarrow \mathsf{Sim}\left(\left(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi_{(i,b)}^{\mathsf{szk}}\right)_{b \in \{0,1\}}\right),$$

(α_i', β_i') is a *partial measurement* that is not necessarily valid. (E.g., $(\perp, \perp)$.) Let $R' := \{(i, \alpha_i', \beta_i') : i \in \mathsf{Reports}'\}$.
3. Invoke $\mathcal{F}_{\mathsf{wHH}}$ on each $(i, \alpha_i', \beta_i') \in R'$ (on behalf of each corrupt client $C_i \in \mathsf{Reports}'$ chosen by $\mathcal{A}_{\mathsf{HH}}$). Each honest client provides its input $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$ to $\mathcal{F}_{\mathsf{wHH}}$. $\mathcal{F}_{\mathsf{wHH}}$ returns the following outputs:
   - T-heavy hitters $\mathsf{HH}^{\leq n}$, and
   - For each heavy-hitting string $p \in \mathsf{HH}^{\leq n}$ $\mathcal{F}_{\mathsf{wHH}}$ also sends weight^p, $\mathsf{weight}^{p\parallel 0}$ and $\mathsf{weight}^{p\parallel 1}$.
4. Send the following to $\mathcal{A}$: Output $\mathsf{HH}^{\leq n}$, and weights weight^p, $\mathsf{weight}^{p\parallel 0}$ and $\mathsf{weight}^{p\parallel 1}$ for each $p \in \mathsf{HH}^{\leq n}$.

Figure 18: Ideal Game for defining robustness of Π_{Mastic}. Let $\mathcal{F}_{\mathsf{wHH}}$ be as defined in 16.

between the two hybrids. Then $\mathcal{B}$ wins the soundness game with at least $\mathbf{Adv}_{0,1}$ probability for a single client. The same attack should be considered for all N' clients. By applying the union bound over all N' and by using triangle inequality we bound $\mathbf{Adv}_{0,1}$ as follows:

$$\mathbf{Adv}_{0,1} \leq \sum_{i \in [\mathsf{N}']} \mathbf{Adv}_{0,1}^i \leq \mathsf{N}' \cdot \mathbf{Adv}_{\mathcal{Z}}^{\mathsf{sound}}(\mathcal{B}).$$

Note: Next, we consider a pair of $2n$ hybrids, each pair for a level. We first replace the hash functions with plain communication and then we rely on the verifiability of $\mathcal{V}$ to argue that each client provides a single non-zero path in the evaluation tree.

Run the following for $j \in [1, \ldots, n-1]$:

- HYB_{2j}: This is the same as HYB_{2j-1}, except the simulator runs both servers on the ith client's VIDPF keys on the empty string at level 0, extracts the y_i at the root level, checks that

it is valid by verifying that the path-verifiability and one-hot verifiability checks verify at level 1. If the checks verify then there will be a single path at level 1 which will be non-zero. The simulator considers that path as the active path $\alpha_{i,j}$. The same process is iteratively repeated by the simulator until $j-1$ level to extract the first j bits of α_i as $\alpha_{i,\leq j}$ for every client. Once all the client's inputs are extracted until the first j, the simulator runs $\mathcal{F}_{\mathsf{wHH}}$ on them to obtain the set of jth-length heavy-hitting prefixes. Then the simulator runs Mastic starting from jth level. The simulated $\mathcal{S}_b$ accumulates its local state as $R_{(i,b)}^j := \left(\big\|_{p \in \mathsf{HH}^j} (p, h_{(i,b)}^p, \pi_{(i,b)}^{p\parallel 0}, \pi_{(i,b)}^{p\parallel 1})\right)$ instead

---

Sim for $\mathcal{G}^{\text{rob-ideal}}_{\text{Sim},\mathcal{F}_{\text{wHH}}}$

Sim simulates the role of the honest aggregators in this protocol.

Malicious Client Computation: Each corrupt client C_i for $i \in$ Reports' sends report share $(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi^{\text{szk}}_{(i,b)})$ to the simulated aggregators $\mathcal{S}_b$ for each $b \in \{0,1\}$.

Simulated Aggregator Computation:

Input: The aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ are run by Sim. They start with a verification key $vk \in \{0,1\}^{\text{vkl}}$ established out-of-band. They receive the set of corrupt clients as Reports'. Each aggregator $\mathcal{S}_b$ obtains $(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi^{\text{szk}}_{(i,b)})$ for $i \in$ Reports'.

1. **For each** client $i \in$ Reports:
   a. If $(\text{nonce}_i, \text{pub}_i, \text{key}_{(i,b)}, \pi^{\text{szk}}_{(i,b)})$ does not follow the correct input-formatting of Mastic then Sim sets $(\alpha'_i, \beta'_i) := (\perp, \perp)$ as the ith client's measurement and skips rest of this loop for this particular value of i. ▸ Input-Formatting check.
   b. Otherwise, Sim computes $\beta'_i := [\![\beta_i]\!]_0 + [\![\beta_i]\!]_1$, where $[\![\beta_i]\!]_b := \mathcal{V}.\text{EvalRoot}^{\text{H}_1}(\text{key}_{(i,b)}, \text{pub}_i)$ for $b \in \{0,1\}$.
      ▸ Weight-computation. $\mathcal{F}_{\text{wHH}}$ performs weight-check on this β'_i in Step 2.
   c. Sim extracts the α'_i as follows. Initialize $r := \epsilon$ and $\alpha'_i := \epsilon$. For $b \in \{0,1\}$, Sim sets $[\![y^r_i]\!]_b := \epsilon, \text{st}^r_{(i,b)} := \epsilon$, $\pi^r_{(i,b)} := \epsilon$ and store them in memory. For $k \in [0,1,\ldots,n-1]$ run the following:
      i. Each $\mathcal{S}_b$ retrieves the state $([\![y^r_i]\!]_b, \text{st}^r_{(i,b)}, \pi^r_{(i,b)})$ from memory corresponding to prefix r.
      ii. Each $\mathcal{S}_b$ runs $([\![y^\gamma_i]\!]_b, \text{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)}) := \mathcal{V}.\text{Eval}^{\text{H}_1}(\text{key}_{(i,b)}, \text{pub}_i, \gamma, \text{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)})$ for each prefix $\gamma \in \{r \| 0, r \| 1\}$ and stores the results in memory.
      iii. For $\gamma \in \{r \| 0, r \| 1\}$, Sim computes $y^\gamma_i := [\![y^\gamma_i]\!]_0 + [\![y^\gamma_i]\!]_1$.
      iv. If any of the three conditions hold then Sim considers (α'_i, β'_i) as the ith client's measurement and skips this inner and the outer loop for this particular value of i:
      - Both $\pi^{r\|0}_{(i,0)} \neq \pi^{r\|0}_{(i,1)} \wedge \pi^{r\|1}_{(i,0)} \neq \pi^{r\|1}_{(i,1)}$,
      - Both $y^{r\|0}_i \neq 0$ and $y^{r\|1}_i \neq 0$, or　　　　　　　　　　　　　　　　　　　　▸ One-hot verifiability.
      - $y^{r\|0}_i + y^{r\|1}_i \neq \beta'_i$.　　　　　　　　　　　　　　　　　　　　　　　　　▸ Path verifiability.
      v. If $(y^{r\|0}_i = \beta'_i)$ then update $r := r \| 0$ and $\alpha'_i := \alpha'_i \| 0$. Otherwise, if $(y^{r\|1}_i = \beta'_i)$ then update $r := r \| 1$ and $\alpha'_i := \alpha'_i \| 1$
   d. After the above loop, Sim stores the ith client's measurement as (i, α'_i, β'_i).
2. Sim returns all the corrupt client measurements as (i, α'_i, β'_i) for $i \in$ Reports'.

Figure 19: Simulator for ideal robustness game $\mathcal{G}^{\text{rob-ideal}}_{\text{Sim},\mathcal{F}_{\text{wHH}}}$.

of $R^j_{(i,b)} := \text{H}(\|_{p \in \text{HH}^j} (p, h^p_{(i,b)}, \pi^{p\|0}_{(i,b)}, \pi^{p\|1}_{(i,b)}))$. The checks are performed over $R^j_{(i,b)}$. We present it in Figs. 21, 22.

An adversary $\mathcal{A}_{2j-1,2j}$ distinguishes between the two hybrids if it finds a collision in H s.t. the verification in HYB_{2j} fails due to $R^j_{(i,0)} \neq R^j_{(i,1)}$ whereas they match in HYB_{2j-1} since to a collision in the hash. Using this adversary, we construct an adversary C that finds a collision in the hash by returning $(\|_{p \in \text{HH}^j} (p, h^p_{(i,0)}, \pi^{p\|0}_{(i,0)}, \pi^{p\|1}_{(i,0)}))$ and $(\|_{p \in \text{HH}^j} (p, h^p_{(i,1)}, \pi^{p\|0}_{(i,b)}, \pi^{p\|1}_{(i,1)}))$. These two values are unequal but their hash values match.

If the advantage of $\mathcal{A}_{2j-1,2j}$ is $\mathbf{Adv}_{2j-1,2j}$ then C wins the game with advantage at least $\mathbf{Adv}_{2j-1,2j}$. The same attack should be considered for all N' clients. By applying the union bound over all N' and by using triangle inequality we bound $\mathbf{Adv}_{2j-1,2j}$ as follows:

$$\mathbf{Adv}_{2j-1,2j} \leq \sum_{i \in [\text{N'}]} \mathbf{Adv}^i_{2j-1,2j} \leq \text{N'} \cdot \mathbf{Adv}^{\text{coll}}_{\text{H}}(C).$$

- HYB_{2j+1} : This is same as HYB_{2j}, except the servers extract $\alpha_{i,j+1}$ by running the extraction algorithm until level j and then the heavy-hitter set $\text{HH}^{\leq j+1}$ is computed based on the extracted values, whereas $\widetilde{\text{HH}^{j+1}}$ is computed following the Mastic protocol; and if $\text{HH}^{j+1} \neq \widetilde{\text{HH}^{j+1}}$, then Sim return

$(\perp, \perp)$ to $\mathcal{F}_{\text{wHH}}$ on behalf of the corrupt clients. We present it in Figures. 23, 24.

An adversary $\mathcal{A}_{2j,2j+1}$ distinguishes between the two hybrids if it finds two non-zero evaluation paths in the $\mathcal{V}$ evaluation at level j such that they lead to two different non-zero values on prefixes $r \| 0$ and $r \| 1$. In such a case, HYB_{2j} fails to detect it and HYB_{2j+1} detects it. HYB_{2j} will consider $\widetilde{\text{HH}^j}$ as the heavy-hitting set and HYB_{2j+1} will return $(\perp, \perp)$ to $\mathcal{F}_{\text{wHH}}$. Using this adversary, we construct an adversary $\mathcal{D}$ that breaks the verifiability of the $\mathcal{V}$. When client i returns the report shares simulated as per HYB_{2j} and if it encounters two such non-zero evaluation paths - $(r \| 0, r \| 1)$ on which the proofs verify, then return the $\mathcal{V}$ keys in the report shares and $(r \| 0, r \| 1)$ to the $\mathcal{V}$ challenger. Since the proofs are verified, it means that in Step 2d the $\mathcal{V}$ keys evaluate to two values $y^{r\|0}_i$ and $y^{r\|1}_i$ such that either $y^{r\|0}_i \neq y^{r\|1}_i \neq 0$ or $y^{r\|0}_i + y^{r\|1}_i \neq \beta_i$. The second case cannot occur in HYB_{2j} or HYB_{2j+1} since the servers verify that by sending $h^p_{(i,b)}$ in the clear. So it must be the case that $y^{r\|0}_i \neq y^{r\|1}_i \neq 0$ and the corresponding proofs verified. Thus, the $\mathcal{V}$ keys sent by $\mathcal{A}_{2j,2j+1}$ can be used to break the verifiability property of $\mathcal{V}$ where $\mathcal{D}$ returns the $\mathcal{V}$ keys as $(\text{key}_0, \text{key}_1)$ and $(r \| 0, r \| 1)$ as (u, v) respectively in the verifiability game.

Hybrid HYB₁

Sim simulates the role of the honest aggregators in this protocol.

Primitives:

1. A VIDPF ($\mathcal{V}$.Gen, $\mathcal{V}$.Eval, $\mathcal{V}$.EvalRoot) as defined in Section 2.3.
2. A shared ZK ($\mathcal{Z}$.Prove, $\mathcal{Z}$.Query, $\mathcal{Z}$.Decide) for a language $\mathcal{L} \subseteq \mathbb{F}^m$ as defined in Section 2.4.
3. Functions H, H₁, H₂ modeled in our analysis as random oracles.

Client Computation:

Input: Each client C_i for $i \in [\mathsf{N}]$ holds measurement $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$ composed of an input α_i and its weight β_i.

1. C_i runs $(\mathsf{pub}_i, \mathsf{key}_{(i,0)}, \mathsf{key}_{(i,1)}) \xleftarrow{\$} \mathcal{V}.\mathsf{Gen}^{H_1}(\alpha_i, \beta_i)$.
2. C_i runs $(\pi^{\mathsf{szk}}_{(i,0)}, \pi^{\mathsf{szk}}_{(i,1)}, \mathsf{nonce}_i) \xleftarrow{\$} \mathcal{Z}.\mathsf{Prove}^{H_2}(\llbracket \beta_i \rrbracket_0, \llbracket \beta_i \rrbracket_1)$ where $\llbracket \beta_i \rrbracket_b := \mathcal{V}.\mathsf{EvalRoot}^{H_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i)$ for $b \in \{0,1\}$.
3. C_i sends report share $(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi^{\mathsf{szk}}_{(i,b)})$ to $\mathcal{S}_b$ for each $b \in \{0,1\}$.

Aggregator Computation:

Input: The simulated aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ start with a verification key $vk \in \{0,1\}^{\mathsf{vkl}}$ established out-of-band. Each sets $\mathsf{HH}^{\leq n} := \{\mathsf{HH}^0, \mathsf{HH}^1, \ldots \mathsf{HH}^n\}$ as $\{\epsilon, \emptyset, \ldots, \emptyset\}$, the initial set of *candidate prefixes* for each level and sets $\mathsf{Reports} := [\mathsf{N}]$, the initial set of *candidate reports*.

1. **For each** client $i \in \mathsf{Reports}$: ▷ Weight check using $\mathcal{Z}$ at the root.
   a. Remove i from Reports if $(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi^{\mathsf{szk}}_{(i,b)})$ does not follow the correct formatting ▷ Input-Formatting check.
   b. Otherwise, Sim computes $\beta'_i := \llbracket \beta_i \rrbracket_0 + \llbracket \beta_i \rrbracket_1$, where $\llbracket \beta_i \rrbracket_b := \mathcal{V}.\mathsf{EvalRoot}^{H_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i)$ for $b \in \{0,1\}$.
   Sim removes i from Reports if $\beta'_i \notin \mathcal{L}$. ▷ Weight-check.
2. **For each** level $k \in [0, \ldots, n-1]$ and prefix $p \in \mathsf{HH}^k$:
   a. **For each** candidate report $i \in \mathsf{Reports}$: ▷ Path & One-hot Verifiability checks.
      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p \rrbracket_b, \mathsf{st}^p_{(i,b)}, \pi^p_{(i,b)})$ from memory corresponding to prefix p and client C_i.
      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma \rrbracket_b, \mathsf{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)}) := \mathcal{V}.\mathsf{Eval}^{H_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}^p_{(i,b)}, \pi^p_{(i,b)})$ for each prefix $\gamma \in \{p \| 0, p \| 1\}$ and stores the results in memory.
      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \| 0$ and $p \| 1$. To do so, each $\mathcal{S}_b$ computes $h^p_{(i,b)} := (-1)^b \cdot (\llbracket y_i^p \rrbracket_b - \llbracket y_i^{p\|0} \rrbracket_b - \llbracket y_i^{p\|1} \rrbracket_b)$. ▷ Observe that $h^p_{(i,0)} = h^p_{(i,1)}$.
      iv. $\mathcal{S}_b$ accumulates its local state as $R^k_{(i,b)} := \mathsf{H}(\|_{p \in \mathsf{HH}^k}(p, h^p_{(i,b)}, \pi^{p\|0}_{(i,b)}, \pi^{p\|1}_{(i,b)}))$. ▷ This is for all heavy-hitter prefixes.
      v. $\mathcal{S}_b$ sends $R^k_{(i,b)}$ to $\mathcal{S}_{1-b}$. If $R^k_{(i,0)} \neq R^k_{(i,1)}$, then $\mathcal{S}_b$ removes i from Reports. ▷ One hash for each client.
   b. **For each** k-bit heavy-hitting prefix $p \in \mathsf{HH}^k$ the aggregators prune on $\gamma \in \{p \| 0, p \| 1\}$ as: ▷ Aggregation & Pruning.
      i. Each $\mathcal{S}_b$ accumulates $\llbracket \mathsf{weight}^\gamma \rrbracket_b := \llbracket \mathsf{weight}^\gamma \rrbracket_b + \llbracket y_i^\gamma \rrbracket_b$. ▷ Each $\llbracket y_i^\gamma \rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.
      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\mathsf{weight}^\gamma := \llbracket \mathsf{weight}^\gamma \rrbracket_0 + \llbracket \mathsf{weight}^\gamma \rrbracket_1$. If $\mathsf{order}(\mathsf{weight}^\gamma) < \mathsf{T}$, then prune γ from the candidate prefix set. Otherwise, accumulate $\mathsf{HH}^{k+1} := \mathsf{HH}^{k+1} \cup \{\gamma\}$. ▷ $\mathsf{order}(\cdot)$ is decided by the aggregators.
3. Finally, the servers output HH^n as the set of weighted T-heavy-hitters.

Figure 20: Hybrid HYB₁ for the Robustness Proof. Changes from HYB₀ are highlighted.

If the advantage of $\mathcal{A}_{2j,2j+1}$ is $\mathbf{Adv}_{2j,2j+1}$ then $\mathcal{D}$ wins the game with advantage at least $\mathbf{Adv}_{2j,2j+1}$. The same attack should be considered for all N' clients. By applying the union bound over all N' and by using the triangular inequality, we bound $\mathbf{Adv}_{2j,2j+1}$ as follows:

$$\mathbf{Adv}_{2j,2j+1} \leq \sum_{i \in [\mathsf{N}']} \mathbf{Adv}^i_{2j,2j+1} \leq \mathsf{N}' \cdot \mathbf{Adv}^{\mathsf{verif}}_{\mathcal{V},k}(\mathcal{D}).$$

$\mathcal{G}^{\mathsf{rob\text{-}ideal}}_{\mathsf{Sim}, \mathcal{F}_{\mathsf{wHH}}}$ as follows:

$$\mathbf{Adv}^{\mathsf{rob}}_{\mathsf{Mastic,Sim}, \mathcal{F}_{\mathsf{wHH}}}(\mathcal{A}) =$$

$$\mathbf{Adv}_{0,1} + \sum_{j \in [1, \ldots n-1]} (\mathbf{Adv}_{2j-1,2j} + \mathbf{Adv}_{2j,2j+1}) \leq$$

$$\mathsf{N}' \cdot \mathbf{Adv}^{\mathsf{sound}}_{\mathcal{Z}}(\mathcal{B}) + \sum_{j \in [1, \ldots n-1]} (\mathsf{N} \cdot (\mathbf{Adv}^{\mathsf{coll}}_{\mathsf{H}}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{verif}}_{\mathcal{V},k}(\mathcal{D}))),$$

which is equal to

$$\mathsf{N}' \cdot (\mathbf{Adv}^{\mathsf{sound}}_{\mathcal{Z}}(\mathcal{B}) + n \cdot (\mathbf{Adv}^{\mathsf{coll}}_{\mathsf{H}}(\mathcal{C}) + \mathbf{Adv}^{\mathsf{verif}}_{\mathcal{V},k}(\mathcal{D}))).$$

This concludes the robustness proof.

It can be observed that HYB_{2n} corresponds to $\mathcal{G}^{\mathsf{rob\text{-}ideal}}_{\mathsf{Sim}, \mathcal{F}_{\mathsf{wHH}}}$. Thus, we bound the advantage of $\mathcal{A}$ distinguishing between $\mathcal{G}^{\mathsf{rob\text{-}real}}_{\mathsf{Mastic}}$ and

---

Hybrid HYB$_{2j}$

Sim simulates the role of the honest aggregators in this protocol.

Primitives:

1. A VIDPF ($\mathcal{V}$.Gen, $\mathcal{V}$.Eval, $\mathcal{V}$.EvalRoot) as defined in Section 2.3.
2. A shared ZK ($\mathcal{Z}$.Prove, $\mathcal{Z}$.Query, $\mathcal{Z}$.Decide) for a language $\mathcal{L} \subseteq \mathbb{F}^m$ as defined in Section 2.4.
3. Functions H, H$_1$, H$_2$ modeled in our analysis as random oracles.

Client Computation:

Input: Each client C_i for $i \in [\mathsf{N}]$ holds measurement $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$ composed of an input α_i and its weight β_i.

1. C_i runs $(\mathrm{pub}_i, \mathrm{key}_{(i,0)}, \mathrm{key}_{(i,1)}) \stackrel{\$}{\leftarrow} \mathcal{V}.\mathrm{Gen}^{\mathrm{H}_1}(\alpha_i, \beta_i)$.
2. C_i runs $(\pi^{\mathrm{szk}}_{(i,0)}, \pi^{\mathrm{szk}}_{(i,1)}, \mathrm{nonce}_i) \stackrel{\$}{\leftarrow} \mathcal{Z}.\mathrm{Prove}^{\mathrm{H}_2}(\llbracket \beta_i \rrbracket_0, \llbracket \beta_i \rrbracket_1)$ where $\llbracket \beta_i \rrbracket_b \coloneqq \mathcal{V}.\mathrm{EvalRoot}^{\mathrm{H}_1}(\mathrm{key}_{(i,b)}, \mathrm{pub}_i)$ for $b \in \{0,1\}$.
3. C_i sends report share $(\mathrm{nonce}_i, \mathrm{pub}_i, \mathrm{key}_{(i,b)}, \pi^{\mathrm{szk}}_{(i,b)})$ to $\mathcal{S}_b$ for each $b \in \{0,1\}$.

Aggregator Computation:

Input: The simulated aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ start with a verification key $vk \in \{0,1\}^{\mathrm{vkl}}$ established out-of-band. Each sets $\mathrm{HH}^{\leq n} \coloneqq \{\mathrm{HH}^0, \mathrm{HH}^1, \ldots \mathrm{HH}^n\}$ as $\{\epsilon, \emptyset, \ldots, \emptyset\}$, the initial set of *candidate prefixes* for each level and sets Reports $\coloneqq [\mathsf{N}]$, the initial set of *candidate reports*.

1. **For each** client $i \in$ Reports: ▷ Weight check using $\mathcal{Z}$ at the root.
   a. Remove i from Reports if $(\mathrm{nonce}_i, \mathrm{pub}_i, \mathrm{key}_{(i,b)}, \pi^{\mathrm{szk}}_{(i,b)})$ does not follow the correct formatting.   ▷ Input-Formatting check.
   b. Otherwise, Sim computes $\beta_i \coloneqq \llbracket \beta_i \rrbracket_0 + \llbracket \beta_i \rrbracket_1$, where $\llbracket \beta_i \rrbracket_b \coloneqq \mathcal{V}.\mathrm{EvalRoot}^{\mathrm{H}_1}(\mathrm{key}_{(i,b)}, \mathrm{pub}_i)$ for $b \in \{0,1\}$. Sim removes i from Reports if $\beta_i \notin \mathcal{L}$

                                                          ▷ Weight-check.

2. **For each** client $i \in$ Reports: Sim extracts the $\alpha_{i,\leq j-1}$ as follows. Initialize $r \coloneqq \epsilon$ and $\alpha_i \coloneqq \epsilon$. For $b \in \{0,1\}$, Sim sets $\llbracket y_i^r \rrbracket_b \coloneqq \epsilon, \mathrm{st}^r_{(i,b)} \coloneqq \epsilon$, $\pi^r_{(i,b)} \coloneqq \epsilon$ and store them in memory. **For** $k \in [0, 1, \ldots, j-1]$ run the following:
   a. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^r \rrbracket_b, \mathrm{st}^r_{(i,b)}, \pi^r_{(i,b)})$ from memory corresponding to prefix r.
   b. Each $\mathcal{S}_b$ runs $(\llbracket y_i^\gamma \rrbracket_b, \mathrm{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)}) \coloneqq \mathcal{V}.\mathrm{Eval}^{\mathrm{H}_1}(\mathrm{key}_{(i,b)}, \mathrm{pub}_i, \gamma, \mathrm{st}^r_{(i,b)}, \pi^r_{(i,b)})$ for each prefix $\gamma \in \{r \,\|\, 0, r \,\|\, 1\}$ and stores the results in memory.
   c. For $\gamma \in \{r \,\|\, 0, r \,\|\, 1\}$, Sim computes $y_i^\gamma \coloneqq \llbracket y_i^\gamma \rrbracket_0 + \llbracket y_i^\gamma \rrbracket_1$.
   d. If any of the following three conditions hold then Sim considers (α_i, β_i) as the ith client's measurement and skips this inner and the outer loop for this particular value of i:
      i. Both $\pi^{r\|0}_{(i,0)} \neq \pi^{r\|0}_{(i,1)} \wedge \pi^{r\|1}_{(i,0)} \neq \pi^{r\|1}_{(i,1)}$, or   ▷ One-hot verifiability.
      ii. Both $y_i^{r\|0} \neq 0$ and $y_i^{r\|1} \neq 0$, or   ▷ One-hot verifiability.
      iii. $y_i^{r\|0} + y_i^{r\|1} \neq \beta_i$.   ▷ Path verifiability.
      (If any of the above conditions hold then the client's input is valid up to r and invalid for both $(r \,\|\, 0, r \,\|\, 1)$.)
   e. If $(y_i^{r\|0} = \beta_i)$ then update $r \coloneqq r \,\|\, 0$ and $\alpha_i \coloneqq \alpha_i \,\|\, 0$. Otherwise, if $(y_i^{r\|1} = \beta_i)$ then update $r \coloneqq r \,\|\, 1$ and $\alpha_i \coloneqq \alpha_i \,\|\, 1$

3. Compute $\mathrm{HH}^{\leq j}$ as follows: For $k \in [0, \ldots, j-1]$ and for each prefix $p \in \mathrm{HH}^k$, consider $\gamma \in \{p \,\|\, 0, p \,\|\, 1\}$ and update $\mathrm{HH}^{k+1} \coloneqq \mathrm{HH}^{k+1} \cup \gamma$ if $\mathrm{order}(\mathrm{weight}^\gamma) > \mathsf{T}$, where weight^γ is sum of weights β_i of each input α_i with prefix γ. More formally:

$$\mathrm{weight}^\gamma \coloneqq \sum_i \beta_i \text{ for } i \in [\mathrm{Reports}] \wedge (\alpha_{i,\leq k+1} = \gamma).$$

Figure 21: Hybrid HYB$_{2j}$ for the Robustness Proof (Cont. in Fig. 22). Changes from HYB$_{2j-1}$ are $\boxed{\text{highlighted}}$.

Hybrid HYB_{2j}

4. At the jth level - For each prefix $p \in \mathsf{HH}^j$:
   a. **For each** candidate report $i \in$ Reports:　　　　　　　　　　　　　　　▷ Path & One-hot Verifiability checks.
      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p \rrbracket_b, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ from memory corresponding to prefix p and client C_i.
      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma \rrbracket_b, \mathsf{st}_{(i,b)}^\gamma, \pi_{(i,b)}^\gamma) \coloneqq \mathcal{V}.\mathsf{Eval}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ for each prefix $\gamma \in \{p \parallel 0, p \parallel 1\}$ and stores the results in memory.
      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \parallel 0$ and $p \parallel 1$. To do so, each $\mathcal{S}_b$ computes $h_{(i,b)}^p \coloneqq (-1)^b \cdot (\llbracket y_i^p \rrbracket_b - \llbracket y_i^{p\parallel 0} \rrbracket_b - \llbracket y_i^{p\parallel 1} \rrbracket_b)$.　　　　　　　▷ Observe that $h_{(i,0)}^p = h_{(i,1)}^p$.
      iv. $\boxed{\mathcal{S}_b \text{ accumulates its local state as } R_{(i,b)}^j \coloneqq (\parallel_{p \in \mathsf{HH}^j} (p, h_{(i,b)}^p, \pi_{(i,b)}^{p\parallel 0}, \pi_{(i,b)}^{p\parallel 1})).}$　　　▷ This is for all heavy-hitter prefixes.
      v. $\mathcal{S}_b$ sends $R_{(i,b)}^j$ to $\mathcal{S}_{1-b}$. If $R_{(i,0)}^k \neq R_{(i,1)}^k$, then $\mathcal{S}_b$ removes i from Reports.　　　　　▷ One hash for each client.
   b. **For each** j-bit heavy-hitting prefix $p \in \mathsf{HH}^j$ the aggregators prune on $\gamma \in \{p \parallel 0, p \parallel 1\}$ as:　　▷ Aggregation & Pruning.
      i. Each $\mathcal{S}_b$ accumulates $\llbracket \mathsf{weight}^\gamma \rrbracket_b \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_b + \llbracket y_i^\gamma \rrbracket_b$.　　　　　▷ Each $\llbracket y_i^\gamma \rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.
      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\mathsf{weight}^\gamma \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_0 + \llbracket \mathsf{weight}^\gamma \rrbracket_1$. If $\mathsf{order}(\mathsf{weight}^\gamma) < \mathsf{T}$, then prune γ from the candidate prefix set. Otherwise, accumulate $\mathsf{HH}^{j+1} \coloneqq \mathsf{HH}^{j+1} \cup \{\gamma\}$.

5. **For each** level $k \in [j+1, \dots, n-1]$ and prefix $p \in \mathsf{HH}^k$:
   a. **For each** candidate report $i \in$ Reports:　　　　　　　　　　　　　　　▷ Path & One-hot Verifiability checks.
      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p \rrbracket_b, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ from memory corresponding to prefix p and client C_i.
      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma \rrbracket_b, \mathsf{st}_{(i,b)}^\gamma, \pi_{(i,b)}^\gamma) \coloneqq \mathcal{V}.\mathsf{Eval}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ for each prefix $\gamma \in \{p \parallel 0, p \parallel 1\}$ and stores the results in memory.
      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \parallel 0$ and $p \parallel 1$. To do so, each $\mathcal{S}_b$ computes $h_{(i,b)}^p \coloneqq (-1)^b \cdot (\llbracket y_i^p \rrbracket_b - \llbracket y_i^{p\parallel 0} \rrbracket_b - \llbracket y_i^{p\parallel 1} \rrbracket_b)$.　　　　▷ Observe that $h_{(i,0)}^p = h_{(i,1)}^p$.
      iv. $\mathcal{S}_b$ accumulates its local state as $R_{(i,b)}^k \coloneqq \mathsf{H}(\parallel_{p \in \mathsf{HH}^k} (p, h_{(i,b)}^p, \pi_{(i,b)}^{p\parallel 0}, \pi_{(i,b)}^{p\parallel 1}))$.　　　▷ This is for all heavy-hitter prefixes.
      v. $\mathcal{S}_b$ sends $R_{(i,b)}^k$ to $\mathcal{S}_{1-b}$.
      vi. $\mathcal{S}_b$ computes $d' \coloneqq R_{(i,0)}^k = R_{(i,1)}^k$. If $d' = False$, then remove i from Reports.
   b. **For each** k-bit heavy-hitting prefix $p \in \mathsf{HH}^k$ the aggregators prune on $\gamma \in \{p \parallel 0, p \parallel 1\}$ as:　　▷ Aggregation & Pruning.
      i. Each $\mathcal{S}_b$ accumulates $\llbracket \mathsf{weight}^\gamma \rrbracket_b \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_b + \llbracket y_i^\gamma \rrbracket_b$.　　　　　▷ Each $\llbracket y_i^\gamma \rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.
      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\mathsf{weight}^\gamma \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_0 + \llbracket \mathsf{weight}^\gamma \rrbracket_1$. If $\mathsf{order}(\mathsf{weight}^\gamma) < \mathsf{T}$, then prune γ from the candidate prefix set. Otherwise, accumulate $\mathsf{HH}^{k+1} \coloneqq \mathsf{HH}^{k+1} \cup \{\gamma\}$.　　　　　　　　　　　　　▷ $\mathsf{order}(\cdot)$ is decided by the aggregators.

6. Finally, the servers output HH^n as the set of weighted T-heavy-hitters.

Figure 22: Hybrid HYB_{2j} for the Robustness Proof (Cont. from Fig. 21).

Hybrid HYB_{2j+1}

Sim simulates the role of the honest aggregators in this protocol.

Primitives:

1. A VIDPF ($\mathcal{V}$.Gen, $\mathcal{V}$.Eval, $\mathcal{V}$.EvalRoot) as defined in Section 2.3.
2. A shared ZK ($\mathcal{Z}$.Prove, $\mathcal{Z}$.Query, $\mathcal{Z}$.Decide) for a language $\mathcal{L} \subseteq \mathbb{F}^m$ as defined in Section 2.4.
3. Functions H, H_1, H_2 modeled in our analysis as random oracles.

Client Computation:

Input: Each client C_i for $i \in [\mathsf{N}]$ holds measurement $(\alpha_i, \beta_i) \in (\{0,1\}^n, \mathcal{L})$ composed of an input α_i and its weight β_i.

1. C_i runs $(\mathsf{pub}_i, \mathsf{key}_{(i,0)}, \mathsf{key}_{(i,1)}) \xleftarrow{\$} \mathcal{V}.\mathsf{Gen}^{\mathsf{H}_1}(\alpha_i, \beta_i)$.
2. C_i runs $(\pi^{\mathsf{szk}}_{(i,0)}, \pi^{\mathsf{szk}}_{(i,1)}, \mathsf{nonce}_i) \xleftarrow{\$} \mathcal{Z}.\mathsf{Prove}^{\mathsf{H}_2}(\llbracket\beta_i\rrbracket_0, \llbracket\beta_i\rrbracket_1)$ where $\llbracket\beta_i\rrbracket_b := \mathcal{V}.\mathsf{EvalRoot}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i)$ for $b \in \{0,1\}$.
3. C_i sends report share $(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi^{\mathsf{szk}}_{(i,b)})$ to $\mathcal{S}_b$ for each $b \in \{0,1\}$.

Aggregator Computation:

Input: The simulated aggregators $\mathcal{S}_0$ and $\mathcal{S}_1$ start with a verification key $vk \in \{0,1\}^{\mathsf{vkl}}$ established out-of-band. Each sets $\mathsf{HH}^{\leq n} := \{\mathsf{HH}^0, \mathsf{HH}^1, \ldots \mathsf{HH}^n\}$ as $\{\epsilon, \emptyset, \ldots, \emptyset\}$, the initial set of *candidate prefixes* for each level and sets $\mathsf{Reports} := [\mathsf{N}]$, the initial set of *candidate reports*.

1. **For each** client $i \in \mathsf{Reports}$:   ▷ Weight check using $\mathcal{Z}$ at the root.
   a. Remove i from Reports if $(\mathsf{nonce}_i, \mathsf{pub}_i, \mathsf{key}_{(i,b)}, \pi^{\mathsf{szk}}_{(i,b)})$ does not follow the correct formatting.   ▷ Input-Formatting check.
   b. Otherwise, Sim computes $\beta_i := \llbracket\beta_i\rrbracket_0 + \llbracket\beta_i\rrbracket_1$, where $\llbracket\beta_i\rrbracket_b := \mathcal{V}.\mathsf{EvalRoot}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i)$ for $b \in \{0,1\}$. Sim removes i from Reports if $\beta_i \notin \mathcal{L}$
                                                                                                   ▷ Weight-check.
2. **For each** client $i \in \mathsf{Reports}$: Sim extracts the $\boxed{\alpha_{i,\leq j}}$ as follows. Initialize $r := \epsilon$ and $\alpha_i := \epsilon$. For $b \in \{0,1\}$, Sim sets $\llbracket y_i^r\rrbracket_b := \epsilon, \mathsf{st}^r_{(i,b)} := \epsilon$, $\pi^r_{(i,b)} := \epsilon$ and store them in memory. **For** $\boxed{k \in [0, 1, \ldots, j]}$ run the following:
   a. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^r\rrbracket_b, \mathsf{st}^r_{(i,b)}, \pi^r_{(i,b)})$ from memory corresponding to prefix r.
   b. Each $\mathcal{S}_b$ runs $(\llbracket y_i^\gamma\rrbracket_b, \mathsf{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)}) := \mathcal{V}.\mathsf{Eval}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}^\gamma_{(i,b)}, \pi^\gamma_{(i,b)})$ for each prefix $\gamma \in \{r \parallel 0, r \parallel 1\}$ and stores the results in memory.
   c. For $\gamma \in \{r \parallel 0, r \parallel 1\}$, Sim computes $y_i^\gamma := \llbracket y_i^\gamma\rrbracket_0 + \llbracket y_i^\gamma\rrbracket_1$.
   d. If any of the following three conditions hold then Sim considers (α_i, β_i) as the ith client's measurement and skips this inner and the outer loop for this particular value of i:
      i. Both $\pi^{r\parallel0}_{(i,0)} \neq \pi^{r\parallel0}_{(i,1)} \wedge \pi^{r\parallel1}_{(i,0)} \neq \pi^{r\parallel1}_{(i,1)}$, or   ▷ One-hot verifiability.
      ii. Both $y_i^{r\parallel0} \neq 0$ and $y_i^{r\parallel1} \neq 0$, or   ▷ One-hot verifiability.
      iii. $y_i^{r\parallel0} + y_i^{r\parallel1} \neq \beta_i$.   ▷ Path verifiability.
      (If any of the above conditions hold then the client's input is valid up to r and invalid for both $(r \parallel 0, r \parallel 1)$.)
   e. If $(y_i^{r\parallel0} = \beta_i)$ then update $r := r \parallel 0$ and $\alpha_i := \alpha_i \parallel 0$. Otherwise, if $(y_i^{r\parallel1} = \beta_i)$ then update $r := r \parallel 1$ and $\alpha_i := \alpha_i \parallel 1$
3. $\boxed{\text{Compute } \mathsf{HH}^{\leq j+1} \text{ as follows: For } k \in [0, \ldots, j]}$ and for each prefix $p \in \mathsf{HH}^k$, consider $\gamma \in \{p \parallel 0, p \parallel 1\}$ and update $\mathsf{HH}^{k+1} := \mathsf{HH}^{k+1} \cup \gamma$ if $\mathsf{order}(\mathsf{weight}^\gamma) > \mathsf{T}$, where weight^γ is sum of weights β_i of each input α_i with prefix γ. More formally:

$$\mathsf{weight}^\gamma := \sum_i \beta_i \text{ for } i \in [\mathsf{Reports}] \wedge (\alpha_{i,\leq k+1} = \gamma).$$

Figure 23: Hybrid HYB_{2j+1} for the Robustness Proof (Cont. in Fig. 24). Changes from HYB_{2j} are $\boxed{\text{highlighted}}$.

Hybrid HYB_{2j+1}

4. At the jth level - For each prefix $p \in \mathsf{HH}^j$:
   a. **For each** candidate report $i \in \mathsf{Reports}$:   ▷ Path & One-hot Verifiability checks.
      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p \rrbracket_b, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ from memory corresponding to prefix p and client C_i.
      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma \rrbracket_b, \mathsf{st}_{(i,b)}^\gamma, \pi_{(i,b)}^\gamma) \coloneqq \mathcal{V}.\mathsf{Eval}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ for each prefix $\gamma \in \{p \parallel 0, p \parallel 1\}$ and stores the results in memory.
      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \parallel 0$ and $p \parallel 1$. To do so, each $\mathcal{S}_b$ computes $h_{(i,b)}^p \coloneqq (-1)^b \cdot (\llbracket y_i^p \rrbracket_b - \llbracket y_i^{p\parallel 0} \rrbracket_b - \llbracket y_i^{p\parallel 1} \rrbracket_b)$.   ▷ Observe that $h_{(i,0)}^p = h_{(i,1)}^p$.
      iv. $\mathcal{S}_b$ accumulates its local state as $R_{(i,b)}^j \coloneqq (\big\Vert_{p \in \mathsf{HH}^j} (p, h_{(i,b)}^p, \pi_{(i,b)}^{p\parallel 0}, \pi_{(i,b)}^{p\parallel 1}))$.   ▷ This is for all heavy-hitter prefixes.
      v. $\mathcal{S}_b$ sends $R_{(i,b)}^j$ to $\mathcal{S}_{1-b}$. If $R_{(i,0)}^j \neq R_{(i,1)}^j$, then $\mathcal{S}_b$ removes i from $\mathsf{Reports}$.
      vi. $\boxed{\text{For } \gamma \in \{p \parallel 0, p \parallel 1\}\text{: simulated } \mathcal{S}_0 \text{ and } \mathcal{S}_1 \text{ reconstruct } y_i^\gamma \coloneqq \llbracket y_i^\gamma \rrbracket_0 + \llbracket y_i^\gamma \rrbracket_1.}$
   b. **For each** j-bit heavy-hitting prefix $p \in \mathsf{HH}^j$ the aggregators prune on $\gamma \in \{p \parallel 0, p \parallel 1\}$ as: ▷ Aggregation & Pruning.
      i. Each $\mathcal{S}_b$ accumulates $\llbracket \mathsf{weight}^\gamma \rrbracket_b \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_b + \llbracket y_i^\gamma \rrbracket_b$.   ▷ Each $\llbracket y_i^\gamma \rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.
      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\mathsf{weight}^\gamma \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_0 + \llbracket \mathsf{weight}^\gamma \rrbracket_1$. If $\mathsf{order}(\mathsf{weight}^\gamma) < \mathsf{T}$, then prune γ from the candidate prefix set.

      $\boxed{\text{Otherwise, accumulate } \widetilde{\mathsf{HH}^{j+1}} \coloneqq \widetilde{\mathsf{HH}^{j+1}} \cup \{\gamma\}.}$

      $\boxed{\text{If } \mathsf{HH}^{j+1} \neq \widetilde{\mathsf{HH}^{j+1}} \text{ then return } (\bot, \bot) \text{ to } \mathcal{F}_{\mathsf{wHH}} \text{ for all corrupt clients.}}$
5. **For each** level $k \in [j+1, \ldots, n-1]$ and prefix $p \in \mathsf{HH}^k$:
   a. **For each** candidate report $i \in \mathsf{Reports}$:   ▷ Path & One-hot Verifiability checks.
      i. Each $\mathcal{S}_b$ retrieves the state $(\llbracket y_i^p \rrbracket_b, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ from memory corresponding to prefix p and client C_i.
      ii. Each $\mathcal{S}_b$ runs as $(\llbracket y_i^\gamma \rrbracket_b, \mathsf{st}_{(i,b)}^\gamma, \pi_{(i,b)}^\gamma) \coloneqq \mathcal{V}.\mathsf{Eval}^{\mathsf{H}_1}(\mathsf{key}_{(i,b)}, \mathsf{pub}_i, \gamma, \mathsf{st}_{(i,b)}^p, \pi_{(i,b)}^p)$ for each prefix $\gamma \in \{p \parallel 0, p \parallel 1\}$ and stores the results in memory.
      iii. The aggregators check that the output for prefix p is equal to the sum of the outputs on prefixes $p \parallel 0$ and $p \parallel 1$. To do so, each $\mathcal{S}_b$ computes $h_{(i,b)}^p \coloneqq (-1)^b \cdot (\llbracket y_i^p \rrbracket_b - \llbracket y_i^{p\parallel 0} \rrbracket_b - \llbracket y_i^{p\parallel 1} \rrbracket_b)$.   ▷ Observe that $h_{(i,0)}^p = h_{(i,1)}^p$.
      iv. $\mathcal{S}_b$ accumulates its local state as $R_{(i,b)}^k \coloneqq \mathsf{H}\big(\big\Vert_{p \in \mathsf{HH}^k} (p, h_{(i,b)}^p, \pi_{(i,b)}^{p\parallel 0}, \pi_{(i,b)}^{p\parallel 1})\big)$.   ▷ This is for all heavy-hitter prefixes.
      v. $\mathcal{S}_b$ sends $R_{(i,b)}^k$ to $\mathcal{S}_{1-b}$. If $R_{(i,0)}^k \neq R_{(i,1)}^k$, then $\mathcal{S}_b$ removes i from $\mathsf{Reports}$.   ▷ One hash for each client.
   b. **For each** k-bit heavy-hitting prefix $p \in \mathsf{HH}^k$ the aggregators prune on $\gamma \in \{p \parallel 0, p \parallel 1\}$ as: ▷ Aggregation & Pruning.
      i. Each $\mathcal{S}_b$ accumulates $\llbracket \mathsf{weight}^\gamma \rrbracket_b \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_b + \llbracket y_i^\gamma \rrbracket_b$.   ▷ Each $\llbracket y_i^\gamma \rrbracket_b$ is a vector of field elements $\mathbb{F}^m$.
      ii. $\mathcal{S}_0$ and $\mathcal{S}_1$ recover $\mathsf{weight}^\gamma \coloneqq \llbracket \mathsf{weight}^\gamma \rrbracket_0 + \llbracket \mathsf{weight}^\gamma \rrbracket_1$. If $\mathsf{order}(\mathsf{weight}^\gamma) < \mathsf{T}$, then prune γ from the candidate prefix set. Otherwise, accumulate $\mathsf{HH}^{k+1} \coloneqq \mathsf{HH}^{k+1} \cup \{\gamma\}$.   ▷ $\mathsf{order}(\cdot)$ is decided by the aggregators.
6. Finally, the servers output HH^n as the set of weighted T-heavy-hitters.

Figure 24: Hybrid HYB_{2j+1} for the Robustness Proof (Cont. from Fig. 23).