

ChameleMon: Shifting Measurement Attention as Network State Changes

Kaicheng Yang[†], Yuhan Wu[†], Ruijie Miao[†], Tong Yang[†], Zirui Liu[†], Zicang Xu[†]
Rui Qiu[†], Yikai Zhao[†], Hanglong Lv[†], Zhigang Ji[¶], Gaogang Xie[§]

[†]National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

[¶]Huawei Technologies Co., Ltd. [§]CNIC CAS, UCAS

ABSTRACT

*Network measurement is critical to many network applications. There are mainly two kinds of flow-level measurement tasks: 1) packet accumulation tasks and 2) packet loss tasks. In practice, the two kinds of tasks are often required at the same time, but existing works seldom handle both. In this paper, we design ChameleMon to support the two kinds of tasks simultaneously. The key design of ChameleMon is to shift measurement attention as network state changes, through two dimensions of dynamics: 1) dynamically allocating memory between the two kinds of tasks; 2) dynamically monitoring the flows of importance. To realize the key design, we propose a key technique, leveraging Fermat's little theorem to devise a flexible data structure, namely FermatSketch. FermatSketch is dividable, additive, and subtractive, supporting the two kinds of tasks. We have implemented a ChameleMon prototype on a testbed with a Fat-tree topology. We conduct extensive experiments and the results show ChameleMon supports the two kinds of tasks with low memory/bandwidth overhead, and more importantly, it can automatically shift measurement attention as network state changes.

CCS CONCEPTS

• **Networks** → **Data path algorithms; Network measurement; Programmable networks; Network monitoring.**

KEYWORDS

Sketch; Programmable Switch; Packet Loss; Measurement Attention

ACM Reference Format:

Kaicheng Yang[†], Yuhan Wu[†], Ruijie Miao[†], Tong Yang[†], Zirui Liu[†], Zicang Xu[†], Rui Qiu[†], Yikai Zhao[†], Hanglong Lv[†], Zhigang Ji[¶], Gaogang Xie[§]. 2023. **ChameleMon: Shifting Measurement Attention as Network State Changes**. In *ACM SIGCOMM 2023 Conference (ACM SIGCOMM '23)*, September 10–14, 2023, New York, NY, USA. ACM, New York, NY, USA, 23 pages. <https://doi.org/10.1145/3603269.3604850>

*Kaicheng Yang, Yuhan Wu, and Ruijie Miao contribute equally to this paper. Tong Yang (yangtongemail@gmail.com) is the corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
ACM SIGCOMM '23, September 10–14, 2023, New York, NY, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0236-5/23/09...\$15.00
<https://doi.org/10.1145/3603269.3604850>

1 INTRODUCTION

Network measurement provides critical statistics for various network operations, such as traffic engineering [1, 2], congestion control [3], network accounting [4], anomaly detection [5–8], and failure troubleshooting [9, 10]. In earlier years, sampling-based solutions [11–14] were widely accepted thanks to their simplicity and ease of use. Recently, sketch-based solutions [15–17] have attracted much more attention than sampling-based ones, as they are designed to approach the ultimate goal of network measurement [18–20]: to support more tasks and achieve higher accuracy with less memory. The emerging programmable switches that can process packets at terabit line rate further make sketches practical for production deployment, and designing novel sketches for flow-level measurement capabilities on programmable switches has become a hot topic [18–20].

There are mainly two kinds of flow-level measurement tasks. The first kind is *packet accumulation tasks* that focus on flow sizes at certain network nodes, including flow size estimation [21], heavy-hitter detection [22], entropy estimation [23], etc.. The second kind is *packet loss tasks* that focus on changes of flow sizes between network nodes, among which the most representative one is packet loss detection [24]. However, the two kinds of tasks are seldom considered and supported simultaneously in one solution. One reason behind is that the two kinds of tasks require very different flow-level statistics.

However, in practice, the two kinds of tasks are often required at the same time, and there are only limited resources for measurement in programmable switches (e.g., O(10MB) SRAM and limited accesses to the SRAM). Therefore, the *first requirement* for a practical measurement system is versatile to support the two kinds of tasks with high accuracy using limited resources, where limited resources refer to sub-linear space complexity.

Based on the first requirement, the *second requirement* is to pay attention to different kinds of tasks for different network states. When the network state is healthy and there are only few packet losses in the network, the system should pay more attention (e.g., allocate more memory) to packet accumulation tasks. When the network state is ill and there are lots of packet losses in the network, the system should pay more attention to packet loss tasks to help diagnose network faults, especially for those flows which experience a great number of packet losses.

In summary, a practical measurement system should meet the following requirements: **[R_{1.1}]** *versatility requirement*: supporting both packet loss tasks and packet accumulation tasks simultaneously; **[R_{1.2}]** *efficiency requirement*: achieving high accuracy with sub-linear space complexity; **[R₂]** *attention requirement*: paying attention to different kinds of tasks for different network states.

Existing solutions can be mainly classified into three categories according to supported measurement tasks:

- (1) *Solutions for packet loss tasks*: Typical solutions including LossRadar [24] based on Invertible Bloom filter [25], Netseer [26] and Dapper [27] based on the advanced features of programmable switches, and more. These solutions are often carefully designed to only obtain the exact difference set of flows/packets to minimize measurement overhead, while packet accumulation tasks require approximate sizes of all flows or simply large flows. Therefore, these solutions can hardly be extended to packet accumulation tasks and fail to meet $[R_{1.1}]$.
- (2) *Solutions for packet accumulation tasks*: These solutions are usually based on sketches, including CM sketch [21], UnivMon [18], ElasticSketch [15], HashPipe [22], and more. To efficiently maintain approximate flow sizes, these solutions choose to embrace hash collisions and select the estimation with least collisions to minimize error. For these solutions, due to their inherent error caused by hash collisions, it is difficult to obtain the exact difference set of flows/packets. Therefore, these solutions can hardly be extended to packet loss tasks and fail to meet $[R_{1.1}]$.
- (3) *Solutions for both kinds of tasks*: These solutions support both kinds of tasks by recording exact IDs and sizes of all flows, including FlowRadar [28], OmniMon [29], Counter Braids [30], Marple [31] and more. However, recording exact IDs and sizes of all flows requires at least memory/bandwidth overhead linear with the number of flows. Therefore, these solutions fail to meet $[R_{1.2}]$.

In summary, the first two categories of solutions cannot meet $[R_{1.1}]$ due to their limited measurement capabilities, and the third category of solutions cannot meet $[R_{1.2}]$ due to their linear space complexities. A naive solution meeting both $[R_{1.1}]$ and $[R_{1.2}]$ is to combine the first two categories of solutions: choosing one solution in the corresponding category for each kind of tasks. However, such a combination fails to achieve $[R_2]$ on programmable switches. The reason behind is that the data structures and operations of different categories of solutions usually differ significantly. For example, LossRadar [24] records the IDs and existences of packets using XOR operation and addition, while ElasticSketch [15] records the IDs and sizes of flows using comparison, substitution, and addition. Therefore, solutions in different categories can only utilize their resources allocated at compile time, which prohibits flexible allocation of memory resources between packet loss tasks and packet accumulation tasks. Therefore, the naive solution cannot pay attention to different kinds of tasks for different network states.

In this paper, we design ChameleMon, which meets all the above requirements simultaneously. ChameleMon can support almost all packet loss tasks and packet accumulation tasks. Compared to the state-of-the-art solutions, for packet loss tasks, ChameleMon reduces the memory overhead from proportional to the number of all flows (FlowRadar) or lost packets (LossRadar), to proportional to the number of flows experiencing packet losses, which we call *victim flows*; for packet accumulation tasks, ChameleMon achieves at least comparable accuracy. Our ChameleMon has a key design and a key technique as follows.

The key design of ChameleMon is to shift measurement attention as network state changes, which is just like the process of the

chameleons changing their skin coloration, through two dimensions of dynamics: 1) dynamically allocating memory between the two kinds of tasks; 2) dynamically monitoring the flows of importance. First, ChameleMon monitors the network state and allocates memory between the two kinds of tasks accordingly. When the network state is healthy and only a few packet losses occur in the network, ChameleMon pays most attention to and allocates most of the memory for packet accumulation tasks. As the network state degrades and packet losses increase, ChameleMon gradually shifts measurement attention to and allocates more and more memory for packet loss tasks to assist in diagnosing network faults. Second, ChameleMon ranks the flows according to their importance, and selects those of most importance to monitor. When the network state is ill and there are too many victim flows, ChameleMon selects those flows experiencing many packet losses (called *heavy-losses*, HLs for short) to monitor, instead of monitoring all victim flows.

Overall, when the network state continuously degrades from the healthy state to the ill state, ChameleMon runs as follows. 1) As the number of victim flows increases, ChameleMon leverages the first dimension of dynamic: gradually shifting measurement attention to and allocating more and more memory for packet loss tasks; 2) When the victim flows are too many to monitor, ChameleMon leverages the second dimension of dynamic: focusing measurement attention on HLs while monitoring a small portion of other packet losses (called *light-losses*, LLs for short) through sampling.

To realize the key design, ChameleMon incorporates a key technique, leveraging Fermat's little theorem¹ to devise a flexible data structure, namely *FermatSketch*. The data structure of FermatSketch is made of many same units. FermatSketch is dividable, additive, and subtractive, supporting packet loss detection and heavy-hitter (HH for short) detection simultaneously. By dividing FermatSketch into three parts to detect HLs, LLs, and HHs, ChameleMon can flexibly move the division points to shift attention and allocate memory between the two kinds of tasks as network state changes. For each incoming packet, we further use a flow classifier (TowerSketch [32]) to determine which of the three parts to insert. For packet loss detection, owing to Fermat's little theorem, FermatSketch only requires memory proportional to the number of victim flows. Differently, the state-of-the-art solutions require memory proportional to the number of all flows (FlowRadar) or lost packets (LossRadar).

Thanks to the visibility to per-flow size provided by TowerSketch, ChameleMon can support five other widely-studied [15, 18, 32, 33] packet accumulation tasks, including flow size estimation, heavy-change detection, flow size distribution estimation, entropy estimation, and cardinality estimation. We have fully implemented a ChameleMon prototype on a testbed with a Fat-tree topology composed of 10 Tofino switches and 8 end-hosts. We conduct extensive experiments and the results show that ChameleMon supports both kinds of tasks with low memory/bandwidth overhead, and more importantly, it can *automatically shift measurement attention as network state changes at run-time without recompilation*. We have released all related source codes at Github [34].

Ethics: This work does not raise any ethical issue.

¹Fermat's little theorem states that if p is a prime, then for any integer a that is indivisible by p , we have $a^{p-1} \equiv 1 \pmod{p}$.

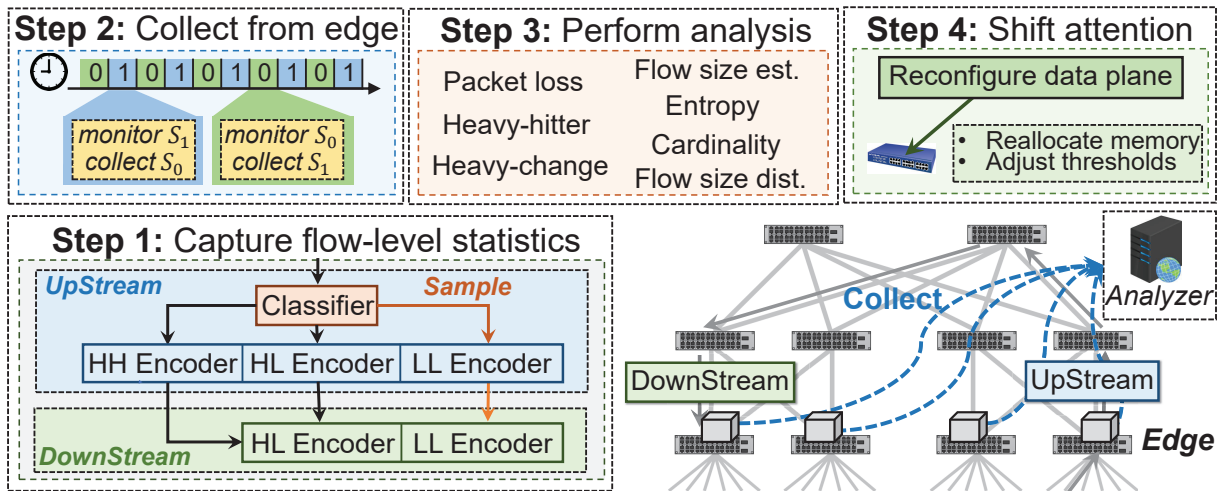


Figure 1: Overview of ChameleMon.

2 OVERVIEW OF CHAMELEMON

ChameleMon monitors the network in four steps (Figure 1).

1) Capturing flow-level statistics on edge switches: To capture desired flow-level statistics, ChameleMon deploys three sketches on the data plane of each edge switch, including a flow classifier (TowerSketch), an upstream flow encoder (our FermatSketch), and a downstream flow encoder (our FermatSketch). To detect HHs, HLs, and LLs, the upstream and downstream flow encoders are divided into multiple parts: 1) the upstream flow encoder is divided into an upstream HH encoder, an upstream HL encoder, and an upstream LL encoder; 2) the downstream flow encoder is divided into a downstream HL encoder and a downstream LL encoder. For every packet with flow ID f entering the network, according to the size of flow f , the flow classifier classifies flow f into one of three hierarchies: 1) HH candidate, 2) HL candidate, or 3) LL candidate. The LL candidate is further classified into sampled LL candidate or non-sampled LL candidate through sampling. Based on the hierarchy of flow f , the packet is then inserted into the corresponding part of the upstream flow encoder and downstream flow encoder when it enters and exits the network, respectively.

2) Collecting sketches from edge switches: A central controller periodically collects sketches from each edge switch to support persistent measurement. To avoid colliding with packet insertion when collecting sketches, each edge switch divides the timeline into consecutive fixed-length time intervals (called *epochs*), and copies a group of sketches for rotation. Every time an epoch ends, the central controller collects the group of sketches monitoring this epoch, and the other group of sketches starts to monitor the current epoch.

3) Performing network-wide analysis: Every epoch, the central controller performs network-wide analysis of the collected sketches to support seven measurement tasks. By analyzing the upstream and downstream flow encoders, the central controller can support packet loss detection. By analyzing the flow classifier and the upstream HH encoder, the central controller can support heavy-hitter detection and five other packet accumulation tasks.

4) Shifting measurement attention as network state changes: Every epoch, the central controller monitors the real-time network

state by analyzing the collected sketches. Then, the central controller reconfigures the data plane of edge switches at run-time according to the real-time network state, shifting measurement attention through two dimensions of dynamics. In the first dimension, the central controller dynamically allocates memory between packet loss tasks and packet accumulation tasks by reallocating the memory of the upstream and downstream encoders between their different parts. In the second dimension, the central controller dynamically selects the most important flows (HH/HL/sampled LL candidates) to monitor by adjusting the thresholds for flow classification and the sample rate for sampling LL candidates.

3 CHAMELEMON DATA PLANE

The ChameleMon data plane consists of the flow classifier, the upstream flow encoder, and the downstream flow encoder deployed on each edge switch. In this section, we detail the design of the ChameleMon data plane. First, we propose the key technique of ChameleMon, namely *FermatSketch*. Second, we detail each component of the ChameleMon data plane in sequence.

3.1 The FermatSketch Algorithm

Rationale: Our primary goal is to detect packet losses with low memory overhead. Existing solutions focus on either per-packet loss (LossRadar [24]) or all-flow visibility (FlowRadar [28]), incurring unacceptable memory overhead. To reduce overhead, we hope to aggregate all the lost packets of the same flow to detect per-flow packet losses. It is very challenging because existing solutions commonly use XOR operation for high memory efficiency and hardware-friendliness, but simply using XOR operation to aggregate flow IDs of lost packets causes every two lost packets of the same flow to cancel each other out. While invertible Bloom lookup table (IBLT) [35] can overcome this challenge as IBLT uses addition to aggregate flow IDs, such design requires computation over large numbers, and thus complicates the implementation of IBLT on programmable switches. To address this challenge while maintaining hardware-friendliness, we devise FermatSketch, which uses modular addition to aggregate flow IDs and leverages Fermat's little theorem to extract flow IDs.

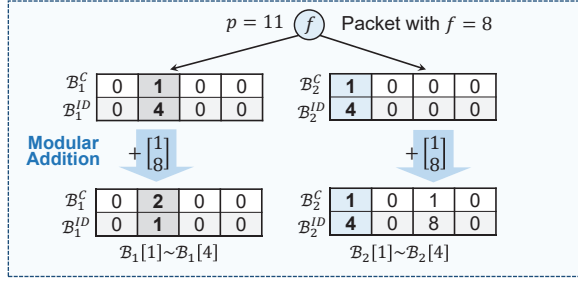


Figure 2: An example of encoding/insertion.

Data structure (Figure 2): FermatSketch has d equal-sized bucket arrays $\mathcal{B}_1, \dots, \mathcal{B}_d$, each of which consists of m buckets. Each bucket array \mathcal{B}_i is associated with a pairwise-independent hash function $h_i(\cdot)$ that maps each incoming packet into one bucket (called mapped bucket) in it. Each bucket $\mathcal{B}_i[j]$ contains two fields: 1) a *count field* $\mathcal{B}_i^c[j]$ recording the number of packets mapped into the bucket; 2) an *IDsum field* $\mathcal{B}_i^{ID}[j]$ recording the result of the sum of flow IDs of packets mapped into the bucket modulo a prime p . At initialization, we set all fields of all buckets in FermatSketch to zero, and p to a prime that must be larger than any available flow ID f and the size of any flow, so as to make use of Fermat's little theorem.

Encoding/Insertion operation (Figure 2): To encode an incoming packet with flow ID f , we first calculate the d hash functions to locate d mapped buckets: $\mathcal{B}_1[h_1(f)], \mathcal{B}_2[h_2(f)], \dots, \mathcal{B}_d[h_d(f)]$. For each mapped bucket $\mathcal{B}_i[h_i(f)]$, we update it as follows. First, we increment its count field $\mathcal{B}_i^c[h_i(f)]$ by one. Second, we update its IDsum field through *modular addition*: $\mathcal{B}_i^{ID}[h_i(f)] \leftarrow ((\mathcal{B}_i^{ID}[h_i(f)] + f) \bmod p)$. The pseudo-code of encoding operation is shown in Algorithm 1.

Decoding operation: The decoding operation, which can extract exact flow IDs and flow sizes from FermatSketch, has two important suboperations: 1) *pure bucket verification* that verifies whether a bucket only records packets of a single flow (pure bucket); 2) *single flow extraction* that extracts and deletes a single flow and its size from all its mapped buckets. Next, we propose the workflow of decoding operation and detail the two suboperations. The pseudo-code of decoding operation is shown in Algorithm 2.

• *Decoding workflow (Figure 3):* decoding proceeds as follows.

- ① Traverse FermatSketch and push all non-zero buckets to decoding queue.
- ② Pop a bucket from queue.
- ③ For the popped bucket, we perform pure bucket verification to verify whether it is a pure bucket. If not, we simply ignore the bucket and go back to step ②.
- ④ If so, we perform single flow extraction to extract and delete a single flow and its size from the pure bucket as well as the other mapped buckets of the single flow.
- ⑤ We insert the extracted single flow and its size into a hash table, namely *Flowset*, which is used to record all the extracted flows and their sizes. We regard all flows recorded in Flowset as the flows previously encoded into FermatSketch.
- ⑥ Except the pure bucket, we push the other mapped non-zero buckets of the extracted flow into queue.

Algorithm 1: Encoding/Insertion operation of FermatSketch

Input: Flow ID f

- 1 **for** $i \in [1, d]$ **do**
- 2 $j = h_i(f)$;
- 3 $\mathcal{B}_i^{ID}[j] = (\mathcal{B}_i^{ID}[j] + f) \bmod p$;
- 4 $\mathcal{B}_i^c[j] ++$;
- 5 **end**

Algorithm 2: Decoding operation of FermatSketch

- 1 **Function** IsPure(i, j):
- 2 $f = (\mathcal{B}_i^{ID}[j] \times \mathcal{B}_i^c[j])^{(p-2)} \bmod p$;
- 3 **return** $j == h_i(f)$;
- 4 **Function** Delete($\mathcal{B}_{i'}[j'], \mathcal{B}_i[j]$):
- 5 $\mathcal{B}_{i'}^c[j'] = \mathcal{B}_{i'}^c[j'] - \mathcal{B}_i^c[j]$;
- 6 $\mathcal{B}_{i'}^{ID}[j'] = (\mathcal{B}_{i'}^{ID}[j'] - \mathcal{B}_i^{ID}[j]) \bmod p$;
- 7 **Function** Decode():
- 8 Queue is an empty queue;
- 9 Flowset is an empty map;
- 10 **for** $i \in [1, d], j \in [1, w]$ **do**
- 11 **if** $\mathcal{B}_i^c[j] \neq 0$ **then**
- 12 Queue.push($\mathcal{B}_i[j]$);
- 13 **end**
- 14 **end**
- 15 **while** !Queue.empty() **do**
- 16 $\mathcal{B}_i[j] = \text{Queue.front}()$;
- 17 Queue.pop();
- 18 **if** IsPure(i, j) **then**
- 19 $f' = (\mathcal{B}_i^{ID}[j] \times \mathcal{B}_i^c[j])^{(p-2)} \bmod p$;
- 20 Flowset[f'] = Flowset[f'] + $\mathcal{B}_i^c[j]$;
- 21 **for** $i' \in [1, d]$ **do**
- 22 Delete($\mathcal{B}_{i'}[h_{i'}(f')], \mathcal{B}_i[j]$);
- 23 **if** $\mathcal{B}_{i'}^c[h_{i'}(f')] \neq 0$ **then**
- 24 Queue.push($\mathcal{B}_{i'}[h_{i'}(f')]$);
- 25 **end**
- 26 **end**
- 27 **end**
- 28 **end**
- 29 **return** Flowset

⑦ Check whether the queue is empty. If so, the decoding stops. Otherwise, go back to step ②. After stopping, if there are still non-zero buckets in FermatSketch, the decoding is considered as failed. Otherwise, the decoding is considered as successful.

• *Pure bucket verification:* The pure bucket verification reports whether one given bucket is pure (i.e., only records a single flow), but it may misjudge a non-pure bucket as a pure one with a small probability $\frac{1}{m}$. Suppose a bucket $\mathcal{B}_i[j]$ only records a single flow f' , it should satisfy that $(\mathcal{B}_i^c[j] \times f') \bmod p = \mathcal{B}_i^{ID}[j]$. Leveraging Fermat's little theorem, we can get that $f' = (\mathcal{B}_i^{ID}[j] \times (\mathcal{B}_i^c[j])^{p-2}) \bmod p$. Considering that bucket $\mathcal{B}_i[j]$ should be one of the d mapped buckets of flow f' , to verify whether $\mathcal{B}_i[j]$ is a pure bucket, we propose a verification method namely *reshaping verification*. First, we calculate the i^{th} hash function $h_i(\cdot)$ to locate

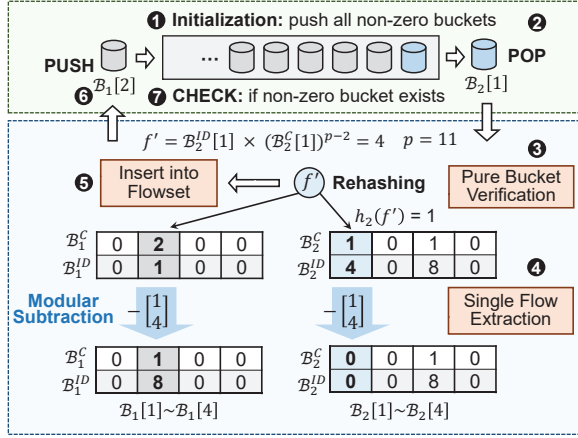


Figure 3: An example of decoding.

the i^{th} mapped bucket of f' , i.e., we calculate $h_i(f')$. Then we check whether $h_i(f')$ is equal to j . If so, we consider $B_i[j]$ as a pure bucket recording flow f' with size $B_i^C[j]$. Note that the false positive rate of pure bucket verification, i.e., the probability of misjudging a non-pure bucket as a pure one, is $\frac{1}{m}$, which is calculated as follows. For any non-pure bucket, we can calculate its flow ID, which should be considered as a random value. The probability that a random ID is hashed to the same bucket is $\frac{1}{m}$. In this section, we further discuss that such false positives can be automatically eliminated during decoding, and prove they have little impact on decoding success rate through mathematical analysis (Theorem 3.1).

- **Single flow extraction:** To extract/delete flow f' from $B_i[j]$ as well as its other mapped buckets, first, we locate its other $(d - 1)$ mapped buckets. Second, for each mapped bucket $B_i[h_i(f')$], we decrement its count field $B_i^C[h_i(f')$ by $B_i^C[j]$, and update its IDsum field to $((B_i^D[h_i(f')] - B_i^D[j]) \bmod p)$ through modular subtraction.

Addition/Subtraction operations: Adding/Subtracting FermatSketch FS_1 to/from FermatSketch FS_2 . FS_1 and FS_2 must use the same parameters including hash functions, number of arrays, number of buckets, and primes. For each bucket of FS_2 , we update it as follows. First, we locate the bucket of FS_1 that is in the same position as it. Second, we add/subtract the count field of the located bucket of FS_1 to/from its count field. Third, we modular add/subtract the IDsum field of the located bucket of FS_1 to/from its IDsum field.

Space complexity: Suppose FermatSketch is large enough, and then the pure bucket verification has negligible false positive rate. The decoding operation is almost the same as that of IBLT [35], which is exactly the procedure used to find the 2-core of a random hypergraph [36, 37]. Therefore, the memory overhead of FermatSketch is proportional to the number of inserted flows M , i.e., $\Theta(M)$. FermatSketch also shares similar properties with IBLT: the number of hash functions, i.e., the number of the bucket arrays d , is recommended to set to 3 for the highest memory efficiency, that on average 1.23 buckets can record a flow and its size.

Time complexity of decoding operation: Suppose FermatSketch is large enough and the false positive rate in pure bucket verification is negligible. In step ①, we traverse FermatSketch and push all non-zero buckets into the decoding queue. The number of these

buckets is at most md , and thus the time complexity of step ① is $O(md)$. In the rest steps, we process all the buckets pushed into the queue, which consists of two parts: 1) the md buckets pushed into in step ①, and 2) the mapped buckets except the popped pure bucket of each extracted flow. Considering that the number of extracted flows is bounded by the number of buckets of FermatSketch, i.e., md , the number of buckets of the second part is $O(md^2)$. Therefore, the time complexity of the rest steps is $O(md^2)$. Adding up the time complexities of all steps, the time complexity of decoding operation is $O(md^2)$.

Eliminating false positives during decoding: Due to hash collisions, the rehashing verification will inevitably misjudge some non-pure buckets as pure buckets with false positive rate $\frac{1}{m}$. Such misjudgement will lead to extraction of flows that are not inserted into FermatSketch, and finally could hinder the decoding. From another point of view, extracting a flow from such a misjudged non-pure bucket, i.e., false positive, equals to inserting a fabricated flow with a negative size into FermatSketch. The decoding operation can automatically eliminate these false positives: in the decoding procedure, these inserted fabricated flows could also be extracted and deleted from FermatSketch, and then the impact caused by the false positives disappears. We use Theorem 3.1 to show that when $md > c_d M + \epsilon$ and M (the number of inserted flows) is not too small, the decoding only fails with an extremely small probability $O(\frac{1}{M^{d-2}})$, where c_d refers to the minimum average number of buckets required to record a flow and its size for a d -array FermatSketch.

THEOREM 3.1. Suppose $md > c_d M + \epsilon$ and $M = \Omega(d^{4d} \log^d(md))$. The decoding of FermatSketch fails with probability $O(\frac{1}{M^{d-2}})$, where both ϵ and c_d are small constants.

$$c_d = \left(\sup \left\{ \alpha \mid \alpha \in (0, 1), \forall x \in (0, 1), 1 - e^{-d\alpha x^{d-1}} \right\} \right)^{-1}$$

For example, $c_3 = 1.23$, $c_4 = 1.30$, $c_5 = 1.43$.

The detailed proof is presented in Appendix A.1.

Packet loss detection: To support packet loss detection, we can deploy a group of FermatSketches on edge switches to encode the packets entering the network, and another group of FermatSketches to encode the packets exiting the network. Thanks to the additivity and subtractivity of FermatSketch, for each group, we add up the FermatSketches in it to obtain a cumulative FermatSketch encoding all the packets entering or exiting the network. Then, we subtract the cumulative FermatSketch encoding all the packets exiting the network from the other one, and the FermatSketch after subtraction just encodes all the victim flows in the network. Therefore, this FermatSketch just requires memory proportional to the number of victim flows for successful decoding. In other words, FermatSketch can support packet loss detection with memory overhead proportional to the number of victim flows.

[Optional] fingerprint verification: To reduce the false positive rate of pure bucket verification, we propose an extra verification method, namely *fingerprint verification*. Please refer to Appendix A.2 for details.

3.2 Data Plane Components

As shown in Figure 1, every packet entering the network undergoes the three components of the Chameleon data plane in sequence:

1) the flow classifier, 2) the upstream flow encoder, and 3) the downstream flow encoder.

3.2.1 Flow Classifier.

Rationale: To detect HHs, HLs, and LLs, ChameleonMon deploys the flow classifier in the ingress of each edge switch, so as to classify flows into different hierarchies. While it is easy to select HHs to monitor according to flow sizes, it is not easy to select HLs to monitor because we can hardly predict how many packets a flow will lose. Our observation is that for each flow, the number of its lost packets cannot exceed its size. Therefore, the sizes of HLs should have a minimum value. ChameleonMon selects flows whose sizes exceed this value to monitor, so as to approximate the monitoring of HLs. In summary, the flow classifier classifies flows purely according to flow sizes. We choose TowerSketch [32], a simple, accurate, and hardware-friendly sketch, as the flow classifier.

Data Structure: The flow classifier consists of l equal-sized arrays. The i^{th} array \mathcal{A}_i consists of w_i δ_i -bit counters, where $w_i \times \delta_i$ is a constant and $\delta_{i-1} < \delta_i$. Also, array \mathcal{A}_i is associated with a pairwise-independent hash function $s_i(\cdot)$. For each δ_i -bit counter, its maximum value $2^{\delta_i} - 1$ is used to represent the state that it is overflowed, and thus be regarded as $+\infty$.

Insertion: To insert a packet with flow ID f , we first calculate the l hash functions to locate l counters: $\mathcal{A}_1[s_1(f)], \mathcal{A}_2[s_2(f)], \dots, \mathcal{A}_l[s_l(f)]$. We call these counters *the l mapped counters*. Then, for each of the l mapped counters, we increment it by one unless it is overflowed.

Online query: To query the size of flow f online, we simply report the minimum value among the l mapped counters.

Packet processing: For a packet with flow ID f entering the network, the flow classifier processes it as follows. First, we insert it into the flow classifier and query the size of flow f . Then, with the queried flow size, we classify flow f into the corresponding hierarchy according to two thresholds T_h and T_l , where T_h is used for selecting HH candidates, and T_l is used for selecting HL candidates. In general, it satisfies that $T_l \leq T_h$. If the flow size is larger than or equal to T_h , flow f is classified as a HH candidate. If the flow size is less than T_l , flow f is classified as a LL candidate. If the flow size is between T_l and T_h , flow f is classified as a HL candidate. The LL candidate is further classified into sampled LL candidate or non-sample LL candidate through sampling.

3.2.2 Upstream Flow Encoder.

Rationale: To support packet loss detection, ChameleonMon deploys the upstream flow encoder in the ingress of each edge switch just after the flow classifier, so as to encode the packets entering the network. Therefore, the upstream flow encoder should contain two FermatSketches to encode HL candidates and sampled LL candidates individually. Here, for better monitoring of the network state, ChameleonMon monitors a portion of LLs to maintain an overview of all victim flows. Besides, to support heavy-hitter detection, the upstream flow encoder should contain a FermatSketch to encode HH candidates. In summary, the upstream flow encoder should consist of three FermatSketches.

Data structure: The upstream flow encoder is a d -array FermatSketch divided into three d -array FermatSketches: 1) an upstream

HH encoder for encoding HH candidates; 2) an upstream HL encoder for encoding HL candidates; 3) an upstream LL encoder for encoding sampled LL candidates. We denote the number of buckets per array of the upstream flow encoder, HH encoder, HL encoder, and LL encoder by m_{uf} , m_{hh} , m_{hl} , and m_{ll} , respectively. Obviously, it satisfies that $m_{uf} = m_{hh} + m_{hl} + m_{ll}$.

Packet processing: For a packet with flow ID f entering the network, the upstream flow encoder processes it by encoding the packet into one of the encoders corresponding to the hierarchy of flow f unless flow f is a non-sampled LL candidate. Here, the hierarchy of flow f can be directly obtained because the upstream flow encoder and the flow classifier are deployed on the same edge switch.

3.2.3 Downstream Flow Encoder.

Rationale: To support packet loss detection, ChameleonMon deploys the downstream flow encoder in the egress of each edge switch, so as to encode the packets exiting the network. As the downstream flow encoder is not responsible for heavy-hitter detection, it should consist of two FermatSketches to encode HL candidates and sampled LL candidates.

Data structure: The downstream flow encoder is a d -array FermatSketch divided into two d -array FermatSketches: 1) a downstream HL encoder; 2) a downstream LL encoder. To support packet loss detection, the number of buckets per array of the downstream HL encoder and LL encoder must also be m_{hl} and m_{ll} , respectively, so as to support addition and subtraction operations with the corresponding upstream encoder. We denote the number of buckets per array of the downstream flow encoder by m_{df} . In general, it satisfies that $m_{df} < m_{uf}$, and therefore satisfies that $m_{df} \geq m_{hl} + m_{ll}$.

Packet processing: For a packet with flow ID f exiting the network, the downstream flow encoder processes it by encoding the packet into one of the encoders corresponding to the hierarchy of flow f unless flow f is a non-sampled LL candidate. Here, packets of HH candidates are also encoded into the downstream HL encoder. Different from the upstream flow encoder, the downstream flow encoder cannot directly obtain the flow hierarchy from the flow classifier, as a flow could enter and exit the network at different edge switches. To address this issue, first, considering that there are four flow hierarchies, we can use $\lceil \log(4) \rceil = 2$ bits in the original packet header to transmit this information. For example, for IPv4 protocol, we can use the unused bits in the type of service (ToS) field. If there are not enough unused bits, second, we can transmit the flow hierarchy in an INT-like [38] manner.

4 CHAMELEMON CONTROL PLANE

The ChameleonMon control plane consists of a central controller, as well as the control plane of each edge switch. In this section, we detail the design of the ChameleonMon control plane. We begin by laying out how the ChameleonMon control plane collects sketches from the ChameleonMon data plane, then introduce how to support seven measurement tasks with the collected sketches, and finally propose how to shift measurement attention as network state changes.

4.1 Collection from Data Plane

The central controller needs to periodically collect sketches, *i.e.*, the flow classifier, the upstream flow encoder, and the downstream flow

encoder, from the ChameleMon data plane, so as to support persistent measurement. However, the collection cannot be completed in an instant, and thus inevitably collide with packet insertion if there is only a group of sketches. Specifically, if the central controller wants to collect sketches at time t , it will inevitably collect some counters inserted by packets after t , which could result in decoding failure of FermatSketch. To address this issue, ChameleMon takes two steps: 1) *timeline division* and 2) *clock synchronization*. Next, we just briefly cover the two steps. We detail the two steps in Appendix B, where we further analyze *the appropriate time* for the central controller to collect sketches.

Timeline division: Each edge switch periodically flips a 1-bit timestamp to divide the timeline into fixed-length time intervals (called epochs) with interleaved 0/1 timestamp, and copies a group of sketches for rotation. Each group of sketches corresponds to a distinct timestamp value, and monitors the epochs with that timestamp value.

Clock synchronization: The central controller also maintains a 1-bit periodically flipping timestamp, and periodically synchronizes its clock with the control plane of each edge switch, so as to make opportunities for collection.

Every time the locally maintained 1-bit timestamp flips, an epoch ends, the central controller starts to collect the group of sketches monitoring this epoch, and the other group of sketches starts to monitor the current epoch.

4.2 Measurement Tasks

With the collected sketches, the central controller can support packet loss detection and six packet accumulation tasks.

Packet loss detection: reporting each victim flow and the number of its lost packets. The central controller can support packet loss detection by analyzing the upstream and downstream flow encoders collected from each edge switch. First, for each edge switch, we decode the upstream HH encoder to obtain the HH Flowset, and then reinsert each flow with its size in the HH Flowset into the upstream HL encoder. Second, we add up the upstream/downstream HL/LL encoder of each edge switch through addition operation to obtain the cumulative upstream/downstream HL/LL encoder. Third, we subtract the cumulative downstream HL/LL encoder from the cumulative upstream HL/LL encoder to obtain the *delta HL/LL encoder*. Fourth, we decode the delta HL/LL encoder to obtain the HL/LL Flowset. Finally, we report the flows in the HL Flowset as HLs, and the flows in the LL Flowset but not in the HL Flowset as LLs. For each of these flows, its estimated number of lost packets is the sum of its size in the HL Flowset and the LL Flowset.

For each edge switch, the central controller can support the following six widely-studied [15, 18, 32, 33] packet accumulation tasks by analyzing the flow classifier and upstream HH encoder collected from it. Then, by synthesizing the results of each edge switch, the central controller can easily support these tasks in a network-wide manner. We detail these six tasks from the perspective of an edge switch.

Heavy-hitter detection: reporting flows whose sizes exceed Δ_h . First, we decode the upstream HH encoder to obtain the HH Flowset, which records flows with ID f_i and size q_i . For any flow f_j in the HH Flowset, if its estimated flow size $\mathcal{T}_h + q_j$ is larger than Δ_h , we

report it as a HH. Note that T_h is the threshold used for selecting HH candidates.

Flow size estimation: reporting flow size of flow f_j . Similarly, we obtain the HH Flowset. If flow f_j is in the HH Flowset, we report its flow size as $\mathcal{T}_h + q_j$. Otherwise, we report its flow size as query result from the flow classifier.

Heavy-change detection: reporting flows whose sizes change beyond Δ_c in two adjacent epochs. Similarly, we obtain the HH Flowset. For any flow f_j in the HH Flowset of either epoch, we estimate its flow size in the two epochs. If the difference between the two estimated flow sizes is larger than Δ_c , we report flow f_j as a heavy-change.

Cardinality estimation: reporting number of flows. We apply linear-counting algorithm [39] to the counter array with most counters in the flow classifier for estimation.

Flow size distribution estimation: reporting distribution of flow sizes. We apply MRAC algorithm [40] to each counter array in the flow classifier. Array \mathcal{A}_i provides distribution of flow size in range $[2^{\delta_i-1} - 1, 2^{\delta_i} - 1)$. The remaining distribution of flow size in range $[2^{\delta_i} - 1, +\infty)$ is estimated from the flows larger than $2^{\delta_i} - 2$ in the HH Flowset.

Entropy estimation: reporting entropy of flow sizes. Based on the estimated flow size distribution, we can easily compute the entropy as follows: $-\sum \left(n_i \cdot \frac{1}{N} \log \frac{1}{N} \right)$, where n_i is the number of flows of size i , and $N = \sum (i \cdot n_i)$.

4.3 Shifting Measurement Attention

A practical measurement system should pay attention to different kinds of tasks for different network states. When there are only rare packet losses in network, the system should pay more attention to and allocate more memory for packet accumulation tasks. In contrast, when there are lots of packet losses in network, the system should pay more attention to and allocate more memory for packet loss detection to help diagnose network faults.

Aiming at this target, ChameleMon decides to shift measurement attention as network changes at run-time without recompilation. Every time all the sketches monitoring the previous epoch are collected, ChameleMon takes two phases to shift measurement attention. First, the central controller monitors the real-time network state, including the number and flow size distribution of flows and victim flows, by analyzing the collected sketches. Second, the central controller reconfigures the ChameleMon data plane according to the real-time network state while *maintaining high memory utilization*. The central controller not only reallocates memory of the upstream and downstream encoders between their different parts, but also adjusts the thresholds for flow classification and the sample rate for sampling LL candidates. To avoid interference with the monitoring of the current epoch, the reconfiguration will not function immediately, but in the next epoch.

For ChameleMon, the network state could be clearly classified into two levels: 1) healthy network state that ChameleMon can allocate sufficient memory to monitor all victim flows; 2) ill network state that ChameleMon cannot allocate sufficient memory to monitor all victim flows, and thus must select HLs to monitor. For each level of network state, ChameleMon behaves similarly in shifting measurement attention, and we detail how it behaves in this section.

4.3.1 Healthy Network State.

Suppose the previously monitored network state is healthy, and now the central controller starts to shift measurement attention. Currently, the LL encoders are not allocated any memory as Chameleon can monitor all victim flows, and T_l must be 1 as no flows should be classified into LL candidates. The memory allocation between the upstream HH encoder and the upstream HL encoder is flexible.

Monitoring real-time network state: The monitoring proceeds as follows. First, for each edge switch, the central controller estimates the number of flows and flow size distribution² as described above (§ 4.2). Second, for each edge switch, the central controller obtains the number of HH candidates by decoding the upstream HH encoder. After all decoding steps, if the decoding of any upstream HH encoder fails, the central controller stops the monitoring as the decoding of the delta HL encoder requires reinserting the decoded HH candidates into the upstream HL encoders. Third, the central controller obtains the number of HLs (equals to victim flows for healthy network state) by decoding the delta HL encoder as described above (§ 4.2). If the decoding fails, the central controller estimates the number of HLs by applying linear-counting algorithm to any bucket array of the delta HL encoder.

Reconfiguring Chameleon data plane: The core idea of reconfiguration is to first ensure the successful decoding of FermatSketches for supporting packet loss detection and heavy-hitter detection, while maintaining high memory utilization. The reconfiguration proceeds as follows.

Step 1: We focus on the successful decoding of the upstream HH encoders as they are decoded first. For each edge switch, if the decoding of the upstream HH encoder fails, the central controller turns up T_h according to the number of flows and flow size distribution, controlling the expected load factor³ of the upstream HH encoder at around 70%⁴, so as to maintain high memory utilization. After turning up T_h , the central controller stops the reconfiguration as the decoding of the delta HL encoder cannot proceed.

Step 2: We focus on the successful decoding and high memory utilization of the delta HL encoder. If the decoding of the delta HL encoder fails, according to the estimated number of HLs, the central controller estimates the required memory for 70% load factor. If the maximum memory that the HL encoders can be allocated to, *i.e.*, all the memory of the downstream flow encoder, cannot cover the required memory, the healthy network state transitions to the ill network state. In this case, the central controller 1) reallocates the memory inside the upstream and downstream flow encoders as the fixed allocation described in the ill network state (§ 4.3.2), 2) sets T_l to T_h , and 3) adjusts the sample rate for 70% load factor of the delta LL encoder assuming that each HL will be a LL. Otherwise, the central controller just expands the HL encoders to the required

memory. If the decoding of the delta HL encoder succeeds and its load factor is lower than 60%, the central controller tries to compress the HL encoders to approach 70% load factor for high memory utilization. Here, we *reserve the minimum memory* for the HL encoders to handle the potential small burst of victim flows.

Step 3: After all the memory reallocation, we focus on the successful decoding and high memory utilization of the upstream HH encoders. For each edge switch, with the number of HH candidates and the memory of the upstream HH encoder, the central controller further estimates the expected load factor of the upstream HH encoder. If the expected load factor of the upstream HH encoder is lower than 60% or larger than 70%, the central controller turns down or up T_h to approach 70% load factor.

4.3.2 Ill Network State.

Suppose the previously monitored network state is ill, and now the central controller starts to shift measurement attention. Currently, all the HH, HL and LL encoders are allocated fixed memory, and T_l must be larger than 1 to select HL candidates. Specifically, the upstream HH encoder is compressed to the minimum memory, which is the memory difference between the upstream flow encoder and the downstream flow encoder.

Monitoring real-time network state: The monitoring proceeds in a similar way to that of the healthy network state. In addition, the central controller obtains the number of LLs by decoding the delta LL encoder as described above (§ 4.2). If the decoding fails, the central controller estimates the number of LLs by applying linear-counting algorithm to the delta LL encoder, and then stops the monitoring. If both decoding of the delta HL and LL encoders succeeds, the central controller estimates the number and flow size distribution of victim flows as follows. First, the central controller samples the HLs with the same sampling method and rate as LLs. Second, the central controller merges sampled HLs and sampled LLs to obtain sampled victim flows. Third, the central controller estimates the flow size distribution of victim flows through querying the flow size of each sampled victim flow, and the number of victim flows through dividing the number of sampled victim flows by sample rate. If the decoding of the delta HL encoder fails, the central controller regards the estimated flow size distribution of sampled LLs, which is also estimated by querying flow sizes, as the flow size distribution of victim flows.

Reconfiguring Chameleon data plane: The core idea of reconfiguration is the same as that of the healthy network state. The reconfiguration proceeds as follows.

Step 1: We focus on the successful decoding of the upstream HH encoders, and the reconfiguration proceeds the same as the first step of the healthy network state. In addition, we focus on the successful decoding of the delta LL encoder. If the decoding of the delta LL encoder fails, according to the estimated number of LLs, the central controller adjusts the sample rate to make the delta LL encoder approach 70% load factor, and then stops the reconfiguration.

Step 2: We focus on the successful decoding of the delta HL encoder. If the decoding of the delta HL encoder fails, according to the estimated flow size distribution of victim flows, assuming that each victim flow larger than T_l will be a HL, the central controller turns up T_l to make the delta HL encoder approach 70% load factor.

Step 3: We focus on the high memory utilization of the HL and LL encoders. If both the decoding of the delta HL and LL encoders

²The estimation of flow size distribution using the MRAC algorithm typically takes several seconds to perform multiple iterations. Therefore, we recommend either 1) monitoring the flow size distribution over time intervals longer than epoch length or 2) reducing the number of iterations to support more real-time monitoring.

³Load factor refers to the ratio of the number of recorded flows to the number of buckets of FermatSketch. FermatSketch achieves the highest memory efficiency when d is set to 3, that on average 1.23 buckets can record a flow and its size. Thus, the maximum load factor of FermatSketch is around $81.3\% = \frac{1}{1.23}$.

⁴Here, we decide not to pursue the maximum load factor for two reasons: 1) the potential increase of HH candidates in the current epoch and 2) the inevitable estimation error in linear-counting.

succeeds, according to the estimated number of victim flows, the central controller estimates the required memory for monitoring all the victim flows with 70% load factor. If the downstream flow encoder can cover the required memory, the ill network state transitions to the healthy network state. In this case, the central controller 1) eliminates the LL encoders, 2) allocates the required memory (at least the reserved minimum memory) to the HL encoders, and 3) sets T_l to 1. If the downstream flow encoder cannot cover the required memory, and the load factor of the delta HL encoder or the delta LL encoder is lower than 60%, the central controller turns up T_l or the sample rate according to the estimated flow size distribution of victim flows or the estimated number of LLs, respectively, so as to approach 70% load factor.

Step 4: After all the memory reallocation, we focus on the successful decoding and high memory utilization of the upstream HH encoders, and the reconfiguration proceeds the same as the third step of the healthy network state.

5 EVALUATION

We conduct various experiments on CPU platform and our testbed, and focus on the following five key questions.

How much memory/time can ChameleMon save in packet loss detection? (Figure 4-6) We implement FermatSketch and its competitors in C++, and use CAIDA dataset [41] to evaluate their memory and time overhead for packet loss detection on CPU platform. Results show that FermatSketch can save memory in all cases and time in most cases.

How accurately can ChameleMon support six packet accumulation tasks? (Figure 11 in Appendix C) We implement the combination of TowerSketch and FermatSketch and its competitors in C++, and use CAIDA dataset to evaluate their accuracy for these six tasks on CPU platform. Results show that the combination can achieve at least comparable accuracy in all six tasks.

Can ChameleMon automatically shift measurement attention? (Figure 7-8) We generate workloads according to widely used traffic distributions (e.g., DCTCP [42]) for evaluation. We use the above workloads to evaluate ChameleMon by generating different network states on our testbed. Results show that ChameleMon can always automatically shift measurement attention as network state changes at run-time, and maintains high memory utilization in most cases.

How fast can ChameleMon shift measurement attention? (Figure 9) We use the above workloads to evaluate ChameleMon over a large time window, in which the network state changes 8 times. Results show that ChameleMon can shift measurement attention within at most 3 epochs.

How fast can ChameleMon monitor the network? (Figure 20-22 in Appendix F) We use the above workloads to evaluate various factors that can affect the epoch length. Results show that ChameleMon can monitor the network every 50ms on our testbed, using only one CPU core and consuming only 320Mbps bandwidth. We believe ChameleMon can easily scale to monitor a much larger network in a faster manner.

5.1 Evaluation on Packet Loss Detection

Dataset: We use the anonymized IP traces collected in 2018 from CAIDA [41] as dataset, and use the 32-bit source IP address as the

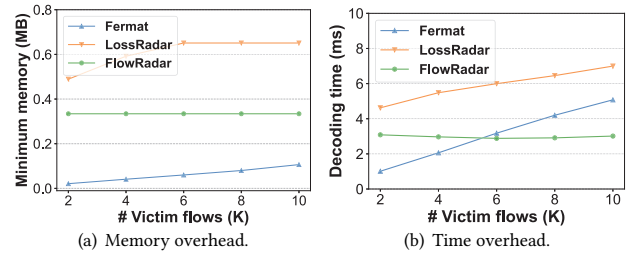


Figure 4: Memory/Time overhead vs. # victim flows.

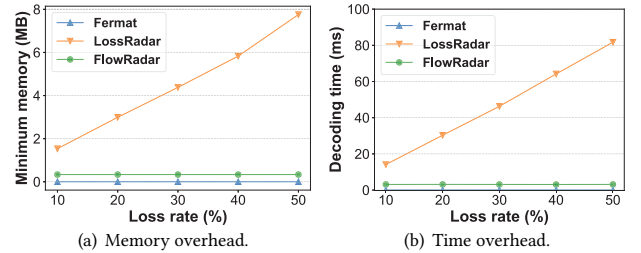


Figure 5: Memory/Time overhead vs. packet loss rate.

flow ID. We use the first 100K flows containing 5.3M packets for evaluation.

Setup: We set up a simulation with a simple topology consisting of only a link on CPU platform. We compare FermatSketch with FlowRadar [28] and LossRadar [24]. For FermatSketch, we set its count field and ID field to 32bits, and the number of hash functions to 3. For FlowRadar, we allocates 10% memory to the flow filter and 90% memory to the counting table. For the flow filter, which is actually a Bloom filter [43], we sets its number of hash functions to 10. For the counting table, we set its FlowXOR field, FlowCount field, and PacketCount field to 32bits, and its number of hash functions to 3. For LossRadar, we set its count field to 32bits, xorSum field to 48bits, and number of hash functions to 3. Here, the xorSum field of LossRadar encodes a 32-bit flow ID as well as a 16-bit packet-specific information that represents the order of a packet in a flow. For each solution, we deploy it upstream and downstream of the link to encode the packets entering and exiting the link.

Memory/Time overhead⁵ vs. number of victim flows (Figure 4): Experimental results show that the memory/time overhead of FermatSketch is proportional to the number of victim flows. We let the largest 10K flows pass through the link, among which a part of flows are victim flows. The packet loss rate of victim flows is set to 1%. As the number of victim flows increases, the memory/time overhead of FlowRadar remains unchanged, while that of FermatSketch increases almost linearly. We find when the number of victim flows exceeds 6000, the decoding time of FermatSketch exceeds that of FlowRadar. This is because the decoding operation of FermatSketch is more complex than FlowRadar. Compared to FlowRadar/LossRadar, FermatSketch saves up to 15.9/23.2 times memory and up to 3.0/4.6 times decoding time.

Memory/Time overhead vs. packet loss rate (Figure 5): Experimental results show that the memory/time overhead of FermatSketch is independent of the number of lost packets. We let the largest 10K flows pass through the link, among which the largest 100 flows

⁵The memory overhead refers to the minimum memory required to achieve 99.9% decoding success rate, and the time overhead refers to the corresponding decoding time with the minimum memory.




Figure 6: Memory/Time overhead vs. # flows.

are victim flows. As the packet loss rate of victim flows increases, the memory/time overhead of FermatSketch and FlowRadar remains unchanged, while that of LossRadar increases linearly. Compared to FlowRadar/LossRadar, FermatSketch saves up to 276.1/6411.2 times memory and up to 64.5/1585.6 times decoding time.

Memory/Time overhead vs. number of flows (Figure 6): Experimental results show that the memory/time overhead of FermatSketch is independent of the number of flows. We let a certain number of flows pass through the link, among which the largest 100 flows are victim flows. The packet loss rate of victim flows is set to 1%. As the number of flows increases, the memory/time overhead of FermatSketch and LossRadar remains unchanged, while that of FlowRadar increases linearly. Compared to FlowRadar/LossRadar, FermatSketch saves up to 1535.0/128.8 times memory and up to 821.3/23.7 times decoding time.

5.2 Evaluation on Testbed

Testbed setup: We have fully implemented a ChameleMon prototype on a testbed with a Fat-tree topology composed of 10 Tofino switches and 8 servers, with 1400 lines of P4 [44] code and 2400 lines of C/C++ code. Each server has 48 2.1GHz CPU cores, 256 GB RAM, and a 40Gb Mellanox Connectx-3 Pro NIC. Switches and servers are interconnected with 40Gb links. We deploy the ChameleMon data plane on all four ToR/edge switches. An additional server linked with a certain edge switch works as the central controller. For implementation details of the ChameleMon data plane and control plane, please refer to Appendix D.

Workloads: We generate workloads consisting of UDP flows according to four widely used distribution: DCTCP [42], HADOOP [45], VL2 [46] and CACHE [47]. We regard the distribution as known input to ChameleMon. We use the 104-bit 5-tuple as the flow ID. For each flow, We choose its source and destination IP address uniformly, and therefore each server sends and receives almost the same number of flows. The packet sender and receiver are integrated into a program written in DPDK [48]. To manually control packet losses, we let switches proactively drop packets whose ECN fields are set to 1. In this way, we can flexibly specify any flow as a victim flow and control its packet loss rate. To avoid packet losses due to congestion, we set the size of every packet to 64 bytes regardless of its original size, so as to significantly reduce the traffic load in the network and eliminate congestion. Such operation does not change the number of packets of each flow, and thus has no impact on the behavior of ChameleMon.

Parameter settings: We set the epoch length to 50ms by default⁶. For the flow classifier, we set it to consist of an 8-bit counter array

⁶For some workloads that cannot run out in 50ms, we extend the epoch length appropriately.

and a 16-bit counter array. We set the number of 8-bit counters w_1 to 32768 and the number of 16-bit counters w_2 to 16384. For the upstream flow encoder and downstream flow encoder, we set them to consist of 3 bucket arrays for the highest memory efficiency. We set the number of buckets per array of the upstream flow encoder m_{uf} to 4096, and that of the downstream flow encoder m_{df} to 3072. For the healthy network state, we fix the minimum memory reserved for HL encoders to a 3-array FermatSketch with 512 buckets per array. For the ill network state, we fix the upstream HH, HL, LL encoders to a 3-array FermatSketch with 1024, 2560, and 512 buckets per array, respectively. The resource usage of ChameleMon is shown in Table 1. Please refer to Appendix D.1 for more details.

Table 1: Resources used by ChameleMon in Tofino.

Resource	Usage	Percentage
Exact Match Input xbar	353	22.98%
Ternary Match Input xbar	33	4.17%
VLIW Instructions	43	11.20%
Map RAM	102	17.71%
SRAM	130	13.54%
TCAM	8	2.78%
Hash Bits	809	16.21%
Stateful ALU	32	66.67%

First, on DCTCP workload, we evaluate whether ChameleMon can automatically shift measurement attention for different network states⁷. For experimental results on the other three workloads, please refer to Appendix E.

Measurement attention vs. number of flows (Figure 7): Experimental results show that ChameleMon can automatically shift measurement attention to packet loss detection while maintaining high memory utilization, as the number of flows increases and the number of victim flows increases. We vary the number of flows in the network from 10K to 100K, and fix the ratio of victim flows to 10%. At first, the network state is healthy. As the number of flows increases from 10K to 20K, ChameleMon can record all flows and victim flows, and therefore sets both T_h and T_l to 1. As the number of flows increases from 30K to 70K, ChameleMon records all victim flows by allocating more and more memory to HL encoders. However, ChameleMon cannot record all flows, and thus increases T_h to decrease the number of HH candidates. As the number of flows increases from 80K to 100K, ChameleMon cannot record all victim flows, and thus the network state transitions to the ill network state. ChameleMon allocates fixed memory to LL encoders, increases T_l , and decreases the sample rate, so as to control the number of HLs and sampled LLs. Meanwhile, ChameleMon keeps increasing T_h to control the number of HH candidates. Throughout the experiment, ChameleMon maintains high memory utilization. The sum of decoded flows (Figure 7(b)) always exceeds 8K unless ChameleMon can record all flows and victim flows, representing a load factor larger than 65% given that the upstream flow encoder has 12288 buckets. It is acceptable considering that the target load factor of ChameleMon is 70% and maximum load factor is $\frac{1}{1.23} = 81.3\%$.

Measurement attention vs. ratio of victim flows (Figure 8): Experimental results show that ChameleMon can automatically shift measurement attention to packet loss detection while maintaining

⁷For each data point of Figure 7-8, we collect it after ChameleMon successfully shifts measurement attention and the configuration of the ChameleMon data plane is stable.

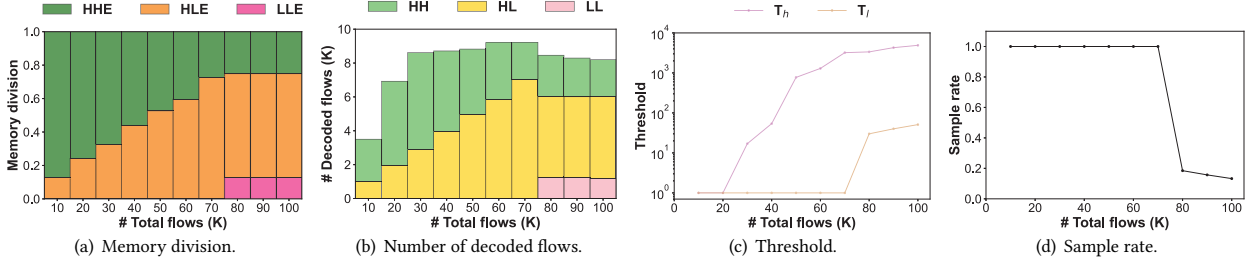


Figure 7: Measurement attention vs. number of flows. Figure 7(a) depicts the memory division of HH encoder (HHE), HL encoder (HLE), and LL encoder (LLE) inside the upstream flow encoder. Figure 7(b) depicts the number of HH candidates of an edge switch, the number of HLs in the network, and the number of sampled LLs in the network.

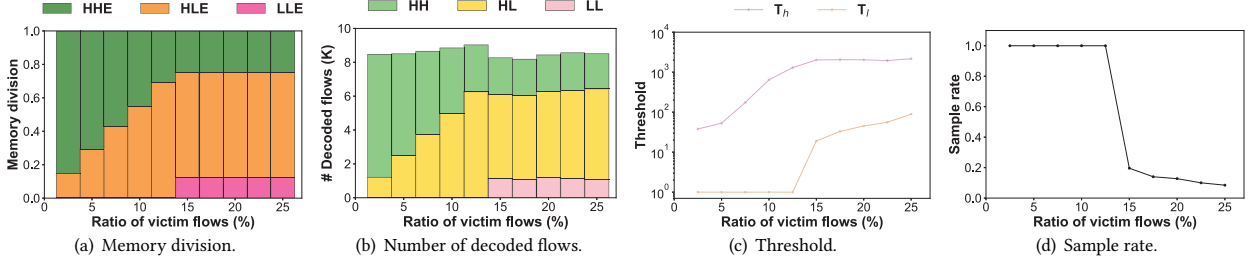


Figure 8: Measurement attention vs. ratio of victim flows.

high memory utilization, as the ratio of victim flows increases and the number of victim flows increases. We fix the number of flows to 50K, and vary the ratio of victim flows from 2.5% to 25%. At first, the network state is healthy. As the ratio of victim flows increases from 2.5% to 12.5%, ChameleMon records all victim flows by allocating more and more memory to HL encoders, and increases T_h to decrease the number of HH candidates. As the ratio of victim flows increases from 15% to 25%, ChameleMon cannot record all victim flows, and thus the network state transitions to the ill network state. ChameleMon allocates fixed memory to LL encoders, increases T_l , and decreases the sample rate, so as to control the number of HLs and sampled LLs. Meanwhile, because the memory of upstream HH encoder and the number of flows remain unchanged, T_h also remains unchanged. Throughout the experiment, ChameleMon maintains high memory utilization. The sum of decoded flows (Figure 8(b)) always exceeds 8K, representing a load factor larger than 65%. It is acceptable considering that the target load factor of ChameleMon is 70% and maximum load factor is 81.3%.

Second, on DCTCP workload, we evaluate how fast can ChameleMon shift measurement attention over a large time window, in which the network state changes 8 times.

Measurement attention vs. epoch (Figure 9): Experimental results show that ChameleMon can shift measurement attention within at most 3 epochs. Figure 9 plots the shift of measurement attention in a large time window consisting of 45 consecutive epochs. We change the network state (either the number of flows or the victim flow ratio) every 5 epochs, and the detailed settings are shown in the top sub-figure. Overall, the network state first degrades from the healthy network state to the ill network state, and then improves back to the healthy network state. For the eight changes, ChameleMon shifts measurement attention within 1 (6-→7), 2 (11-→13), 3 (16-→19), 2 (21-→23), 2 (26-→28), 1 (31-→32), 1 (36-→37), and 1 (41-→42) epochs, respectively.

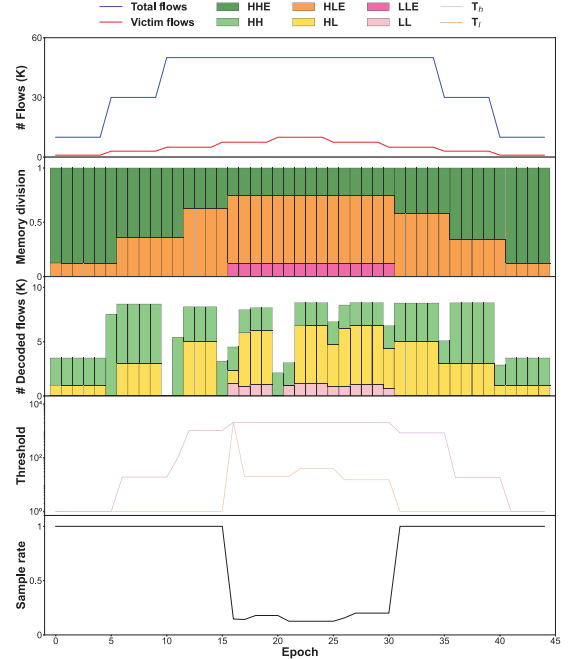


Figure 9: Measurement attention vs. epoch.

To evaluate how fast can ChameleMon monitor the network, we evaluate various factors that could affect the setting of epoch length: 1) the time and bandwidth required to collect sketches from edge switches, 2) the time required to respond to different network states, and 3) the time required to reconfigure the ChameleMon data plane. We detail the corresponding experimental results in Appendix F. In summary, we find that ChameleMon can monitor the network every 50ms with only one CPU core while consuming only 0.8% bandwidth of a 40Gb NIC. Thus, we believe ChameleMon can easily scale to monitor a much larger network with a shorter epoch length, requiring only one server as the central controller.

6 RELATED WORK

First, we discuss prior art for packet loss tasks and packet accumulation tasks. Then, we discuss prior art for resource management on switches.

Prior art for packet loss tasks: They can be classified into two kinds of solutions. The first kind is algorithm-oriented solutions, including LossRadar [24] based on Invertible Bloom filter (IBF) [25]. LossRadar can pinpoint the location of every lost packet and infer the root causes of packet losses by deploying IBF to monitor every link in the network. The second kind is system-oriented solutions, including Netseer [26], PacketScope [49] and Dapper [27] that are based on programmable switches. Among them, Netseer utilizes the programmable data plane to detect both intra-switch and inter-switch packet losses, associate packet losses with direct causes, and batch lost packets to further reduce bandwidth overhead. Both the above kinds of solutions are designed for only obtaining the exact difference set of flows/packets. Therefore, they fail to meet versatility requirement as they can hardly be extended to packet accumulation tasks that require approximate flow sizes of all flows or simply large flows.

Prior art for packet accumulation tasks: They can be classified into two kinds of solutions. The first kind is sketches designed for specific packet accumulation task, including HashPipe [22], R-HHH [50], and more [51–54]. Among them, HashPipe designs a multi-stage data structure and kicks out small flows through comparison. The second kind is sketches that support many packet accumulation tasks, including UnivMon [18], ElasticSketch [15], CocoSketch [20], and more [4, 16, 17, 19, 21, 55, 56]. Among them, CocoSketch proposes a key technique, namely stochastic variance minimization technique, to provide unbiased estimation for arbitrary partial key. Both the above kinds of solutions choose to embrace hash collisions and provide approximate flow sizes for higher memory efficiency. Therefore, they fail to meet versatility requirement as they can hardly be extended to packet loss tasks that require exact difference set of flows/packets.

Prior art for both kinds of tasks: These solutions record the IDs and sizes of all flows in a zero-error manner. Typical solutions include FlowRadar [28], OmniMon [29], Counter Braids [30], Marple [31], and more [57]. Among them, FlowRadar encodes the IDs and sizes of all flows into a variant of IBLT [35] in switches, and then executes well-designed decoding schemes to retrieve exact flow IDs and sizes. Marple designs a query language for a wide range of network measurement tasks, which relies on the programmable key-value store in switch hardware. Marple requires an additional backing store to handle evicted flows. These solutions fail to meet efficiency requirement as they record the exact IDs and sizes of all flows, incurring memory/bandwidth overhead linear with the number of flows. Besides, *INT-based solutions* that carry desired statistics in packet headers can potentially support both tasks given packet-level visibility. Typical solutions include INT [38], PINT [58], NetSight [9], and more [32, 59–62]. However, INT-based solutions suffer from granularity-cost trade-off, and thus fail to meet either versatility requirement or efficiency requirement.

Prior art for resource management: Due to the limited resources in hardware such as programmable switches, many solutions focus on resource management in measurement. Some solutions [63–65] target at compile-time resource management. Among them,

HeteroSketch [65] optimizes network-wide measurement by automatically optimizing the placement of sketches on heterogeneous devices. These solutions differ from ChameleMon as ChameleMon executes memory reallocation at run-time. Other solutions focus on run-time resource management [31, 66–70]. Among them, FlyMon [67] achieves run-time reconfiguration of measurement tasks and resources. However, these solutions does not focus on the resource management between packet loss tasks and packet accumulation tasks, which is our main focus.

Here, we further discuss other prior art relevant to network measurement.

Sampling-based solutions: These solutions collect desired statistics from a subset of network traffic through packet sampling, including Csamp [14], NetFlow [12], sFlow [13], EverFlow [11], and more [45, 71–78]. While sampling solutions significantly reduce the bandwidth overhead through sampling, they cannot well handle packet loss tasks as only sampled packets are measured, and thus fail to meet versatility requirement.

Programmable-switch-assisted solutions: Besides packet loss detection, some solutions leverage the advanced features and capabilities of programmable switches to monitor micro-bursts [79], perform queue measurement [80–82], and more [83–86].

Host-based solutions: Due to the flexibility, abundant resources, and high visibility to flow-level statistics of end-hosts, these solutions are typically designed for inferring the existences, locations, and root causes of specific network events or network failures. Typical solutions either send tailored probes into the network [10, 87–93] or analyze the performance of protocol stack or other I/O [94–103]. Besides, some solutions further leverage switches to perform measurement [104, 105] or record forwarding paths [106, 107]. ChameleMon can complement these solutions as ChameleMon provides flow-level statistics with high accuracy. Take 007 [102] as an instance. Network operators can replace the TCP monitoring agent that detects TCP retransmissions in 007 with ChameleMon. After the replacement, 007 can monitor packet losses of TCP flows as well as packet losses of flows of other protocols. Such extra visibility can help 007 better locate the link failures.

7 CONCLUSION

In this paper, we present ChameleMon, which can automatically shift measurement attention as network state changes at run-time. To achieve this, ChameleMon designs FermatSketch to support both packet loss tasks and packet accumulation tasks simultaneously. We have fully implemented a ChameleMon prototype on a testbed consisting of 10 Tofino switches and 8 end-hosts. Experimental results on our testbed verify that 1) ChameleMon can achieve high accuracy in packet loss detection and six packet accumulation tasks; 2) ChameleMon can monitor the network every 50ms and shift measurement attention within at most 3 epochs as network state changes.

ACKNOWLEDGEMENT

We would like to thank the anonymous reviewers for their valuable suggestions that helped improve the paper. This work is supported by National Key R&D Program of China (No. 2022YFB2901504), and National Natural Science Foundation of China (NSFC) (No. U20A20179).

REFERENCES

- [1] Theophilus Benson, Ashok Anand, Aditya Akella, and Ming Zhang. Microte: Fine grained traffic engineering for data centers. In *Proceedings of the seventh conference on emerging networking experiments and technologies*, pages 1–12, 2011.
- [2] Anja Feldmann, Albert Greenberg, Carsten Lund, Nick Reingold, Jennifer Rexford, and Fred True. Deriving traffic demands for operational ip networks: Methodology and experience. *IEEE/ACM Transactions On Networking*, 9(3):265–279, 2001.
- [3] Yuliang Li, Rui Miao, Hongqiang Harry Liu, Yan Zhuang, Fei Feng, Lingbo Tang, Zheng Cao, Ming Zhang, Frank Kelly, Mohammad Alizadeh, et al. Hppc: High precision congestion control. In *Proceedings of the ACM Special Interest Group on Data Communication*, pages 44–58, 2019.
- [4] Cristian Estan and George Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems (TOCS)*, 21(3):270–313, 2003.
- [5] Ying Zhang. An adaptive flow counting method for anomaly detection in sdn. In *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, pages 25–30, 2013.
- [6] Jianning Mai, Chen-Nee Chuah, Ashwin Sridharan, Tao Ye, and Hui Zang. Is sampled data sufficient for anomaly detection? In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 165–176, 2006.
- [7] Cristian Estan, Ken Keys, David Moore, and George Varghese. Building a better netflow. *ACM SIGCOMM Computer Communication Review*, 34(4):245–256, 2004.
- [8] Nick Duffield, Carsten Lund, and Mikkel Thorup. Estimating flow distributions from sampled flow statistics. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 325–336, 2003.
- [9] Nikhil Handigol, Brandon Heller, Vimalkumar Jeyakumar, David Mazières, and Nick McKeown. I know what your packet did last hop: Using packet histories to troubleshoot networks. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, pages 71–85, 2014.
- [10] Cheng Tan, Ze Jin, Chuanxiong Guo, Tianrong Zhang, Haitao Wu, Karl Deng, Dongming Bi, and Dong Xiang. Netbouncer: Active device and link failure localization in data center networks. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 599–614, 2019.
- [11] Yibo Zhu, Nanxi Kang, Jiabin Cao, Albert Greenberg, Guohan Lu, Ratul Mahajan, Dave Maltz, Lihua Yuan, Ming Zhang, Ben Y Zhao, et al. Packet-level telemetry in large datacenter networks. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 479–491, 2015.
- [12] Benoit Claise, Ganesh Sadasivan, Vamsi Valluri, and Martin Djernaes. Cisco systems netflow services export version 9. 2004.
- [13] Peter Phaal, Sonia Panchen, and Neil McKee. Rfc3176: Inmon corporation's sflow: A method for monitoring traffic in switched and routed networks, 2001.
- [14] Vyas Sekar, Michael K Reiter, Walter Willinger, Hui Zhang, Ramana Rao Kompella, and David G Andersen. csamp: A system for network-wide flow monitoring. 2008.
- [15] Tong Yang, Jie Jiang, Peng Liu, Qun Huang, Junzhi Gong, Yang Zhou, Rui Miao, Xiaoming Li, and Steve Uhlig. Elastic sketch: Adaptive and fast network-wide measurements. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 561–575, 2018.
- [16] Zaoxing Liu, Ran Ben-Basat, Gil Einziger, Yaron Kassner, Vladimir Braverman, Roy Friedman, and Vyas Sekar. Nitrosketch: Robust and general sketch-based monitoring in software switches. In *Proceedings of the ACM Special Interest Group on Data Communication*, pages 334–350, 2019.
- [17] Qun Huang, Patrick PC Lee, and Yungang Bao. Sketchlearn: relieving user burdens in approximate measurement with automated statistical inference. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 576–590, 2018.
- [18] Zaoxing Liu, Antonis Manousis, Gregory Vorsanger, Vyas Sekar, and Vladimir Braverman. One sketch to rule them all: Rethinking network flow monitoring with univmon. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 101–114, 2016.
- [19] Xiaoqi Chen, Shir Landau-Feibish, Mark Braverman, and Jennifer Rexford. Beau-coup: Answering many network traffic queries, one memory update at a time. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 226–239, 2020.
- [20] Yinda Zhang, Zaoxing Liu, Ruixin Wang, Tong Yang, Jizhou Li, Ruijie Miao, Peng Liu, Ruwen Zhang, and Junchen Jiang. Cocosketch: high-performance sketch-based measurement over arbitrary partial key query. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 207–222, 2021.
- [21] Graham Cormode and Shan Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [22] Vibhaalakshmi Sivaraman, Srinivas Narayana, Ori Rottenstreich, Shan Muthukrishnan, and Jennifer Rexford. Heavy-hitter detection entirely in the data plane. In *Proceedings of the Symposium on SDN Research*, pages 164–176, 2017.
- [23] Yu Gu, Andrew McCallum, and Don Towsley. Detecting anomalies in network traffic using maximum entropy estimation. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, pages 32–32, 2005.
- [24] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. Lossradar: Fast detection of lost packets in data center networks. In *Proceedings of the 12th International Conference on emerging Networking EXperiments and Technologies*, pages 481–495, 2016.
- [25] David Eppstein, Michael T Goodrich, Frank Uyeda, and George Varghese. What's the difference?: efficient set reconciliation without prior context. *ACM SIGCOMM Computer Communication Review*, 41(4):218–229, 2011.
- [26] Yu Zhou, Chen Sun, Hongqiang Harry Liu, Rui Miao, Shi Bai, Bo Li, Zhilong Zheng, Lingjun Zhu, Zhen Shen, Yongqing Xi, et al. Flow event telemetry on programmable data plane. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 76–89, 2020.
- [27] Mojgan Ghasemi, Theophilus Benson, and Jennifer Rexford. Dapper: Data plane performance diagnosis of tcp. In *Proceedings of the Symposium on SDN Research*, pages 61–74, 2017.
- [28] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. Flowradar: A better netflow for data centers. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 311–324, 2016.
- [29] Qun Huang, Haifeng Sun, Patrick PC Lee, Wei Bai, Feng Zhu, and Yungang Bao. Omnimon: Re-architecting network telemetry with resource efficiency and full accuracy. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 404–421, 2020.
- [30] Yi Lu, Andrea Montanari, Balaji Prabhakar, Sarang Dharmapurikar, and Abdul Kabbani. Counter braids: a novel counter architecture for per-flow measurement. *ACM SIGMETRICS Performance Evaluation Review*, 36(1):121–132, 2008.
- [31] Srinivas Narayana, Anirudh Sivaraman, Vikram Nathan, Prateesh Goyal, Venkat Arun, Mohammad Alizadeh, Vimalkumar Jeyakumar, and Changhoon Kim. Language-directed hardware design for network performance monitoring. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 85–98, 2017.
- [32] Kaicheng Yang, Yuanpeng Li, Zirui Liu, Tong Yang, Yu Zhou, Jintao He, Tong Zhao, Zhengyi Jia, Yongqiang Yang, et al. Sketchint: Empowering int with tow-ersketch for per-flow per-switch measurement. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*, pages 1–12. IEEE, 2021.
- [33] Cha Hwan Song, Pravein Govindan Kannan, Bryan Kian Hsiang Low, and Mun Choon Chan. Fcm-sketch: generic network measurements with data plane support. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*, pages 78–92, 2020.
- [34] Source code of ChameleMon. <https://github.com/ChameleMoncode/ChameleMon>.
- [35] Michael T Goodrich and Michael Mitzenmacher. Invertible bloom lookup tables. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 792–799. IEEE, 2011.
- [36] Martin Dietzfelbinger, Andreas Goerdt, Michael Mitzenmacher, Andrea Montanari, Rasmus Pagh, and Michael Rink. Tight thresholds for cuckoo hashing via xorst. In *International Colloquium on Automata, Languages, and Programming*, pages 213–225. Springer, 2010.
- [37] Michael Molloy. The pure literal rule threshold and cores in random hypergraphs. 2004.
- [38] Changhoon Kim, Anirudh Sivaraman, Naga Katta, Antonin Bas, Advait Dixit, and Lawrence J Wobker. In-band network telemetry via programmable data-planes. In *SIGCOMM*, 2015.
- [39] Kyu-Young Whang, Brad T Vander-Zanden, and Howard M Taylor. A linear-time probabilistic counting algorithm for database applications. *ACM Transactions on Database Systems (TODS)*, 15(2):208–229, 1990.
- [40] Abhishek Kumar, Minh Sung, Jun Xu, and Jia Wang. Data streaming algorithms for efficient and accurate estimation of flow size distribution. *ACM SIGMETRICS Performance Evaluation Review*, 32(1):177–188, 2004.
- [41] The CAIDA Anonymized Internet Traces. <http://www.caida.org/data/overview/>.
- [42] Mohammad Alizadeh, Albert Greenberg, David A Maltz, Jitendra Padhye, Parveen Patel, Balaji Prabhakar, Sudipta Sengupta, and Murari Sridharan. Data center tcp (dctcp). In *Proceedings of the ACM SIGCOMM 2010 conference*, pages 63–74, 2010.
- [43] Burton H Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [44] Pat Bosshart, Dan Daly, Glen Gibb, Martin Izzard, Nick McKeown, Jennifer Rexford, Cole Schlesinger, Dan Talayco, Amin Vahdat, George Varghese, et al. P4: Programming protocol-independent packet processors. *ACM SIGCOMM Computer Communication Review*, 44(3):87–95, 2014.
- [45] Arjun Roy, Hongyi Zeng, Jasmeet Bagga, George Porter, and Alex C Snoeren. Inside the social network's (datacenter) network. In *Proceedings of the 2015 ACM*

- Conference on Special Interest Group on Data Communication*, pages 123–137, 2015.
- [46] Albert Greenberg, James R Hamilton, Navendu Jain, Srikanth Kandula, Changhoon Kim, Parantap Lahiri, David A Maltz, Parveen Patel, and Sudipta Sengupta. VI2: A scalable and flexible data center network. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, pages 51–62, 2009.
- [47] Berk Atikoglu, Yuehai Xu, Eitan Frachtenberg, Song Jiang, and Mike Paleczny. Workload analysis of a large-scale key-value store. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, pages 53–64, 2012.
- [48] Data plane development kit. <https://www.dpdk.org/>.
- [49] Ross Teixeira, Rob Harrison, Arpit Gupta, and Jennifer Rexford. Packetscope: Monitoring the packet lifecycle inside a switch. In *Proceedings of the Symposium on SDN Research*, pages 76–82, 2020.
- [50] Ran Ben Basat, Gil Einziger, Roy Friedman, Marcelo C Luizelli, and Erez Waisbard. Constant time updates in hierarchical heavy hitters. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 127–140, 2017.
- [51] Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Heavy hitters in streams and sliding windows. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*, pages 1–9. IEEE, 2016.
- [52] Ran Ben-Basat, Xiaoqi Chen, Gil Einziger, and Ori Rottenstreich. Efficient measurement on programmable switches using probabilistic recirculation. In *2018 IEEE 36th International Conference on Network Protocols (ICNP)*, pages 313–323. IEEE, 2018.
- [53] Tong Yang, Haowei Zhang, Jinyang Li, Junzhi Gong, Steve Uhlig, Shigang Chen, and Xiaoming Li. Heavykeeper: An accurate algorithm for finding top-k elephant flows. *IEEE/ACM Transactions on Networking*, 27(5):1845–1858, 2019.
- [54] Robert Schweller, Ashish Gupta, Elliot Parsons, and Yan Chen. Reversible sketches for efficient and accurate change detection over network data streams. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*, pages 207–212, 2004.
- [55] Qun Huang, Xin Jin, Patrick PC Lee, Runhui Li, Lu Tang, Yi-Chao Chen, and Gong Zhang. Sketchvisor: Robust network measurement for software packet processing. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pages 113–126, 2017.
- [56] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002.
- [57] Qun Huang, Siyuan Sheng, Xiang Chen, Yungang Bao, Rui Zhang, Yanwei Xu, and Gong Zhang. Toward {Nearly-Zero-Error} sketching via compressive sensing. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 1027–1044, 2021.
- [58] Ran Ben Basat, Sivaramakrishnan Ramanathan, Yuliang Li, Gianni Antichi, Minian Yu, and Michael Mitzenmacher. Pint: Probabilistic in-band network telemetry. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 662–680, 2020.
- [59] Yikai Zhao, Kaicheng Yang, Zirui Liu, Tong Yang, Li Chen, Shiyi Liu, Naiqian Zheng, Ruixin Wang, Hanbo Wu, Yi Wang, et al. {LightGuardian}: A {Full-Visibility}, lightweight, in-band telemetry system using sketches. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21)*, pages 991–1010, 2021.
- [60] Siyuan Sheng, Qun Huang, and Patrick PC Lee. Deltaint: Toward general in-band network telemetry with extremely low bandwidth overhead. In *2021 IEEE 29th International Conference on Network Protocols (ICNP)*, pages 1–11. IEEE, 2021.
- [61] Tal Mizrahi, Gidi Navon, Giuseppe Fioccola, Mauro Cociglio, Mach Chen, and Greg Mirsky. Am-pm: Efficient network telemetry using alternate marking. *IEEE Network*, 33(4):155–161, 2019.
- [62] John Sonchack, Oliver Michel, Adam J Aviv, Eric Keller, and Jonathan M Smith. Scaling hardware accelerated network monitoring to concurrent and dynamic queries with* flow. In *2018 {USENIX} Annual Technical Conference ({USENIX} {ATC} 18)*, pages 823–835, 2018.
- [63] Mary Hogan, Shir Landau-Feibish, Mina Tahmasbi Arashloo, Jennifer Rexford, and David Walker. Modular switch programming under resource constraints. In *USENIX NSDI*, pages 1–15, 2022.
- [64] Minlan Yu, Lavanya Jose, and Rui Miao. Software defined traffic measurement with opensketch. In *10th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 13)*, pages 29–42, 2013.
- [65] Anup Agarwal, Zaoxing Liu, and Srinivasan Seshan. {HeteroSketch}: Coordinating network-wide monitoring in heterogeneous and dynamic networks. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 719–741, 2022.
- [66] Chris Misa, Walt O'Connor, Ramakrishnan Durairajan, Reza Rejaie, and Walter Willinger. Dynamic scheduling of approximate telemetry queries. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 701–717, 2022.
- [67] Hao Zheng, Chen Tian, Tong Yang, Huiping Lin, Chang Liu, Zhaochen Zhang, Wanchun Dou, and Guihai Chen. Flymon: enabling on-the-fly task reconfiguration for network measurement. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 486–502, 2022.
- [68] Arpit Gupta, Rob Harrison, Marco Canini, Nick Feamster, Jennifer Rexford, and Walter Willinger. Sonata: Query-driven streaming network telemetry. In *Proceedings of the 2018 conference of the ACM special interest group on data communication*, pages 357–371, 2018.
- [69] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. Dream: dynamic resource allocation for software-defined measurement. In *Proceedings of the 2014 ACM conference on SIGCOMM*, pages 419–430, 2014.
- [70] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. Scream: Sketch resource allocation for software-defined measurement. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, pages 1–13, 2015.
- [71] Chris Hare. Simple network management protocol (snmp)., 2011.
- [72] Nick G Duffield and Matthias Grossglauser. Trajectory sampling for direct traffic observation. *IEEE/ACM transactions on networking*, 9(3):280–292, 2001.
- [73] Vyas Sekar, Michael K Reiter, and Hui Zhang. Revisiting the case for a minimalist approach for network flow monitoring. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pages 328–341, 2010.
- [74] Junho Suh, Ted Taekyoung Kwon, Colin Dixon, Wes Felter, and John Carter. Opensample: A low-latency, sampling-based measurement platform for commodity sdn. In *2014 IEEE 34th International Conference on Distributed Computing Systems*, pages 228–237. IEEE, 2014.
- [75] Fangfan Li, Arian Akhavan Niaki, David Choffnes, Phillippa Gill, and Alan Mislove. A large-scale analysis of deployed traffic differentiation practices. In *Proceedings of the ACM Special Interest Group on Data Communication*, pages 130–144, 2019.
- [76] Jeff Rasley, Brent Stephens, Colin Dixon, Eric Rozner, Wes Felter, Kanak Agarwal, John Carter, and Rodrigo Fonseca. Planck: Millisecond-scale monitoring and control for commodity networks. *ACM SIGCOMM Computer Communication Review*, 44(4):407–418, 2014.
- [77] Pavlos Nikolopoulos, Christos Pappas, Katerina Argyraki, and Adrian Perrig. Retroactive packet sampling for traffic receipts. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(1):1–39, 2019.
- [78] Da Yu, Yibo Zhu, Behnaz Arzani, Rodrigo Fonseca, Tianrong Zhang, Karl Deng, and Lihua Yuan. {dShark}: A general, easy to program and scalable framework for analyzing in-network packet traces. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 207–220, 2019.
- [79] Raj Joshi, Ting Qu, Mun Choon Chan, Ben Leong, and Boon Thau Loo. Bursttrader: Practical real-time microburst monitoring for datacenter networks. In *Proceedings of the 9th Asia-Pacific Workshop on Systems*, pages 1–8, 2018.
- [80] Xiaoqi Chen, Shir Landau Feibish, Yaron Koral, Jennifer Rexford, Ori Rottenstreich, Steven A Monetti, and Tzoo-Yi Wang. Fine-grained queue measurement in the data plane. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*, pages 15–29, 2019.
- [81] Yiran Lei, Liangcheng Yu, Vincent Liu, and Mingwei Xu. Printqueue: performance diagnosis via queue measurement in the data plane. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 516–529, 2022.
- [82] John Sonchack, Adam J Aviv, Eric Keller, and Jonathan M Smith. Turboflow: Information rich flow record generation on commodity switches. In *Proceedings of the Thirteenth EuroSys Conference*, pages 1–16, 2018.
- [83] Abir Laraba, Jérôme François, Shihabur Rahman Chowdhury, Isabelle Chrisment, and Raouf Boutaba. Mitigating tcp protocol misuse with programmable data planes. *IEEE Transactions on Network and Service Management*, 18(1):760–774, 2021.
- [84] Weitao Wang, Xinyu Crystal Wu, Praveen Tammana, Ang Chen, and TS Eugene Ng. Closed-loop network performance monitoring and diagnosis with {SpiderMon}. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 267–285, 2022.
- [85] Edgar Costa Molero, Stefano Vissicchio, and Laurent Vanbever. Fast in-network gray failure detection for isps. In *Proceedings of the ACM SIGCOMM 2022 Conference*, pages 677–692, 2022.
- [86] Thomas Holterbach, Edgar Costa Molero, Maria Apostolaki, Alberto Dainotti, Stefano Vissicchio, and Laurent Vanbever. Blink: Fast connectivity recovery entirely in the data plane. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 161–176, 2019.
- [87] Chuanxiong Guo, Lihua Yuan, Dong Xiang, Yingnong Dang, Ray Huang, Dave Maltz, Zhaoyi Liu, Vin Wang, Bin Pang, Hua Chen, et al. Pingmesh: A large-scale system for data center network latency measurement and analysis. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 139–152, 2015.
- [88] Amogh Dhamdhere, David D Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky KP Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C Snoeren, and Kc Claffy. Inferring persistent interdomain congestion. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 1–15, 2018.

- [89] Aijay Adams, Petr Lapukhov, and J Hongyi Zeng. Netnorad: Troubleshooting networks via end-to-end probing. *Facebook White Paper*, 2016.
- [90] François Aubry, David Lebrun, Stefano Vissicchio, Minh Thanh Khong, Yves Deville, and Olivier Bonaventure. Scmon: Leveraging segment routing to improve network monitoring. In *IEEE INFOCOM 2016-The 35th Annual IEEE International Conference on Computer Communications*, pages 1–9. IEEE, 2016.
- [91] Amogh Dhamdhare, Renata Teixeira, Constantine Dovrolis, and Christophe Diot. Netdiagnoser: Troubleshooting network unreachabilities using end-to-end probes and routing data. In *Proceedings of the 2007 ACM CoNEXT conference*, pages 1–12, 2007.
- [92] Yanghua Peng, Ji Yang, Chuan Wu, Chuanxiong Guo, Chengchen Hu, and Zongpeng Li. {deTector}: a topology-aware monitoring system for data center networks. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, pages 55–68, 2017.
- [93] Yilong Geng, Shiyu Liu, Zi Yin, Ashish Naik, Balaji Prabhakar, Mendel Rosenblum, and Amin Vahdat. {SIMON}: A simple and scalable method for sensing, inference and measurement in data center networks. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 549–564, 2019.
- [94] Arjun Roy, Hongyi Zeng, Jasmeet Bagga, and Alex C Snoeren. Passive realtime datacenter fault detection and localization. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*, pages 595–612, 2017.
- [95] David R Choffnes, Fabián E Bustamante, and Zihui Ge. Crowdsourcing service-level network event monitoring. In *Proceedings of the ACM SIGCOMM 2010 Conference*, pages 387–398, 2010.
- [96] Radhika Niranjan Mysore, Ratul Mahajan, Amin Vahdat, and George Varghese. Gestalt: Fast, {Unified} fault localization for networked systems. In *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, pages 255–267, 2014.
- [97] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. Trumpet: Timely and precise triggers in data centers. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 129–143, 2016.
- [98] Behnaz Arzani, Selim Ciraci, Boon Thau Loo, Assaf Schuster, and Geoff Outhred. Taking the blame game out of data centers operations with netpoirot. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 440–453, 2016.
- [99] Anurag Khandelwal, Rachit Agarwal, and Ion Stoica. Confluo: Distributed monitoring and diagnosis stack for high-speed networks. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 421–436, 2019.
- [100] Yang Wu, Ang Chen, and Linh Thi Xuan Phan. Zeno: Diagnosing performance problems with temporal provenance. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*, pages 395–420, 2019.
- [101] Yifeng Yuan, Dong Lin, Ankit Mishra, Sajal Marwaha, Rajeev Alur, and Boon Thau Loo. Quantitative network monitoring with netqre. In *Proceedings of the conference of the ACM special interest group on data communication*, pages 99–112, 2017.
- [102] Behnaz Arzani, Selim Ciraci, Luiz Chamon, Yibo Zhu, Hongqiang Harry Liu, Jitu Padhye, Boon Thau Loo, and Geoff Outhred. 007: Democratically finding the cause of packet drops. In *15th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 18)*, pages 419–435, 2018.
- [103] Junzhi Gong, Yuliang Li, Bilal Anwer, Aman Shaikh, and Minlan Yu. Microscope: Queue-based performance diagnosis for network functions. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*, pages 390–403, 2020.
- [104] Vimalkumar Jeyakumar, Mohammad Alizadeh, Yilong Geng, Changhoon Kim, and David Mazieres. Millions of little minions: Using packets for low latency network programming and visibility. *ACM SIGCOMM Computer Communication Review*, 44(4):3–14, 2014.
- [105] Xuemei Liu, Meral Shirazipour, Minlan Yu, and Ying Zhang. Mozart: Temporal coordination of measurement. In *Proceedings of the Symposium on SDN Research*, pages 1–12, 2016.
- [106] Praveen Tammana, Rachit Agarwal, and Myungjin Lee. Simplifying datacenter network debugging with {PathDump}. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 233–248, 2016.
- [107] Praveen Tammana, Rachit Agarwal, and Myungjin Lee. Distributed network monitoring and debugging with {SwitchPointer}. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 453–456, 2018.
- [108] Gaoxiong Zeng, Wei Bai, Ge Chen, Kai Chen, Dongsu Han, Yibo Zhu, and Lei Cui. Congestion control for cross-datacenter networks. In *2019 IEEE 27th International Conference on Network Protocols (ICNP)*, pages 1–12. IEEE, 2019.
- [109] Amit Goyal, Hal Daumé III, and Graham Cormode. Sketch algorithms for estimating point queries in nlp. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, pages 1093–1103, 2012.
- [110] David L Mills. Internet time synchronization: the network time protocol. *IEEE Transactions on communications*, 39(10):1482–1493, 1991.
- [111] Yilong Geng, Shiyu Liu, Zi Yin, Ashish Naik, Balaji Prabhakar, Mendel Rosenblum, and Amin Vahdat. Exploiting a natural network effect for scalable, fine-grained clock synchronization. In *15th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 18)*, pages 81–94, 2018.
- [112] Pravein Govindan Kannan, Raj Joshi, and Mun Choon Chan. Precise time-synchronization in the data-plane using programmable switching asics. In *Proceedings of the 2019 ACM Symposium on SDN Research*, pages 8–20, 2019.
- [113] Barefoot p4 studio. <https://www.barefootnetworks.com/products/brief-p4-studio>.
- [114] Hun Namkung, Daehyeok Kim, Zaoxing Liu, Vyas Sekar, and Peter Steenkiste. Telemetry retrieval inaccuracy in programmable switches: Analysis and recommendations. In *Proceedings of the ACM SIGCOMM Symposium on SDN Research (SOSR)*, pages 176–182, 2021.

APPENDIX

Appendices are supporting material that has not been peer-reviewed.

A THE FERMATSKETCH ALGORITHM

A.1 Proof of Theorem A.1

THEOREM A.1. *Let FermatSketch consist of d bucket arrays, each of which consists of m buckets. Let M be the number of flows inserted into that FermatSketch. Suppose $md > c_d M + \epsilon$ and $M = \Omega(d^{4d} \log^d(md))$. the decoding of FermatSketch fails with probability $O(\frac{1}{M^{d-2}})$, where both ϵ and c_d are small constants,*

$$c_d = \left(\sup \left\{ \alpha \mid \alpha \in (0, 1), \forall x \in (0, 1), 1 - e^{-d\alpha x^{d-1}} \right\} \right)^{-1}$$

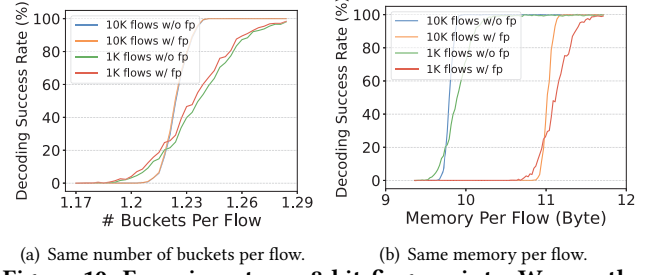
For example, $c_3 = 1.23$, $c_4 = 1.30$, $c_5 = 1.43$.

PROOF. This is an analysis based on the theory of the 2-core in random hypergraph [36, 37] and IBLT. Compared with 2-core or IBLT, we only introduce a kind of additional error, which is the false positives when we use pure bucket verification to verify the pure buckets. The IBLT assumes there is no error when verifying buckets because IBLT uses additional hashkeySum field that can be long enough. The results of 2-core and IBLT show that the failure probability without wrong verification is $O(\frac{1}{M^{d-2}})$. Here we aim at proving that the consequences of our false positives are negligible when M is not too small.

In the decoding procedure, the pure bucket verification runs at most $O(Md)$ times, and the false positive rate is $O(\frac{1}{m})$ with only rehashing verification. By Chernoff bound, when $M = O(md)$ and $\delta = O(\frac{1}{M^{d-2}})$, the number of false positives will not exceed $F = O(d^3 \log(md))$ in most cases (i.e., $1 - O(\delta)$). A false positive will incur a wrong flow ID with a wrong single flow deletion that influences d buckets. There is at most Fd buckets can be influenced, called poisoned buckets. The existing study [35] of poisoned bucket shows that a small number of poisoned buckets will be automatically recovered, and the probability of failure due to poisoned bucket is $O(\left(\frac{Fd}{M}\right)^d) = O(\frac{d^{4d} \log^d(md)}{M^{d-1}})$. When it satisfies that $M = \Omega(d^{4d} \log^d(md))$, the overall failure probability is $\delta = O(\frac{1}{M^{d-2}})$. In practice, $M = \Omega(d^{4d} \log^d(md))$ is easy to meet because M is large and d is a small constant. Here, we only use rehashing verification for pure bucket verification. The theorem can also be easily extended if we further use fingerprint verification. \square

A.2 Fingerprint Verification

To reduce the false positive rate of pure bucket verification, we can perform an extra verification method, namely *fingerprint verification*, by extending the IDsum field in each bucket by w bits and using the extra w bits as a fingerprint. For each incoming packet with flow ID f , a new hash function $h_{fp}(\cdot)$ gives it a w -bit fingerprint $h_{fp}(f)$ for checking whether a bucket is pure. For encoding operation, instead of inserting flow ID f , we insert an extended ID concatenated by flow ID f and fingerprint $h_{fp}(f)$, and the extended IDsum field stores the result of the sum of the extended IDs modulo prime p . Note that p must be a prime larger than any available extended ID. For decoding operation, obviously, we can still perform rehashing verification with the extended ID. Our fingerprint



(a) Same number of buckets per flow. (b) Same memory per flow.
Figure 10: Experiments on 8-bit fingerprints. We use the anonymized IP traces collected in 2018 from CAIDA [41] as dataset. We use the 32-bit source IP address as the flow ID, and choose the first 10K flows for experiments. Here, fp represents 8-bit fingerprint.

verification works as follows. Suppose a bucket is pure. First, we reuse the the extended ID of the single flow calculated in rehashing verification. Then, we divide the extended ID to get the flow ID and its fingerprint. If the divided fingerprint equals to the fingerprint of the divided flow ID, we consider the bucket passes fingerprint verification. Only buckets pass both rehashing and fingerprint verification will be considered as pure. The false positive rate of only fingerprint verification is obviously $\frac{1}{2^w}$. Considering that rehashing verification and fingerprint verification are independent, the false positive rate of pure bucket verification could be reduced to $\frac{1}{2^w m}$ with w -bit fingerprint.

We conduct experiments to demonstrate the effect of 8-bit fingerprint on improving the decoding success rate. As shown in Figure 10(a), when the number of flows is 1K, with the same number of buckets, 8-bit fingerprint can improve the decoding success rate by at most 6.73%. However, when the number of flows is 10K, the improvement falls to at most 2.26%. This is because as the number of buckets increases, m increases, and the false positive rate of pure bucket verification quickly drops, and thus further reducing the false positive rate with fingerprint yields less improvement on the decoding success rate. As shown in Figure 10(b), under the same memory usage, 8-bit fingerprint actually reduces the decoding success rate. This is because fingerprint consumes additional memory, while this memory could be used as buckets to reduce the probability of 2-core of the random hypergraph and improve the decoding success rate. Figure 10(a)-(b) also demonstrate that the memory overhead of FermatSketch is proportional to the number of inserted flows.

In summary, for simplicity and accuracy, we recommend implementing FermatSketch without fingerprints in most cases. Only if there is some memory that can hardly be utilized due to hardware constraints unless used as fingerprints, we recommend implementing FermatSketch with fingerprints.

B COLLECTION FROM DATA PLANE

Timeline split: For each edge switch, we maintain a 1-bit timestamp in its ingress, which is periodically flipped by the switch control plane, so as to split the timeline into consecutive fixed-length epochs with interleaved 0/1 timestamp value. Further, we

copy an additional group of sketches in the switch data plane for rotation. Each group of sketches corresponds to a distinct timestamp value (0/1), and monitors the epochs with that timestamp value. Specifically, at each edge switch, every packet entering the network first obtains the current timestamp value, and then is inserted into the flow classifier and upstream flow encoder corresponding to the obtained timestamp value. When the packet exits the network, it is also inserted into the downstream flow encoder corresponding to the timestamp value it obtained when it entered the network. To maintain the timestamp value during the packet transmission, we can use one unused bit in the original packet header as discussed above (§ 3.2.3).

Clock synchronization: Through maintaining a 1-bit timestamp and copying a group of sketches, we successfully split the timeline and insert packets of different epochs to their corresponding groups of sketches, laying a solid foundation for subsequent collection. However, if the clocks of the control planes of edge switches are out of synchronization to some extent, we still can not find opportunities to collect the sketches without colliding with packet insertion. Considering such an extreme situation. There are three edge switches in a given network, and the transmission time between any two edge switches is the same. The time offset between the control planes of two of the edge switches is exactly the size of the epoch, *i.e.*, at any time, the flipping timestamps of the two edge switches are different ($0 < \delta < 1$). There are continuous packets entering the network at the above two edge switches, and exiting the network at the third switch. As a result, both groups of sketches of the third switch are continuously inserted, and thus can never be collected. To address this issue, the central controller synchronizes the clocks of the control planes of all edge switches with itself, trying to keep only a group of sketches being inserted at any time, so as to make opportunities to collect the other group.

Then, we further discuss the appropriate time for the central controller to collect sketches.

Appropriate time for collection: The central controller also maintains a 1-bit timestamp, trying to collect the group of sketches monitoring the previous epoch after it ends. Before collection, the central controller should ensure that all the packets in the previous epoch have been inserted into sketches or lost in the network, so as to guarantee the correctness of measurement. First, we analyze an ideal situation, that the clock synchronization is zero-error. For ingress sketches, *i.e.*, the flow classifier and the upstream flow encoder, as soon as the locally maintained 1-bit timestamp flips, the central controller can collect the group of ingress sketches monitoring the previous epoch from each edge switch, because all the packets in that epoch have already been inserted into ingress sketches. For egress sketches, *i.e.*, the downstream flow encoder, every time the locally maintained timestamp flips, the central controller must first wait an additional period of time, so as to ensure that all the packets in the previous epoch have either been lost in the network, or passed through the network and been inserted into egress sketches. Then, the central controller can collect the group of egress sketches monitoring the previous epoch from each edge switch. Obviously, the additional period of time should be longer than the maximum time for packet transmission in the network. Considering that the buffer sizes of DCN switches are at 10MB-level [108], with 100Gb link speed, the queuing delay in a single switch

is at most 1ms in most cases. Therefore, for typical data center networks that usually have at most five hops, setting the additional time to 10ms can cope with most cases. However, in practice, the clock synchronization can never be zero-error. Therefore, before collecting both ingress and egress sketches, the central controller needs to wait for another additional period of time, which should be longer than the precision of synchronization, so as to guarantee the correctness of measurement. In addition, the central controller should end the collection some time before its 1-bit timestamp flips again, which should also be longer than the precision of synchronization, in case the packets in the next epoch are inserted into the group of sketches being collected.

C EVALUATION ON PACKET ACCUMULATION TASKS

Metrics:

- **Average Relative Error (ARE):** $\frac{1}{|\Omega|} \sum_{f_i \in \Omega} \frac{|v_i - \hat{n}_i|}{v_i}$, where Ω is the set including all flows, v_i is the true size of flow f_i , and \hat{v}_i is the estimated size of flow f_i .
- **F1 Score:** $\frac{2 \cdot PR \cdot RR}{PR + RR}$, where *PR* (Precision Rate) refers to the ratio of the number of the correctly reported instances to the number of all reported instances, and *RR* (Recall Rate) refers to the ratio of the number of the correctly reported instances to the number of all correct instances.
- **Relative Error (RE):** $\frac{|True - Est|}{True}$, where *True* and *Est* are the true and estimated statistics, respectively.
- **Weighted Mean Relative Error (WMRE) [55]:** $\frac{\sum_{i=1}^z |n_i - \hat{n}_i|}{\sum_{i=1}^z \left(\frac{n_i + \hat{n}_i}{2}\right)}$, where z

is the maximum flow size, n_i and \hat{n}_i are the true and estimated numbers of the flows of size i , respectively.

Dataset: We also use the IP traces from CAIDA [41] as our dataset, and use the 32-bit source IP address as the flow ID. We use four traces for evaluation, each of which monitors the traffic in five seconds. Each trace contains 63K flows and 2.3M packets on average. We report the average accuracy that each algorithm achieves on each CAIDA trace.

Setup: We compare the combination of TowerSketch and FermatSketch (Tower+Fermat) with 9 algorithms: CM [21], CU[4], CountHeap [56], UnivMon [18], ElasticSketch [15], FCM-sketch [33], HashPipe [22], CocoSketch [20], and MRAC [40]. We do not compare with FlowRadar because FlowRadar can hardly perform successful decoding with the memory sizes we used for evaluation (200KB-600KB). For heavy-hitter detection and heavy-change detection, we set their thresholds Δ_h and Δ_c to about 0.02% and 0.01% of the total packets, *i.e.*, 500 and 250, respectively. We configure Tower+Fermat and its competitors as follows. Overall, the configurations of these competitors are recommended in literature.

- **Tower+Fermat:** For Tower, we set it to consist of an 8-bit counter array and a 16-bit counter array. For Fermat, We set its count field and ID field to 32bits, and allocate 2500 buckets to it for 99.9% decoding success rate. we set the threshold T_h for identifying heavy-hitter candidates to the heavy-change threshold Δ_c , *i.e.*, 250, for detecting most heavy-hitters and heavy-changes.
- **CM/CU/CountHeap:** We use 3 hash functions as recommended in [109]. We set the counter size to 32bits. For CountHeap, we additionally set its heap capacity to 4096 for heavy-hitter detection.

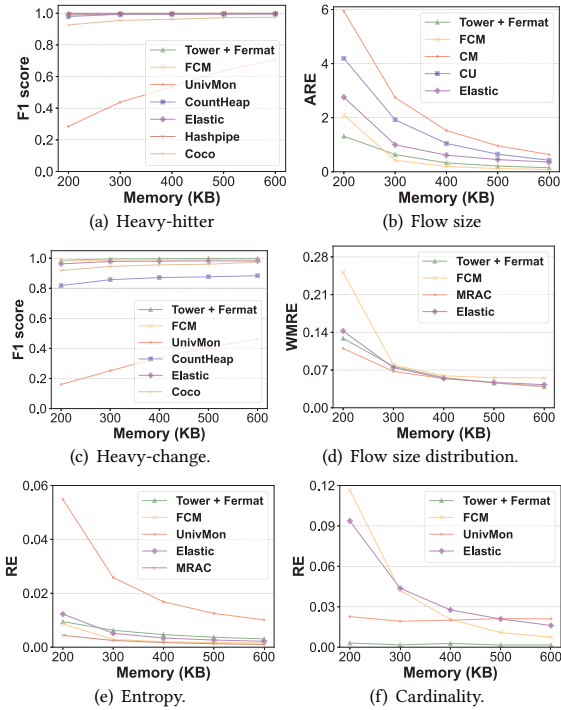


Figure 11: Accuracy comparison for six tasks.

- *UnivMon*: We use 14 levels and each level can record 1000 heavy hitters.
- *Elastic*: We use the hardware version of Elastic. For the heavy part, we use 4 stages and each stage has 3072 buckets. For the light part, we use a one-layer CM with 8-bit counters.
- *FCM*: We use the top- k version of FCM. It is almost the same as Elastic except the light part is substituted by a 16-ary FCM whose depth is set to 2.
- *Hashpipe*: We set the number of stages to 6.
- *Coco*: We use the hardware version of Coco that only uses one hash function.

Heavy-hitter detection (Figure 11(a)): Experimental results show that Tower+Fermat achieves comparable accuracy with HashPipe, and higher accuracy than other algorithms. When using only 200KB memory, the F1 score of Tower+Fermat is 99.8%, while that of Elastic and FCM is lower than 99%.

Flow size estimation (Figure 11(b)): Experimental results show that Tower+Fermat achieves comparable accuracy with FCM, and higher accuracy than other algorithms. When using only 200KB memory, the ARE of Tower+Fermat is 4.51 times, 3.19 times, 2.09 times, and 1.59 times smaller than that of CM, CU, Elastic, and FCM, respectively.

Heavy-change detection (Figure 11(c)): Experimental results show that the Tower+Fermat achieves higher accuracy than other algorithms. Tower+Fermat achieves 99.6% F1 score when using only 400KB memory, while that of the other algorithms is below 99.0%.

Flow size distribution estimation (Figure 11(d)): Experimental results show that Tower+Fermat achieves higher accuracy than Elastic and FCM, and comparable accuracy with MRAC. When




Figure 12: Implementation logic of ChameleonMon.

using 600KB memory, the WMRE of Tower+Fermat is 0.039, 1.09 and 1.42 times smaller than that of Elastic and FCM, respectively. **Entropy estimation (Figure 11(e))**: Experimental results show that Tower+Fermat achieves higher accuracy than UnivMon, and comparable accuracy with Elastic and FCM. When using 600KB memory, the ARE of Tower+Fermat is 0.003, 3.3 times smaller than that of UnivMon.

Cardinality estimation (Figure 11(f)): Experimental results show that the Tower+Fermat achieves higher accuracy than other algorithms. When using 600KB memory, the RE of Tower+Fermat is 0.0016, 13.125 times, 10.08 times, and 4.57 times smaller than that of UnivMon, Elastic, and FCM, respectively.

D PROTOTYPE IMPLEMENTATION

In this section, we present the important details of ChameleonMon prototype. We lay out important implementation details of the ChameleonMon data plane and control plane in sequence.

D.1 Data Plane Implementation

We have fully implemented the ChameleonMon data plane on the switch data planes of four edge Tofino switches in P4 [44]. In this section, we detail the implementation logic of data plane along the workflow (Figure 12).

Hash: First, a packet with flow ID f enters the network at an edge switch. With its flow ID (5-tuple) as input, the packet is hashed to multiple indexes through pairwise-independent hash functions generated from different CRC polynomials, which are deployed at stage 0 in ingress. These hash indexes are either used as base indexes for locating the mapped counters/buckets in the subsequent insertions, or used for sampling LL candidates, or used as fingerprints for improving decoding success rate of FermatSketch. Note that due to the limitation of Tofino switches, each hash index is uniformly distributed on $[0, 2^t - 1]$, where t is an arbitrary positive integer.

1-bit flipping timestamp: Second, the packet reads the current 1-bit flipping timestamp and from a match-action table, which is deployed at stage 1 in ingress. The 1-bit timestamp is used to indicate the corresponding group of sketches for the subsequent insertions.

Flow classifier: Third, the packet is inserted into the flow classifier, which is deployed at stage 2-3 in ingress. The flow classifier is a TowerSketch consisting of an 8-bit counter array and a 16-bit counter array. The 8-bit and 16-bit counter arrays consist of w_1 8-bit and w_2 16-bit counters, respectively. Each counter array is built on a register and accessed by a stateful arithmetic logic unit (SALU). To save SALU resources, we simulate the two flow classifiers by doubling the number of counters of the 8-bit and 16-bit counter arrays instead of building additional registers. The left/right w_1 8-bit and w_2 16-bit counters form the flow classifier corresponding to timestamp 0/1, respectively. We use the base indexes calculated by hash functions as the relative positions of the mapped counters in the flow classifier, and add offsets corresponding to the 1-bit

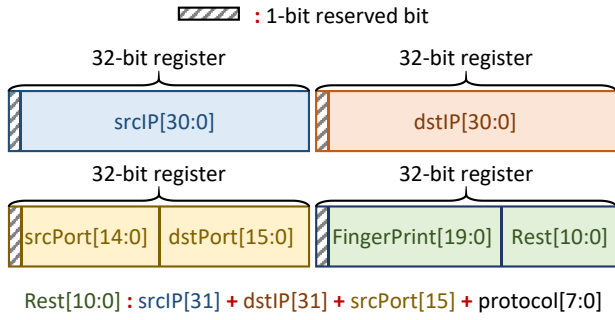


Figure 13: Division of the 5-tuple.

Algorithm 3: Simulated modular addition.

Input: An ID fragment f , a counter reg for encoding the ID fragment, and a prime p .

```

1  $inv \leftarrow p - f$            ▷ Get the additive inverse of  $f$  in  $Z_p$ 
2 if  $reg + f < p$  then
3    $reg \leftarrow reg + f$ 
4 else
5    $reg \leftarrow reg - inv$ 
6 end

```

timestamp to the base indexes, so as to locate the mapped counters. Specifically, when the timestamp is 0, the offset is 0; when the timestamp is 1, the offset is w_1 for 8-bit counter array or w_2 for 16-bit counter array. During insertion, the SALU adopts saturated addition operation for each mapped counter, which can increment the counter to its maximum value but never overflow it, and reports the value recorded in the counter, so as to simulate the behavior of TowerSketch. After insertion, we take the minimum value among the reported values as the size of flow f , and then input the flow size to a match-action Table that uses range matching on the flow size, so as to obtain the hierarchy of flow f .

Sampling: Fourth, if flow f is classified as a LL candidate, the packet reads a value from a match-action table, which is deployed at stage 4 in ingress. We then compare the read value with a 16-bit value, which is calculated by a hash function with the 5-tuple of the flow and a random seed as input. If the read value is equal to or larger than the 16-bit value, flow f is classified as a sampled LL candidate. Otherwise, the flow is classified as a non-sampled LL candidate. Obviously, to simulate a sample rate R , the value should be set to $\lceil 65536 \times R \rceil$.

FermatSketch: Before detailing the implementation of upstream and downstream flow encoders, we present the implementation of FermatSketch that they are based on. To encode the 104-bit flow ID (5-tuple) of each packet, an ideal bucket in FermatSketch should contain a 105-bit IDsum field and a 32-bit count field. However, because each SALU can access up to a pair of 32-bit counters, the IDsum field cannot be directly built in Tofino. To address this issue, we divide the IDsum field into multiple counters. Rather than encoding complete flow IDs, each counter only encodes specific ID fragments. Considering that a 32-bit counter can support at most 32-bit primes, and thus can encode at most 31-bit ID fragment, we need four 32-bit counters to simulate the IDsum field. Specifically, the division of the IDsum field is shown in Figure 13. The first three 32-bit counters encode the lower 31-bits of the source IP address,

the destination IP address, and the concatenation of the source port and destination port, respectively. The last one 32-bit counter encodes the rest 11-bit ID fragment (1-bit source IP address + 1-bit destination IP address + 1-bit source port + 8-bit protocol), and the other unused 20 bits are used to encode a fingerprint to improve decoding success rate. In summary, each bucket of FermatSketch consists of five 32-bit counters: four counters to encode the IDsum field and the fingerprint field, and a counter to encode the count field. Considering that there is no dependency between the five counters in any bucket of FermatSketch, a bucket array of FermatSketch can be built with five 32-bit counter arrays, each of which is built on a register and accessed by a SALU. During insertion, for any of the four counter arrays encoding the IDsum field and fingerprint field, the SALU needs to insert the specific ID fragment into its counter through modular addition. As shown in Algorithm 3, the SALU simulates the modular addition with logic consisting of a conditional judgement and two branches. Such logic is naturally supported by SALUs. For the other counter array encoding the count field, the SALU simply increments its counter by one. In this way, the SALUs simulate the behavior of FermatSketch. By duplicating these five registers and SALUs d times, we can easily build a d -array FermatSketch. Note that we use registers consisting of 32-bit counters, but not registers consisting of pairs of 32-bit counters that can further save SALU resources, to simulate the buckets of FermatSketch. This is because the logic used to simulate the modular addition requires two 32-bit metadata (f and inv) as input, which is just the maximum number that a SALU can support. However, encoding two ID fragments with a SALU requires four 32-bit metadata as input, which is beyond the capabilities of SALU. **Upstream flow encoder:** Fifth, unless flow f is a non-sampled LL candidate, the packet is inserted into the upstream flow encoder, which is deployed at stage 8-11 in ingress. The upstream flow encoder consists of three bucket arrays for the highest memory efficiency. Each bucket array is built as described above, and consists of m_{uf} buckets. The left m_{ll} buckets, the right m_{hh} buckets, and the middle m_{hl} buckets in each array form the upstream LL encoder, HH encoder, and HL encoder, respectively. Similarly, we simulate the two upstream flow encoders by doubling the number of buckets in each array. Based on the hierarchy of flow f , we can easily determine the encoder that the packet should be inserted into. We denote the number of buckets of a bucket array of that encoder by m' . Different from the flow classifier, the base indexes calculate by hash functions cannot be directly used to locate the relative positions of the mapped buckets in the encoder. This is because a base index is uniformly distributed on $[0, 2^l - 1]$, while m' , which could be any of m_{ll} , m_{hl} , and m_{hh} , may not be powers of two, as they are required to vary for supporting dynamic memory allocation. To address this issue, we decide to use the results of base indexes modulo m' as the relative positions of the mapped buckets. To simulate modulo operation in data plane, we input the hierarchy of flow f and a base index h_b to a match-action table that uses exact matching on flow hierarchy and range matching on index. The table first determines m' based on the input flow hierarchy, then outputs the largest number that is divisible by m' and less than h_b , and finally subtracts that number from h_b . Obviously, the result equals to h_b modulo m' . In this way, we locate the relative positions of the mapped buckets at the cost of TCAM

resources, and can finally locate the mapped buckets by adding offsets corresponding to the 1-bit timestamp and the flow hierarchy to the relative positions. Considering that the width of base index is fixed at run-time, if its width is too long compared to the width of m' , the match-action table will need a lot of entries to support range matching, and thus consumes lots of TCAM resources; if its width is just a bit longer than the width of m' , the uniformity of the calculated relative positions will be quite poor, leading to reduction of the decoding success rate of FermatSketch. To address this issue, before we input the base index to the match-action table, we bitwise-AND the base index with a mask to guarantee that the value range of the index is between $4m'$ and $8m'$, so as to make great trade-off between the uniformity of relative positions and the consumption of TCAM resources. Note that due to the inherent features of TCAM, when TCAM is used for range matching, different value range would require different number of TCAM entries for supporting modulo operation.

Downstream flow encoder: Sixth, unless flow f is a non-sampled LL candidate, the packet is inserted into the upstream flow encoder, which is deployed at stage 4-7 in egress. The implementation of downstream flow encoder is almost the same as that of upstream flow encoder, except it omits the heavy-hitter encoder. Note that the flow hierarchy and 1-bit timestamp are obtained from the edge switch where the packet enters the network, and carried by recording them in three bits of the ToS field of the IPv4 protocol.

Resources Usage: As shown in Table 1, under the parameter settings in Section 5.2, the Chameleon data plane consumes SALUs most, achieving 66.7%. This is because the flow classifier, the upstream flow encoder, and the downstream flow encoder all need SALUs to access memory. For resources other than SALUs, Chameleon consumes no more than 25%. Overall, the resource usage of Chameleon is moderate. Although Chameleon indeed consumes a lot of SALUs, the consumption of SALUs will not increase when we further enlarge the above sketches. With the advent of Tofino 2 switches and even Tofino 3 switches, we believe the resource usage will be much more acceptable on these more advanced programmable switches.

D.2 Control Plane Implementation

Central controller: The central controller integrates three modules into a DPDK [48] program: 1) a packet receiver module responsible for collecting sketches; 2) an analyzer module for decoding sketches, monitoring real-time network state, and generating reconfiguration packets for reconfiguring the Chameleon data plane; 3) a packet sender module responsible for sending reconfiguration packets to the control plane of each edge switch.

Switch control plane: The control plane of each edge switch runs a C++ program to load the P4 program to the Tofino ASIC. Every time the switch control plane receives a reconfiguration packet, it first extracts the packet to obtain the reconfiguration. Then, based on the reconfiguration, it generates corresponding table entries and update them to the corresponding match-action tables in the data plane to reconfigure the switch data plane. The time consumption in this step is shown in Figure 22 in Appendix F. To avoid the updated entries to interfere with the monitoring of the current epoch, those corresponding match-action tables further use exact matching on

the 1-bit timestamp. Those newly updated entries match the 1-bit timestamp in the next epoch, so as to function in the next epoch.

Epoch length: On our testbed, we set the epoch length to 50ms by default, and the additional time for all traffic passing through the network is set to 10ms (described in Appendix B).

Clock synchronization: On our testbed, we use the well-known software time synchronization protocol NTP [110] to synchronize the clocks between the control planes of edge switches and the central controller. Every 10s, every edge switch synchronizes its clock with the central controller. The precision of synchronization is around 0.3ms~0.5ms, and thus NTP can already satisfy the precision requirement for epochs of 50ms. We can also use more advanced solutions [111, 112] for us-level or even ns-level precision.

Data plane collection: To collect sketches from data planes of edge switches, a naive solution is to directly use the C++ control plane APIs provided by the Tofino SDK [113]. Currently, the most efficient strategy for this solution is to first use bulk DMA transfer to read data plane counter arrays into control plane buffer, and then read the counter arrays from control plane buffer [114]. However, on our testbed, such strategy takes about 338ms to collect only the upstream flow encoder, which seriously limits the setting of epoch length, and thus degrades the accuracy and timeliness of measurement. To address this issue, we fully exploit the capabilities and features of programmable data plane, including SALUs, mirror, and recirculate ports. Specifically, the switch control plane just needs to send several tailored packets to data plane for collecting sketches. The tailored packet is forwarded to the recirculate port, so as to access the counters of each sketches in turn. Every time a tailored packet accesses a counter, leveraging the SALU, it reads the value and inserts the value into its payload. Every time a tailored packet reaches the maximum transmission unit (MTU, e.g., 1514 Bytes), the switch data plane forwards it to the central controller, and mirrors a new truncated packet (e.g., 64 Bytes) to read the remaining counters. In this way, collecting the upstream flow encoder from the switch data plane only takes 0.44ms, which is 775 times faster than the straightforward solution. To ensure that the tailored packets will not be lost during the transmission, we reserve idle ports in their forwarding paths. Overall, the central controller takes 11.33ms to collect sketches from the data plane each edge switch, which consists five parts: 1) every time the timestamp flips, the central controller first sleeps 1ms to eliminate the impact caused by the error in clock synchronization, ensuring that all the edge switches have started the current epoch; 2) the central controller takes 2.68ms to collect the flow classifier; 3) the central controller takes 0.44ms to collect the upstream flow encoder; 4) the central controller sleeps 6.88ms to ensure that all the packets in the previous epoch have passed through or lost in the network; 5) the central controller takes 0.33ms to collect the downstream flow encoder.

E EVALUATION ON DIFFERENT WORKLOADS

In this section, we show that on workloads other than DCTCP, how Chameleon shifts measurement attention with the change of the number of flows or ratio of victim flows. For the measurement attention under different number of flows, we vary the number of flows in the network from 10K to 100K, and fix the ratio of victim

flows to 10%. For the measurement attention under different ratios of victim flows, we vary the ratio of victim flows from 2.5% to 25%, and fix the number of flows to 50K.

E.1 CACHE Workload

Measurement attention vs. number of flows (Figure 14): As the number of flows increases from 10K to 20K, ChameleMon can record all flows and victim flows, and therefore sets both T_h and T_l to 1. As the number of flows increases from 30K to 70K, ChameleMon allocates more memory to HL encoders and raises T_h higher than 1. As the number of flows increases from 80K to 100K, the healthy network state transitions to the ill network state. ChameleMon allocates memory to LL encoders, increases T_l and decreases the sample rate, so as to control the number of HLs and sampled LLs. Meanwhile, ChameleMon raises T_h to control the number of HH candidates. The relatively low load factor when the number of flows is between 80K and 100K is because of the high skewness of CACHE workload: lower thresholds will lead to a huge increase of the number of recorded flows, thus causing decoding failure. In fact, when the number of flows is between 80K and 100K, the T_h is set to 3, and the T_l is set to 2. ChameleMon has tried its best to select thresholds to maximize the load factor.

Measurement attention vs. ratio of victim flows (Figure 15): As the ratio of victim flows increases from 2.5% to 12.5%, ChameleMon records all victim flows by allocating more and more memory to HL encoders. T_h is not adjusted because of the high skewness of CACHE workload: setting T_h to 2 already makes a fairly small portion of flows as HH candidates, and lower T_h leads to decoding failure. As the ratio of victim flows increases from 15% to 25%, the healthy network state transitions to the ill network state. ChameleMon allocates memory to LL encoders, increases T_l to 2 and decreases the sample rate so as to control the number of HLs and sampled LLs. Meanwhile, because the memory of upstream heavy-hitter encoder and the number of flows remain unchanged, T_h also remains unchanged. The reason why ChameleMon suffers low load factor when the ratio of victim flows is between 15% to 25% is also due to high skewness of CACHE workload. Both T_h and T_l are set to 2, and decrease of thresholds will lead to decoding failure. ChameleMon has tried its best to select thresholds to maximize the load factor.

E.2 VL2 Workload

Measurement attention vs. number of flows (Figure 16): As the number of flows increases from 10K to 20K, ChameleMon can record all flows and victim flows, and therefore sets both T_h and T_l to 1. As the number of flows increases from 30K to 60K, ChameleMon allocates more and more memory to HL encoders, and increases T_h to decrease the number of HH candidates to avoid decoding failure. As the number of flows increases from 70K to 100K, the healthy network state transitions to the ill network state. ChameleMon allocates memory to LL encoders, increases T_l , and decreases the sample rate, so as to control the number of HLs and sampled LLs. Meanwhile, ChameleMon keeps increasing T_h to control the number of HH candidates. Throughout the experiment, ChameleMon maintains the load factor higher than 51%. The load factor is slightly lower, and it is because the distribution of VL2 is highly skewed.

Decreasing the thresholds by 1 will lead to huge increase in the number of recorded flows, and thus causing decoding failure.

Measurement attention vs. ratio of victim flows (Figure 17): As the ratio of victim flows increases from 2.5% to 12.5%, ChameleMon records all victim flows by allocating more and more memory to HL encoders, and increases T_h to decrease the number of HH candidates. As the ratio of victim flows increases from 15% to 25%, ChameleMon cannot record all victim flows and thus the healthy network state transitions to the ill network state. ChameleMon allocates memory to LL encoders, increases T_l , and decreases the sample rate so as to control the number of HLs and sampled LLs. Meanwhile, because the memory of upstream HH encoders and the number of flows remain unchanged, T_h also remains unchanged. Throughout the experiment, ChameleMon maintains the load factor higher than 53%. The load factor is slightly lower, and the reason is the same as the former experiment of the the number of flow.

E.3 HADOOP Workload

Measurement attention vs. number of flows (Figure 18): As the number of flows increases from 10K to 20K, ChameleMon can record all flows and victim flows, and therefore sets both T_h and T_l to 1. As the number of flows increases from 30K to 60K, ChameleMon allocates more and more memory to HL encoders, and increases T_h to decrease the number of HH candidates to avoid decoding failure. As the number of flows increases from 70K to 100K, the healthy network state transitions to the ill network state. ChameleMon allocates memory to LL encoders, increases T_l , and decreases the sample rate to control the number of HLs and sampled LLs. Meanwhile, ChameleMon keeps increasing T_h to control the number of HH candidates. Throughout the experiment, ChameleMon maintains the load factor higher than 47%. The load factor is slightly lower, and it is because the distribution of HADOOP is highly skewed. Decreasing the thresholds by 1 will lead to huge increase in the number of recorded flows, and thus causing decoding failure.

Measurement attention vs. ratio of victim flows (Figure 19): As the ratio of victim flows increases from 2.5% to 12.5%, ChameleMon records all victim flows by allocating more and more memory to HL encoders, and increases T_h to decrease the number of HH candidates. As the ratio of victim flows increases from 15% to 25%, ChameleMon cannot record all victim flows and thus transitions to ill network state. ChameleMon allocates memory to LL encoders, increase T_l , and decreases sample rate, so as to control the number of HHs and HLs. Throughout the experiment, ChameleMon maintains the load factor higher than 48%. The load factor is slightly lower, and the reason is the same as the former experiment of the the number of flows.

F EVALUATION ON TIME/BANDWIDTH OVERHEAD

To evaluate how fast can ChameleMon monitor the network, we evaluate various factors that could affect the setting of epoch length: 1) the time and bandwidth required to collect sketches from edge switches, 2) the time required to respond to different network states, and 3) the time required to reconfigure the ChameleMon data plane. The central controller only uses one CPU core in evaluation.

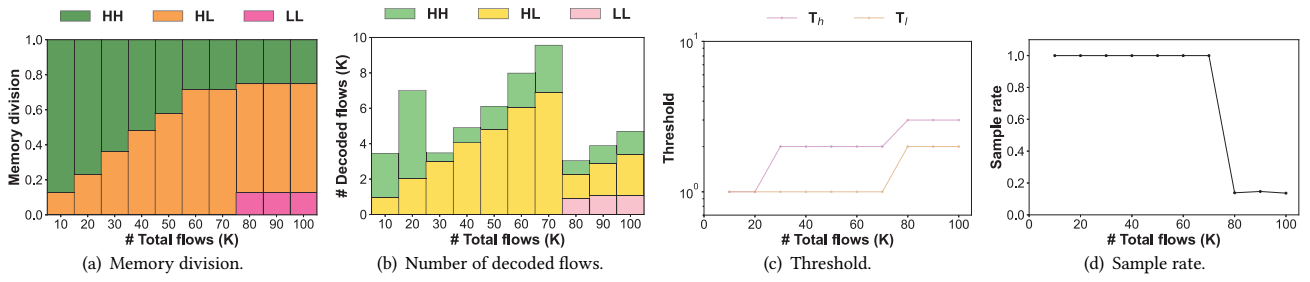


Figure 14: Measurement attention vs. number of flows on CACHE workload.

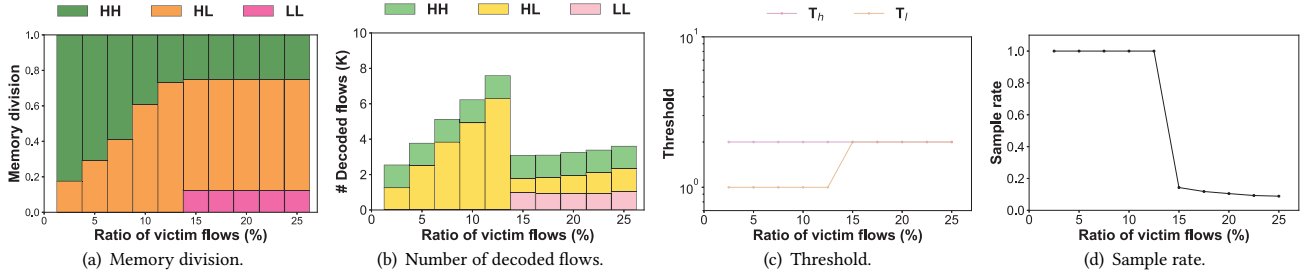


Figure 15: Measurement attention vs. ratio of victim flows on CACHE workload.

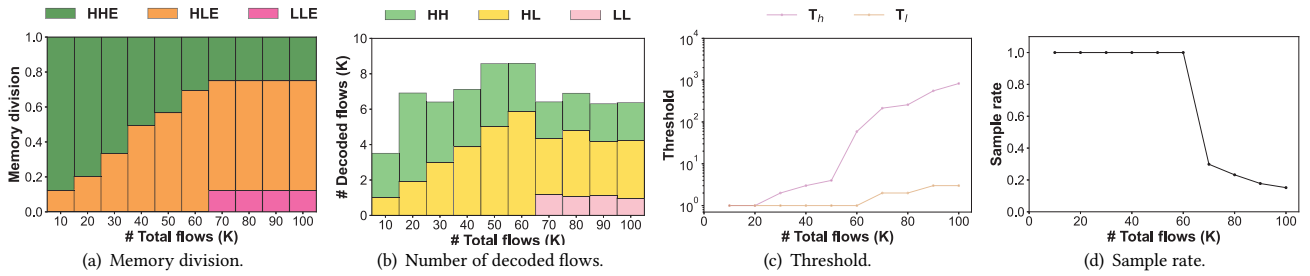


Figure 16: Measurement attention vs. number of flows on VL2 workload.

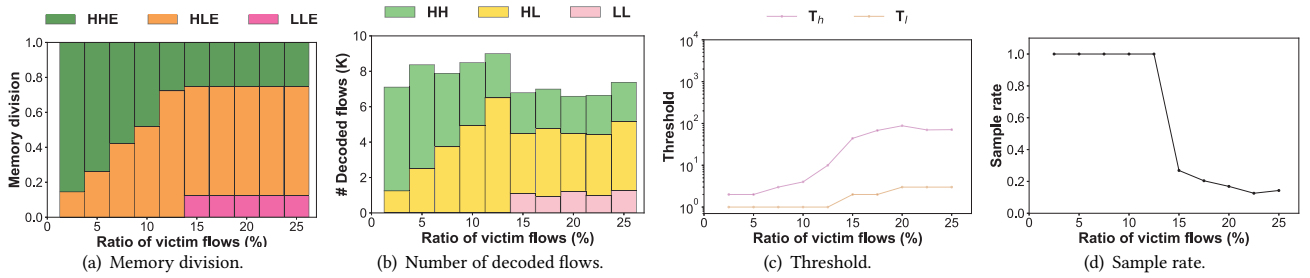


Figure 17: Measurement attention vs. ratio of victim flows on VL2 workload.




Figure 18: Measurement attention vs. number of flows on HADOOP workload.

Time/Bandwidth overhead for collection (Figure 21): Experimental results show that Chameleon consumes only a small amount of time and bandwidth in collecting all the data structures deployed on edge switches. Chameleon takes a total of 11.33ms

to collect sketches (refer to Appendix D.2 for details). As for bandwidth, when the epoch length is set to 50ms, the bandwidth overhead for collection is 317Mbps, consuming only 0.8% bandwidth for the central controller equipped with a 40Gb NIC.




Figure 19: Measurement attention vs. ratio of victim flows on HADOOP workload.

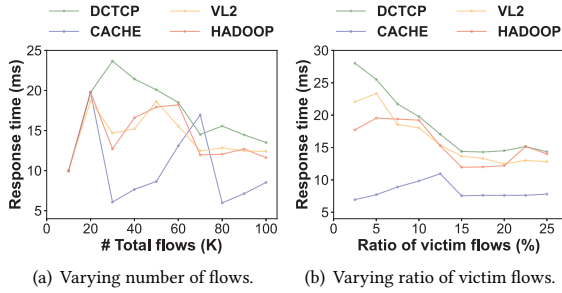


Figure 20: Response time to different network states.




Figure 21: Bandwidth.

Figure 22: Reconfiguration time.

Response time to different network states (Figure 20): Experimental results show that ChameleMon can always respond to different network states within 30ms. We count the response time of ChameleMon to each network state previously appeared in Figure 7-8, where the response time refers to the time interval between the central controller finishing the collection of sketches and the central controller generating the reconfiguration packet⁸ for the ChameleMon data plane. Although the response time does not seem to show a clear trend with the network state, it is mainly determined by the number of HH candidates, because the central controller needs to first extract them from the upstream HH encoders and then reinsert them to the upstream HL encoders. As shown in Figure 20(b), as the ratio of victim flows increases, the response time on all the four workloads decreases because the number of HH candidates decreases. The response time finally stabilizes because the fixed memory allocation in the ill network state always decodes a similar number of flows.

CDF of reconfiguration time (Figure 22): Experimental results show that it takes 2~7ms to reconfigure the ChameleMon data plane. The central controller sends 10K reconfiguration packets with random configuration of the ChameleMon data plane to each edge switch, and we count the time for an edge switch to execute

the reconfiguration. We find 60% of reconfigurations take less than 5ms. The difference in time consumption is mainly because different reconfigurations require updating different numbers of TCAM entries to the switch data plane for supporting dynamic memory allocation (refer to Appendix D.1 for details).

Adding up the above all time consumption, we find that the overall time consumption is less than 50ms. This verifies that ChameleMon can monitor the network every 50ms on our testbed. Considering that 1) the central controller only uses one CPU core in experiments and 2) monitoring the network every 50ms only consumes 0.8% bandwidth of a 40Gb NIC, we believe ChameleMon can easily scale to monitor a much larger network with a shorter epoch length, requiring only one server as the central controller.

⁸The central controller sends the reconfiguration packets to edge switches to reconfigure their data planes. Please refer to Appendix D.2 for details.