

OmniMon: Re-architecting Network Telemetry with Resource Efficiency and Full Accuracy

Qun Huang^{1,2}, Haifeng Sun^{1,2}, Patrick P. C. Lee³, Wei Bai⁴, Feng Zhu², and Yungang Bao²

¹Peking University ²Institute of Computing Technology, Chinese Academy of Sciences

³The Chinese University of Hong Kong ⁴Microsoft Research

ABSTRACT

Network telemetry is essential for administrators to monitor massive data traffic in a network-wide manner. Existing telemetry solutions often face the dilemma between *resource efficiency* (i.e., low CPU, memory, and bandwidth overhead) and *full accuracy* (i.e., error-free and holistic measurement). We break this dilemma via a network-wide architectural design OmniMon, which simultaneously achieves resource efficiency and full accuracy in flow-level telemetry for large-scale data centers. OmniMon carefully coordinates the collaboration among different types of entities in the whole network to execute telemetry operations, such that the resource constraints of each entity are satisfied without compromising full accuracy. It further addresses *consistency* in network-wide epoch synchronization and *accountability* in error-free packet loss inference. We prototype OmniMon in DPDK and P4. Testbed experiments on commodity servers and Tofino switches demonstrate the effectiveness of OmniMon over state-of-the-art telemetry designs.

CCS CONCEPTS

• **Networks** → **Network measurement**.

KEYWORDS

Network measurement; Distributed systems

ACM Reference Format:

Qun Huang, Haifeng Sun, Patrick P. C. Lee, Wei Bai, Feng Zhu, and Yungang Bao. 2020. OmniMon: Re-architecting Network Telemetry with Resource Efficiency and Full Accuracy. In *Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication (SIGCOMM '20)*, August 10–14, 2020, Virtual Event, NY, USA. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3387514.3405877>

1 INTRODUCTION

Modern data centers host numerous critical applications, thereby imposing high demands on network management. Administrators need efficient *network telemetry* solutions to monitor massive traffic and understand its network-wide behaviors. Ideally, we aim for both *resource efficiency* and *full accuracy* when developing network telemetry. For resource efficiency, network entities (e.g., end-hosts

and switches) face heterogeneous resource constraints in computation, memory, and bandwidth, so telemetry systems should limit their performance overhead and never exhaust any specific resources. In particular, the performance overhead of network telemetry must be much lower than that of normal packet processing in routing, NAT, and firewall. Note that achieving resource efficiency does not mean that the resource usage remains constant as the network scale grows; instead, the goal of resource efficiency here is to keep the resource overhead of network telemetry much lower than that of other network operations. For full accuracy, telemetry systems should have in-network, always-on visibility that covers all flows and all entities with error-free flow-level statistics.

Conventional wisdom views resource efficiency and full accuracy as a design trade-off. At one extreme, per-flow monitoring (e.g., Trumpet [49] and Cisco Tetration [15]) aims for fine-grained, error-free measurement but incurs possibly unbounded resource usage; at another extreme, coarse-grained monitoring (e.g., SNMP [12], sFlow [64], and NetFlow [53]) provides best-effort measurement without accuracy guarantees. Between the extremes, many telemetry systems adopt approximation algorithms [5, 26, 28, 29, 38, 48, 62, 63, 65, 72, 75] that achieve high resource efficiency with provably bounded errors, or event matching techniques [25, 33, 52, 69, 74, 79] that only focus on the traffic of interest.

We pose a fundamental question: *Can we break the dilemma between resource efficiency and full accuracy in network telemetry?* We observe that there exist delicate design trade-offs between resource availability and in-network visibility among different network entities, including end-hosts and switches in the data plane as well as a centralized controller in the control plane (§2.1). By carefully re-architecting the collaboration among network entities subject to their resource constraints, we can design a new telemetry architecture that achieves both resource efficiency and full accuracy. Note that recent telemetry proposals also take a collaborative approach [24, 52, 67]. However, they address different design goals (e.g., expressiveness [24, 52] or in-network visibility [67]) rather than the resource-accuracy trade-off. Their architectures cannot readily achieve full accuracy due to the tight switch resources (§2.2). We fill this void through a complementary collaborative telemetry architecture that treats both resource efficiency and full accuracy as “first-class citizens”.

We present OmniMon, a novel network-wide architecture that carefully coordinates telemetry operations among all entities (i.e., end-hosts, switches, and the controller), with a primary goal of simultaneously achieving both resource efficiency and full accuracy in flow-level network telemetry for large-scale data centers. Specifically, OmniMon re-architects network telemetry in a *split-and-merge* fashion. By *split*, OmniMon decomposes network telemetry into partial operations and schedules these partial operations

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 ACM 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.

SIGCOMM '20, August 10–14, 2020, Virtual Event, NY, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7955-7/20/08.

<https://doi.org/10.1145/3387514.3405877>

among different entities. This design fully utilizes the available resources of the entire network and bypasses the resource constraints of individual entities. By *merge*, OmniMon coordinates all entities to collaboratively execute the partial operations. It combines the strengths of both end-hosts and switches (in the data plane) for error-free per-flow tracking across the entire network; meanwhile, its controller (in the control plane) conducts network-wide resource management and collective analysis of all end-hosts and switches. Note that OmniMon is incrementally deployable on a subset of end-hosts and switches across the network.

As in any practical distributed system, the network-wide architectural design of OmniMon needs to address the reliability concerns in network telemetry. Thus, OmniMon incorporates two reliability guarantees, namely *consistency* and *accountability*, to preserve both resource efficiency and full accuracy in practical deployment. For consistency, we propose a network-wide epoch synchronization mechanism with a *hybrid* consistency model, such that (i) all entities reside at the same epoch at most times and their epoch boundaries only deviate by a small time difference (up to 60 μ s; see §8) and (ii) every packet is synchronized at the same epoch during its transmission even under cross-epoch network delays. For accountability, we formulate a system of linear equations for per-switch, per-flow packet loss inference, and ensure that a unique solution exists in the common case by taking into account data center network characteristics.

We prototype OmniMon in DPDK [16] and P4 [56]. We also integrate OmniMon with Sonata [24] to support basic query-driven telemetry tasks. Our testbed experiments on commodity servers and Barefoot Tofino switches [70], as well as large-scale simulation, justify the effectiveness of OmniMon over 11 state-of-the-art telemetry designs. For example, OmniMon reduces the memory usage by 33-96% and the number of actions by 66-90%, while achieving zero errors, compared to sketch-based solutions (which have non-zero errors) [29, 38, 43, 72] in P4 switch deployment.

The source code of our OmniMon prototype is now available at: <https://github.com/N2-Sys/OmniMon>. This work does not raise any ethical issues.

2 BACKGROUND AND MOTIVATION

We first pose the resource and accuracy requirements in network telemetry and identify the use cases (§2.1). We then discuss the root causes of having the resource-accuracy trade-off in existing telemetry systems (§2.2).

2.1 Resource and Accuracy Requirements

We focus on flow-level network telemetry in a data center network, in which we identify a flow by a *flowkey* formed by a combination of packet fields (e.g., 5-tuples). We measure flow-level statistics, in the form of *values*, across fixed-length time intervals called *epochs*. We represent an epoch as an integer, by dividing the current timestamp by the epoch length. In particular, we address two requirements in network telemetry: *resource efficiency* and *full accuracy*.

Resource efficiency. We target a data-center-scale network telemetry framework that includes three types of entities: multiple *end-hosts*, multiple *switches*, and a centralized *controller*. Each type of entities imposes different resource constraints. Thus, the frame-

work should satisfy the resource constraints of the corresponding entities to maintain the overall packet forwarding performance. Also, it should only leverage the existing functionalities of commodity hardware. We elaborate each type of entities as follows.

(i) *End-host*: End-hosts process network traffic in the network edge in software (e.g., kernel space [41] and user space [16]). They provide ample memory space for holistic flow tracking, and high flexibility for realizing various telemetry approaches. However, as end-hosts reside in the edge, they have poor in-network visibility for network-wide measurement. Also, software packet processing on a commodity host is computationally expensive (e.g., requiring multiple CPU cores to achieve a rate of 40 Gbps [17]), yet end-hosts cannot dedicate all CPU resources only for network telemetry as the CPU resources are also shared by other co-locating applications.

(ii) *Switch*: Commodity switch ASICs achieve high forwarding throughput (e.g., several Tbps [10, 70]) and ultra low processing latency (e.g., sub-microseconds). However, they have scarce on-chip memory space (e.g., tens of MB to keep states [34, 45]), and provide limited programmability due to overheating concerns and high manufacturing costs.

(iii) *Controller*: The controller is a *logical* entity that centrally coordinates all end-hosts and switches to obtain a global network view. It can comprise multiple servers, so as to combine their CPU and memory resources for sophisticated analytics and provide fault tolerance (e.g., via Zookeeper [30]). However, it has limited network bandwidth to receive a substantial amount of traffic from all end-hosts and switches.

Full accuracy. We target network telemetry with *full accuracy*, which implies both *completeness* and *correctness*. By completeness, we mean to track all traffic (and hence all flows) in all entities and epochs without missing any information; by correctness, we mean to track the flowkey and complete values for each flow without errors. Full accuracy is critical for many telemetry tasks. We show three use cases that greatly benefit from full accuracy. Note that such use cases are also supported in existing telemetry systems (e.g., [4, 24]), yet any measurement error can degrade their efficiency and usability in practical deployment.

(i) *Performance evaluation*: Administrators often propose new techniques to improve network performance, and verify their correctness and performance gain in real deployment. One example is to evaluate the benefits of flowlet-based load balancing [3] over conventional flow-based load balancing. Administrators can make refinements using complete *per-flow* trajectories, which unveil both per-link utilization and top flows in each link. Full accuracy in per-flow trajectories avoids measurement variations or flawed refinements, and hence allows administrators to focus on evaluating the proposed techniques without concerning any measurement error.

(ii) *Anomaly detection*: Administrators identify network anomalies (e.g., misbehaviors or attacks) for reliability and security purposes. Such anomalies are hard to predict before they actually occur. Thus, administrators require *per-epoch* statistics of *all traffic* to respond to any detected anomalies. Full accuracy in per-epoch statistics eliminates false alarms or undetected anomalies due to any measurement error, and hence prevents any anomaly from compromising the network robustness. Here, we focus on collecting accurate flow statistics, while providing precise anomaly definitions

and configurations is beyond our scope.

(iii) *Network diagnosis*: Performance degradations (e.g., sudden high packet loss) are common in production for various reasons, such as misconfigurations and hardware failures. Administrators require *per-switch*, *per-flow* statistics to localize and repair problematic switches and lossy flows. Full accuracy in per-switch, per-flow statistics implies that any performance issue can be correctly pinpointed in a timely manner without any interference from inaccurate measurement.

2.2 Motivation

Existing telemetry systems (§9) pose resource efficiency and full accuracy as a design trade-off: they either provision enough resources (i.e., CPU, memory, and bandwidth) to achieve full accuracy, or sacrifice full accuracy to satisfy the resource constraints of individual entities. We argue that the trade-off exists due to two root causes.

Root cause 1: Tight coupling between flow tracking and resource management. Existing approaches realize telemetry operations *monolithically*, in which an end-host or a switch extracts the values of each observed flow and provisions resources to track the values (referenced by flowkeys). To resolve any resource conflict among flows, the amount of provisioned resources must be sufficiently large, yet it also heavily depends on the required level of accuracy and the resource constraints of individual entities.

Consider an example that we use a hash table of counters (e.g., as in Trumpet [49]) to track the flowkeys and values of all flows. To resolve hash collisions, we can leverage chaining or Cuckoo hashing [57] to organize counters for colliding flows. Such a design can be feasibly realized in software, but not in switch ASICs as the complex memory management of the hash table is hard to implement due to limited switch programmability. If we solely increase the hash table size and rely on uniform hash functions to resolve hash collisions in switches, the number of counters is huge. Figure 1(a) shows the memory usage (for 4-byte counters) versus the number of flows for different hash collision rates (in general, the collision rate is $1 - \frac{m!}{n!m^n}$ for m counters and n flows). Even for 1,000 flows, the memory usage reaches 200 MB for a collision rate of 1%, which is far beyond the switch capacity.

Existing switch-based telemetry systems sacrifice accuracy to fit all operations into switch ASICs [47, 51]. One solution is to leverage approximation algorithms (e.g., sampling [62, 63], top- k counting [5, 26, 65], and sketches [28, 29, 38, 43, 48, 75]) to fit compact structures in switch memory. However, full accuracy is inevitably violated due to limited provisioned resources. Another solution is to offload traffic to the controller to mitigate switch-side resource usage [25, 69, 74, 79], but it only monitors the traffic of interest to avoid overwhelming the link capacity of the controller.

Root cause 2: Limited network-wide collaboration. Several telemetry designs adopt some forms of collaboration among entities, yet they target different design goals and remain bounded by the resource-accuracy trade-off. For example, some systems deploy a controller to assemble measurement results from end-hosts or switches [28, 29, 38, 49] and/or tune resources among switches [43, 47, 48, 62]. Some hybrid architectures [24, 52] allow the collaboration among switches and the controller to empower efficient and expressive telemetry query engines, while others [32, 42, 67]

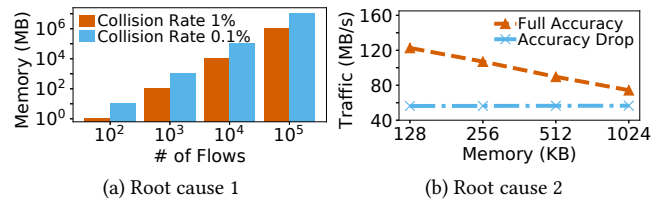


Figure 1: Root causes of resource-accuracy trade-off: (a) Memory usage using simple hashing with full accuracy; (b) Switch resource overheads in Sonata with controller-switch collaboration.

allow the collaboration among end-hosts and switches to combine programmability in end-hosts and visibility in switches. However, existing collaborative approaches do not coordinate the resources for full accuracy among *all* entities (i.e., end-hosts, switches, and the controller). The limited resources of some entities can make full accuracy infeasible.

We use Sonata [24], a collaborative telemetry architecture, to motivate our claims. Sonata targets expressiveness and scalability. It offloads dataflow operators from the controller to switches to exploit fast packet processing in switches. However, it cannot mitigate the memory and bandwidth usage of switch-side operators without sacrificing accuracy. We take the open-source Sonata prototype and evaluate its resource consumption using the 11 applications reported in [24]. We vary per-application memory from 128 KB to 1,024 KB, and measure the total switch traffic to the controller on Barefoot Tofino switches [70]. We compare two cases: (i) “Accuracy drop”, the default setup in Sonata’s prototype that does not deal with hash collisions; and (ii) “Full accuracy”, in which we achieve full accuracy by mapping each flow to d (set to 2 here) counters and evicting a flow to the controller if it has hash collisions in all the d counters (c.f. §3.1 of [24]). We configure 10⁵ active flows at each epoch. Figure 1(b) shows that around 56 MB/s of traffic is triggered in all cases with accuracy degradation. For full accuracy, the amount of triggered traffic substantially increases due to hash collisions (e.g., 2.2× the accuracy degradation case for 128 KB of per-application memory). Note that the handling of hash collisions also incurs excessive usage of other types of switch resources (§8).

3 OmniMon DESIGN

We state our design goals and assumptions (§3.1), and show how OmniMon achieves resource efficiency and full accuracy (§3.2). We pose the consistency and accountability issues in our design (§3.3).

3.1 Design Overview

OmniMon is a network-wide telemetry architecture spanning different entities (i.e., end-hosts, switches, and the controller) in a data center network. It aims for both resource efficiency and full accuracy from an architectural perspective, by re-architecting the collaboration among all entities subject to their resource constraints. It is compatible with existing collaborative telemetry designs; as a case study, we add a query engine based on Sonata [24] as an application atop OmniMon for expressive telemetry (§7).

Coordinating the collaboration among different types of entities is non-trivial due to their heterogeneity in resource availability and

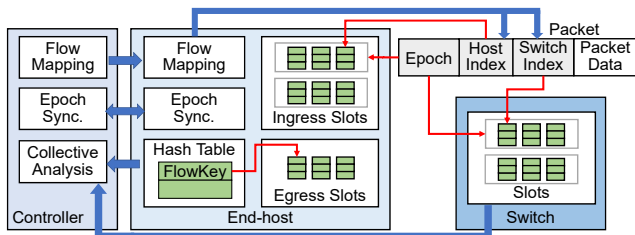


Figure 2: OmniMon architecture.

in-network visibility (§2.1). To this end, we address three classical issues adapted from building practical distributed systems: (i) How do we assign telemetry operations across different entities (i.e., end-hosts, switches, and the controller)? (ii) How should different entities coordinate among themselves? (iii) How should we achieve reliable coordination in the face of unreliable events?

Assumptions. We make several design assumptions.

- We focus on a data center network under a single administrative domain, a common deployment setting in modern data centers. Administrators have access to all entities to apply configurations. The controller has full knowledge of the packet processing policies (e.g., routing tables, access control lists, and path changes due to link failures) in each entity.
- We only consider the packets that actually enter the network (i.e., packets dropped in end-hosts are ignored). This can be done by realizing measurement in proper positions of the network stack in each end-host (e.g., `qdisc` in the Linux kernel). Measuring and dissecting an end-host’s network stack is out of our scope.
- The packet format is extensible. During transmissions, packets can embed new information in various unused fields (e.g., reserved bits in VXLAN) in existing packet headers [3, 20]. Recent advances in user-space networking stacks [16] and programmable switch ASICs [70] also allow administrators to readily define customized packet fields. Note that the embedded information should be of small size to limit transmission overhead.

3.2 Split-and-Merge Telemetry

Figure 2 presents the OmniMon architecture. Recall that OmniMon follows a split-and-merge approach for network telemetry (§1). Specifically, it performs flow tracking and resource management via four partial operations, namely *flowkey tracking*, *value updates*, *flow mapping*, and *collective analysis*. It coordinates these partial operations among end-hosts, switches, and the controller, subject to the heterogeneous resource requirements of different entities.

In the following, we describe each partial operation and show how both resource efficiency and full accuracy are simultaneously achieved. We first assume that there is a global clock that precisely synchronizes all entities and there is no packet loss. We later relax these assumptions in §3.3.

Flowkey tracking. OmniMon holds a copy of the flowkey of each active flow in the network. When a new flow starts, OmniMon records its flowkey in a hash table in the source end-host. It removes the flowkey when the flow terminates, determined by FIN/RST packets in TCP flows or a long idle time (e.g., 30 minutes [27]).

Value updates. OmniMon tracks flow values in both end-hosts

(with ample memory space) and switches (with in-network visibility). Each end-host or switch maintains a set of *slots* with multiple counters each, and further partitions the slots into groups for different epochs. As end-hosts and switches have different resource requirements, they perform value updates differently.

With sufficient memory space, an end-host dynamically allocates one *dedicated* slot for each flow per epoch. When a flow terminates, its slots will be recycled. An end-host holds two types of slots, namely *ingress* and *egress* slots. It stores the values of an incoming flow destined for itself in an ingress slot, and also stores the values of an outgoing flow originated from itself in an egress slot.

On the other hand, since switches have limited memory space, we allow a slot to be shared by multiple flows. Each switch maps a flow to a subset of slots and updates the counters in all those slots by the flow values. Thus, a slot holds aggregate values from multiple flows. We later recover the per-flow values via collective analysis (see below).

Flow mapping. OmniMon coordinates all end-hosts and switches to properly map flows to slots for flowkey tracking and value updates. Its flow mapping design builds on both in-band packet embedding and out-band messages.

(i) *Flow mapping in end-hosts.* OmniMon maps a flow to its slots in both source and destination end-hosts via an index number called the *host index*. When the source end-host applies flowkey tracking for a new flow, it first computes the host index, which denotes the entry position of the flowkey in the hash table. Then it associates the host index with the egress slot of the flow, and embeds the host index and the epoch into each outgoing packet of the flow. The destination end-host extracts the host index and epoch from each received packet, and associates them with the ingress slot. The use of the host index eliminates another hash lookup, which is the major overhead for end-hosts (Exp#1 in §8).

(ii) *Flow mapping in switches.* OmniMon uses an index number called the *switch index* for each flow to determine how the flow maps to multiple slots in a switch. The switch index is collaboratively selected by the controller and source end-hosts. Specifically, the controller assigns a unique switch index to each possible flow mapping to the slots in a switch. It then pre-assigns a disjoint subset of switch indexes to each end-host. A source end-host chooses one switch index for a new flow (detailed in §5.2) and embeds the switch index (together with the host index and the epoch) in each outgoing packet of the flow. When a switch receives a packet, it retrieves the switch index and epoch from the packet and updates the corresponding mapped slots. Such a collaborative approach not only exploits the visibility of the controller to mitigate the hash collisions in switch slots, but also allows the source end-host to quickly choose switch slots without querying the controller.

Collective analysis. For each epoch, the controller collects the hash tables (flowkeys) and the slots (values) from all end-hosts and switches. It performs collective analysis to associate each flow with the values recorded in its source and destination end-hosts and each switch the flow traverses, so as to recover per-flow values from each end-host and switch. Since OmniMon keeps all flowkeys in their source end-hosts, the controller can directly obtain per-flow values (via the host index) from the dedicated slots in the source and destination end-hosts. Per-flow values in switches can also be

recovered from the shared slots based on the flowkeys and values in end-hosts. Specifically, the controller knows all flows mapped to a shared slot (via the switch index). If there is no packet loss, the value of the slot is the aggregation of the flows that are mapped to the slot (e.g., for additive values, the value in the slot equals the sum of per-flow values); otherwise, with packet losses, OmniMon recovers per-flow values via packet loss inference (§5).

Summary. OmniMon achieves both resource efficiency and full accuracy. For full accuracy, the controller exactly recovers per-flow values in all end-hosts and switches on a per-epoch basis (i.e., both correctness and completeness are addressed). OmniMon addresses resource efficiency via three aspects: (i) for memory, it maintains flowkeys and dedicated per-flow values in end-hosts only, while keeping only flow values in the shared slots in switches; (ii) for computation, it performs flow mapping operations collaboratively in the controller and all end-hosts without involving switches, while incurring limited overhead (Exp#1 and Exp#7 in §8); (iii) for bandwidth, the controller collects hash tables and slots without mirroring all traffic.

Note that such resource usage is insignificant compared to other network functionalities. For end-hosts, per-packet hashing is affordable as shown in prior work [28, 49]; in contrast, network stacks in OS kernels or virtualization frameworks perform much more complicated per-packet processing than per-packet hashing in OmniMon. For switches, the partial operations of OmniMon incur much less resource usage than standalone telemetry solutions, as shown in our evaluation (§8).

3.3 Consistency and Accountability

OmniMon faces two unreliability concerns in its network-wide architectural design: (i) lack of a global clock and (ii) packet losses. We extend OmniMon with *consistency* and *accountability* guarantees, so as to preserve both resource efficiency and full accuracy in the face of unreliability events.

- *Consistency:* Without a global clock, OmniMon needs to ensure that (i) all entities reside at the same epoch at most times, thereby providing correctness guarantees for per-epoch applications (e.g., anomaly detection (§2.1)); and (ii) each packet is always monitored at the same epoch during its transmission (even under cross-epoch network delays), thereby permitting collective analysis to correctly associate per-flow values across entities in a consistent network-wide view of traffic. See §4 for details.
- *Accountability:* If packet losses happen (e.g., due to network congestion or switch failures), OmniMon needs to correctly infer per-switch, per-flow packet losses. It is critical to network diagnosis (§2.1), and maintains full accuracy in the collective analysis. See §5 for details.

4 CONSISTENCY

OmniMon achieves consistency via network-wide epoch synchronization, with the goal of synchronizing all entities at the same epoch and monitoring each packet at the same epoch in its transmission. We state the challenge of realizing consistency in OmniMon (§4.1), and propose a new consistency model that closely matches our objective yet can be feasibly realized (§4.2).

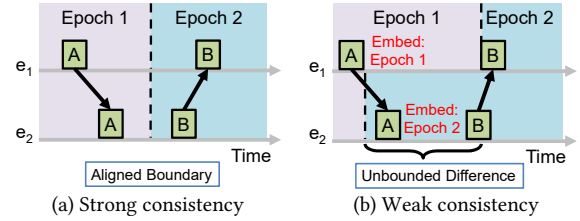


Figure 3: Existing consistency models.

4.1 Challenge

While consistency has been well studied in distributed computing research, providing consistency guarantees for network telemetry with both resource efficiency and full accuracy remains unexplored. Existing consistency protocols mainly target either *strong consistency* or *weak consistency*. We argue why each of them is difficult to achieve in our context via the examples in Figure 3, in which two entities (end-hosts or switches) e_1 and e_2 exchange packets.

Strong consistency. Strong consistency aligns epoch boundaries and ensures that every packet is observed by *all* entities at the same epoch. For example, in Figure 3(a), both e_1 and e_2 see Packet A at Epoch 1 and Packet B at Epoch 2. However, strong consistency protocols (e.g., Paxos [37]) are complex to implement in switches and trigger multiple rounds of out-band messages for synchronization.

Weak consistency. Weak consistency synchronizes per-packet epochs in a *best-effort* manner, so the epochs among entities may diverge and a packet may be monitored at different epochs. In weak consistency, each entity renews its epoch via a local clock and embeds the current epoch in every outgoing packet (as in Lamport Clock [36] and Distributed Snapshots [13, 73]). When an entity receives a packet, it updates its current epoch with the packet’s embedded epoch if the latter is newer. Otherwise, the entity updates the packet’s embedded epoch with its current epoch before forwarding.

Such in-band synchronization incurs limited costs. However, entities may observe a packet at different epochs as the embedded epoch can be modified along the packet transmission path. Also, a heavily clock-delayed entity never updates its epoch until it receives a packet with a newer epoch, causing *unbounded* deviations among epoch boundaries. For example, in Figure 3(b), both e_1 and e_2 synchronize their epochs via in-band packets. However, e_1 ’s local clock is delayed and remains at Epoch 1 until it receives Packet B from e_2 . Furthermore, e_1 embeds Epoch 1 in Packet A, while e_2 embeds Epoch 2. Thus, Packet A may be inconsistently treated as at Epoch 1 or Epoch 2 by different entities in its transmission.

4.2 Hybrid Consistency

OmniMon realizes a *hybrid* consistency model, which ensures strong consistency at most times and is downgraded to weak consistency for only a small bounded time period (e.g., tens of μ s). Our idea is that OmniMon not only embeds epoch information in in-band packets (as in weak consistency), but also leverages the controller’s coordination for network-wide epoch synchronization.

Design overview. OmniMon decides the epoch for each packet *once* during its transmission. Each end-host maintains a local clock and decides its current epoch. When a source end-host sends a

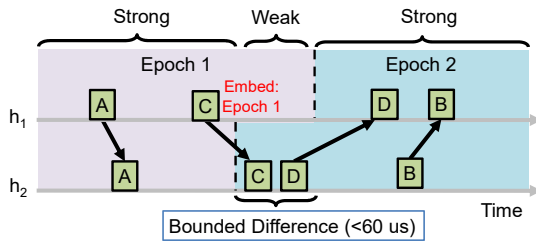


Figure 4: Hybrid consistency in OmniMon.

packet, it embeds its current epoch into the packet. All switches along the path and the destination end-host monitor the packet in the corresponding slots based on the embedded epoch. The destination end-host further updates its current epoch if the embedded epoch is newer. Note that switches do not maintain local epochs to keep their processing pipelines simple.

When an end-host updates its epoch via its local clock, it sends the new epoch to the controller, which relays the new epoch to all other end-hosts via out-band messages. Thus, all end-hosts can keep up-to-date epochs even though their local clocks are delayed or they do not receive any packets with newer embedded epochs. With this controller-assisted approach, it takes at most two hops to propagate a new epoch from one end-host to another. Thus, the time difference between the epoch updates of any two end-hosts is roughly twice the path delay between the controller and the end-hosts, which is around tens to hundreds of microseconds in practice (e.g., less than $60 \mu\text{s}$ for 1,024 end-hosts in our evaluation) and is much smaller than the epoch length (e.g., tens or hundreds of milliseconds [38, 49]).

The controller maintains its local epochs. It updates and relays the epoch only when receiving an out-band message with a newer epoch. We can show that each end-host triggers only $O(1)$ out-band messages at each epoch (see the proof in Appendix). Also, the variations of in-band traffic have limited impact on the consistency model (Exp#3 in §8).

Figure 4 depicts the hybrid consistency model. At most times, all end-hosts (e.g., h_1 and h_2 in the figure) reside at the same epoch. However, end-hosts may still reside at different epochs at times (e.g., when h_2 receives Packet C), in which case OmniMon only ensures weak consistency. Nevertheless, the duration of weak consistency is limited by the controller-assisted synchronization. Also, OmniMon ensures that every packet is monitored at the same epoch along the path, even when its source and destination end-hosts are not yet synchronized. For example, Packet C is sent by h_1 at Epoch 1 but received at Epoch 2 by h_2 . However, both h_1 and h_2 monitor Packet C at Epoch 1 (its embedded epoch). In Appendix, we formally prove the properties of hybrid consistency.

Algorithm. Algorithm 1 shows the epoch update operations among the controller and all end-hosts for realizing hybrid consistency. An end-host triggers an update of its epoch by its local clock (lines 1-3), a newer embedded epoch of a received packet (lines 4-6), or a newer epoch from an out-band message from the controller (lines 7-9). If the end-host triggers its epoch update by its local clock, it also sends the new epoch to the controller. For the controller, if it receives a newer epoch from an end-host, it updates its local epoch and sends it to all other end-hosts (lines 10-14).

Algorithm 1 Epoch updates in hybrid consistency

End-host: local epoch e_h

```

1: function ON_NEW_EPOCH_LOCAL_CLOCK(epoch e)
2:   if  $e > e_h$  then
3:      $e_h = e$ , and send  $e_h$  to the controller
4: function ON_NEW_EPOCH_FROM_EMBEDDED_PACKET(epoch e)
5:   if  $e > e_h$  then
6:      $e_h = e$ 
7: function ON_NEW_EPOCH_FROM_CONTROLLER(epoch e)
8:   if  $e > e_h$  then
9:      $e_h = e$ 

```

Controller: local epoch e_c

```

10: function ON_NEW_EPOCH(epoch e)
11:   if  $e > e_c$  then
12:      $e_c = e$ 
13:   for all end-host  $h$  do
14:     Send  $e_c$  to  $h$ 

```

Remark. Algorithm 1 is in essence a variant of Lamport Clock (LC) [36], but differs from LC in two aspects. First, LC solely relies on packet embedding for epoch updates, while Algorithm 1 relies on both packet embedding and the controller’s out-band messages to speed up epoch updates. Second, LC involves all entities to embed new epochs for value updates, while Algorithm 1 determines the epochs via the controller and all end-hosts, but not switches. Thus, LC only provides weak consistency, while Algorithm 1 provides hybrid consistency.

Distributed controller. As the network size grows, OmniMon can deploy multiple instances of the controller, each of which manages a finite number of end-hosts. We extend our epoch update mechanism to enable OmniMon to synchronize the epochs among the controller instances. Specifically, in addition to the controller instances that connect to end-hosts (referred to as the *leaf instances*), we also configure a controller instance, called the *root instance*, that communicates with all leaf instances. Once any leaf instance enters a new epoch (on receiving a message from one of its connected end-hosts), it notifies the root instance, which then broadcasts this new epoch to other leaf instances. Each leaf instance also sends the epoch update to its connected end-hosts. This two-layer mechanism takes at most four hops (i.e., two between a leaf instance and an end-host, and another two between a leaf instance and the root instance) to propagate a new epoch. Thus, the time difference between the epoch updates of any two end-hosts remains small and bounded. Our evaluation (Exp#4 in §8) shows that a controller instance can manage 1,024 end-hosts, so our two-layer epoch update mechanism can support the scale of modern data centers.

5 ACCOUNTABILITY

OmniMon achieves accountability by inferring the exact traffic losses on a per-switch, per-flow basis, such that it can recover per-flow values from the *shared slots* in each switch. Our loss inference is related to classical network tomography [8, 22], yet the latter often focuses on only end-host information due to limited in-network visibility and aims to minimize accuracy loss. In contrast, our inference model is network-wide (covering all end-hosts and switches) and aims for full accuracy. Here, we target *additive* flow values, and describe how to convert non-additive flow statistics to additive values in Appendix E.

anisms [2] and failure mitigation mechanisms [23, 68, 79] greatly suppress packet losses. Such loss sparsity implies that packet losses likely occur in exactly one switch within a local region (e.g., a small group of racks or clusters). In other words, for all local lossy flows that traverse solely within a region, their packet losses are attributed to a single switch. Thus, we can infer the variables in \vec{x} for such local flows (recall that we can solve for the unique solution with $m = 1$ for such local flows based on the accountability condition). This further relaxes the rank requirement of A .

Design. The above two observations greatly simplify the linear system by reducing the number of variables in \vec{x} . For the remaining variables in \vec{x} , we aim to construct via flow mapping a new matrix M that has a maximum rank r , so that the simplified linear system can return a unique solution with a high likelihood. However, this problem remains non-trivial, since lossy flows are unpredictable and we cannot tell in advance how lossy flows are mapped to different slots.

OmniMon addresses this issue by *mapping a flow to $d > 1$ slots in each switch*, inspired by Bloom filters [7]. Such a design has two benefits. First, the number of slots mapped by lossy flows increases and has an expected value equal to nd (Appendix). Second, the number of possible flow mappings increases to $O(K^d)$, where K is the total number of available slots in a switch, so that the likelihood that two lossy flows are mapped to the same d slots decreases. Thus, the rank of A increases, as the number of mapped slots (i.e., rows of A) increases and the number of overlapped flow mappings (i.e., linearly dependent columns) decreases. For example, the extreme case in Figure 5 can be avoided.

OmniMon takes a two-phase collaborative approach to ensure that any new flow is mapped to a different set of d slots. In the first phase, the controller generates all $\binom{K}{d}$ possible flow mappings and partitions the flow mappings for all end-hosts during their startups, such that each end-host receives a distinct list of flow mappings. In the second phase, a source end-host selects an unassigned flow mapping when it sees a new flow, and recycles any assigned flow mapping upon flow termination, during flowkey tracking (§3.2). Note that if there are too many active flows, the source end-host may run out of unassigned flow mappings. In this case, we choose the least recently used flow mapping for a new flow.

Summary. OmniMon simplifies the linear system through the characteristics of data center networks to infer the exact packet losses. We summarize the steps below.

- *Step 1:* Identify any lossy flow whose difference of flow values in source and destination end-hosts is non-zero.
- *Step 2:* Construct M_t (and hence A) and \vec{b} .
- *Step 3:* Refine the variables in \vec{x} . With traffic locality, if a flow does not traverse a switch, we set the corresponding variable as zero. With loss sparsity, we examine every local region (by racks or clusters). If a region has exactly one switch with packet losses, we identify all local flows in this region and solve for the unique solution for the corresponding variables of these local flows.
- *Step 4:* Solve $A \cdot \vec{x} = \vec{b}$ for the remaining variables.

6 DISCUSSION

We discuss some practical deployment issues of OmniMon.

Scalability. As the network size increases, OmniMon needs more

resources (e.g., memory in switches, CPU power in end-hosts and the controller, etc.) to achieve full accuracy. However, its resource usage required for full accuracy grows at a much lower rate than the growth of the network scale. For end-hosts, since flows are independent, OmniMon can process flows in different CPU cores when a single CPU core cannot process all flows. It remains much more resource-efficient compared to the processing overhead of network stacks, as its hash computations and value updates incur low overhead in commodity CPUs [49]. For switches, as K slots provide $\binom{K}{d}$ flow mappings, a small increment of K can handle much more flows when $d > 1$. OmniMon remains accurate even when the actual number of flows exceeds $\binom{K}{d}$ (Exp#5 in §8). For the controller, it is a logical entity that can support multiple physical servers (§2.1). Also, it only relays limited out-band traffic even for a large number of end-hosts (§4.2), and can accelerate the linear system solving with multiple CPU cores [50, 66].

Parameterization. OmniMon allows simple parameter configurations. For end-hosts, we can adjust resources on demand. For switches, we reuse the slots for different epochs. We need to configure the maximum number of epochs being tracked and the number of slots per epoch. The number of tracked epochs is derived based on the epoch length and the tolerable network delay, both of which are configurable. For example, if we use four epochs of length 100 ms each, we can track the packets with a network delay of up to 400 ms; any packet that is delayed by more than 400 ms is not monitored. The number of slots per epoch depends on the maximum number of flows being tracked (§5.2). Recall that K slots imply $\binom{K}{d}$ possible flow mappings if we map each flow to d slots. Given the maximum number of flows n , we select K and d such that $\binom{K}{d} \geq n$. We now choose $d = 2$ (a larger d implies more computational overhead in switches).

Generality. OmniMon supports various combinations of packet fields as flowkey definitions, as well as different types of flow statistics (e.g., the 344 statistics in [35, 40, 46, 71]). However, since OmniMon works at the granularity of epochs, it cannot perform fine-grained collective analysis for packet timestamps. Instead, it now provides coarse-grained timestamp measurement. For example, it estimates the flow completion time in number of epochs by tracking the first and last epochs where a flow occurs.

Incremental deployment. OmniMon now assumes full deployment across end-hosts and programmable switches. We can also incrementally deploy OmniMon on a subset of end-hosts and switches of interest (e.g., core switches) for partial measurement (i.e., only the traffic that traverses the deployed devices is measured). Characterizing the trade-off between incremental deployment and full accuracy is our future work.

Failure handling. OmniMon needs to handle the failures for different types of entities during measurement. For end-hosts, OmniMon discards the switch slots of any failed end-host, as the flow mappings of those slots are assigned and known by the controller (§5.2). Recovering end-host information is out of our scope, yet it can be done with OS-level fault-tolerance mechanisms [9, 18]. For switches, OmniMon localizes switch failures via loss inference (Exp#8 in Appendix). For the controller, its fault tolerance can be achieved via multiple servers (§2.1).

7 IMPLEMENTATION

We implement a prototype of OmniMon. We present the pseudo-code of all its operations in Appendix.

End-host. We leverage DPDK [16] to remove kernel-space overhead. We use mTCP [31], a user-space network stack, for the communications between end-hosts and the controller. We implement different operations (e.g., packet processing, result reporting, and epoch synchronization) in distinct threads, and encapsulate them as a library that can be plugged into any platforms (e.g., OVS-DPDK [54] and PktGen [58]).

Switch. We implement switch operations in P4 [56] and place the embedded fields between the Ethernet and IP headers. We support both *per-switch* and *per-port* deployments; in *per-port* deployment, each port can be viewed as a logical switch. We allocate a region of registers for all slots, such that a packet can locate its slots in a region based on its embedded switch index, epoch, and switch port (in *per-port* deployment). For efficient slot location, we pre-compute the partial results of some dependent operations (e.g., modulo, left-shift, and add) and store them in match-action tables. For each packet, we perform a table lookup for partial results and compute the final results in ALUs.

Controller. We connect the controller to each end-host and switch with mTCP [31] and ZeroMQ [77], respectively. The controller parallelizes its operations via multi-threading. It solves the linear system (§5) with the eigen library [19].

Query engine. To support expressive network telemetry, we implement a basic query engine based on Sonata [24]. Specifically, the query engine transforms a telemetry task with an SQL-like expression in Sonata. It then invokes Sonata’s compiler to translate the task into different operators. It extracts the relevant flowkeys and statistics from the operators. Finally, it generates a specific configuration for OmniMon that monitors only the extracted flowkeys and statistics.

8 EVALUATION

We conduct testbed experiments that compare OmniMon with 11 state-of-the-art telemetry designs in various aspects. We summarize our findings for OmniMon:

- It achieves higher throughput than Trumpet [49], an end-host-based telemetry system (Exp#1).
- It consumes less switch resource usage than four sketch solutions (which have non-zero errors) [29, 38, 43, 72], while achieving zero errors (Exp#2).
- It completes controller operations in limited time even with 10^6 flows (Exp#3).
- It incurs less switch resource usage and smaller epoch boundary deviations than Speedlight [73], a synchronization solution for network telemetry (Exp#4).
- It infers 99.7% of packet losses in an extreme case (Exp#5).
- It consumes much less resource usage, while achieving better scalability, than SwitchPointer [67], even when the number of flows increases to 10^6 (Exp#6).
- It supports 11 anomaly detection applications, in which it incurs less overhead than Marple [52] and Sonata [24], while achieving zero errors (Exp#7).

In Appendix, we further evaluate OmniMon via simulation in large-scale deployment. It achieves perfect diagnosis in four types of packet losses and outperforms another diagnosis system 007 [4] (Exp#8). It supports performance evaluation for load balancing algorithms, and locates per-link top flows more accurately than HashPipe [65] (Exp#9). Finally, we provide a breakdown of switch resource usage for Exp#2, and per-application resource-accuracy trade-off for Exp#7.

8.1 Setup

Testbed. We deploy our OmniMon prototype in eight servers and three Barefoot Tofino switches [70]. Each server has two 12-core 2.2 GHz CPUs, 32 GB RAM, and a two-port 40 Gbps NIC, while each switch has 32 100 Gb ports. We deploy an end-host in each server and co-locate the controller with one end-host server. We deploy two switches as edge switches connected by four servers each and the remaining switch as the core switch connected by both edge switches.

Workloads. We pick a one-hour trace in CAIDA 2018 [11]. To simulate the locality of data center traffic, we map the IP addresses in the trace to our end-hosts and generate a workload in which 85% of flows are edge-local, as reported in [60]. We use different epoch lengths, such that the number of active flows per end-host ranges from 10^3 to 10^6 (much more than reported in field studies [2, 60]). We use PktGen [58] to generate trace workloads in end-hosts. To remove disk I/O overhead, each end-host loads traces (with modified IP addresses) into the memory buffer of PktGen, and emits traffic as fast as possible to stress-test our prototype.

Methodology. OmniMon counts nine types of per-flow statistics: packet counts, byte counts, SYN counts, FIN/RST counts, SYN-ACK counts, ACK counts, 40-byte packet counts, number of packets with special strings (e.g., “zorro”), and the latest epoch of a flow. Each switch deploys four groups of slots. Each group has 3,072 slots for one epoch, and is recycled when an epoch completes. We map each flow to $d = 2$ slots (§6). Thus, there are $\binom{3072}{2} \approx 4.7 \times 10^6$ flow mappings per epoch, which sufficiently accommodate all active flows (note that a flow appears in both source and destination end-hosts, so there are at most 4×10^6 active flows in the network assuming that each of the eight end-hosts has at most 10^6 active flows). Each plot shows the average results over 100 runs, and omits error bars as they are negligible.

8.2 Results

(Exp#1) End-host overhead. We first evaluate the end-host overhead by comparing OmniMon with Trumpet [49], a zero-error end-host-based telemetry system. Note that Trumpet is designed for only end-hosts, while OmniMon also provides switch-side information. The open-source Trumpet prototype (called “Trumpet-Full (TF)”) realizes 4,096 triggers for different network event matchings. We also implement a simplified variant (called “Trumpet-Simple (TS)”) that uses a single trigger to track the same nine flow statistics as OmniMon for fair comparisons. Figure 6(a) shows the throughput of OmniMon (OM), the two Trumpet variants, and the packet generation using PktGen on a single CPU core. TF achieves only 6.5 Mpps as it performs matching on 4,096 triggers. TS is faster than TF, but is slower than OmniMon due to event matching over-

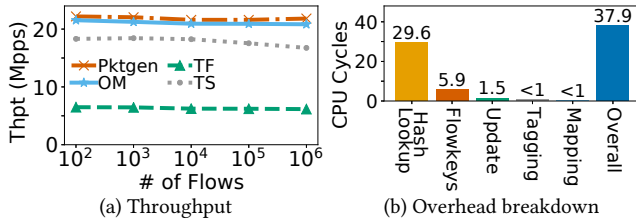


Figure 6: (Exp#1) End-host overhead.

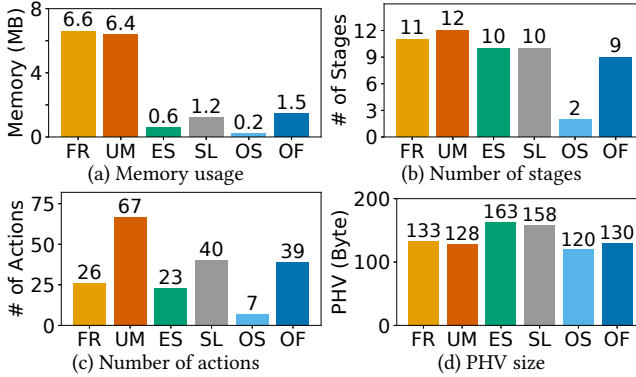


Figure 7: (Exp#2) Switch resource usage.

head. OmniMon only has a slight throughput drop of 5.4% (for 10⁶ flows) compared to PktGen. Figure 6(b) shows a breakdown of the number of CPU cycles for per-packet processing in OmniMon. The throughput drop mainly comes from per-packet hash lookups.

(Exp#2) Switch resource usage. We compare OmniMon in switch resource usage with four sketch solutions: FlowRadar (FR) [38], UniMon (UM) [43], ElasticSketch (ES) [72], and SketchLearn (SL) [29]. All sketch solutions perform approximate measurement using a compact data structure to fit in limited switch memory. We re-build them to measure per-flow packet counts based on their published designs in P4 in our switches. For fair comparisons, we build a simplified OmniMon (OS) that measures only per-flow packet counts. We compare all solutions with our full OmniMon version (OF) that tracks nine flow statistics.

Figure 7 depicts the switch resource usage, in memory usage, numbers of stages and actions (which measure computational resources), and packet header vector (PHV) size (which measures the message size passed across stages). Both FR and UM maintain extra data structures (e.g., bloom filters and heaps) for flowkey recovery, and hence incur high memory usage (Figure 7(a)). Also, sketch solutions perform multiple hash computations that trigger many stages and actions (Figures 7(b) and 7(c)). ES and SL update different parts of counters and incur high PHV sizes (Figure 7(d)). For OmniMon, OS is resource-efficient since its collaborative design imposes only value updates in switches; even for OF, which measures more statistics, its switch resource usage is comparable to the sketch solutions. Note that OmniMon achieves full accuracy, but sketch solutions cannot.

(Exp#3) Controller overhead. We evaluate the controller overhead for flow mapping, epoch synchronization, and collective analysis in our testbed. Recall that we run these operations in distinct threads (§7). Here, we configure each thread to run in a single

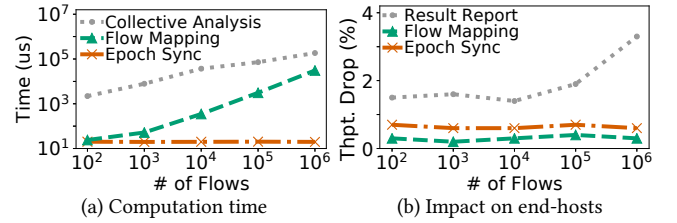


Figure 8: (Exp#3) Controller overhead.

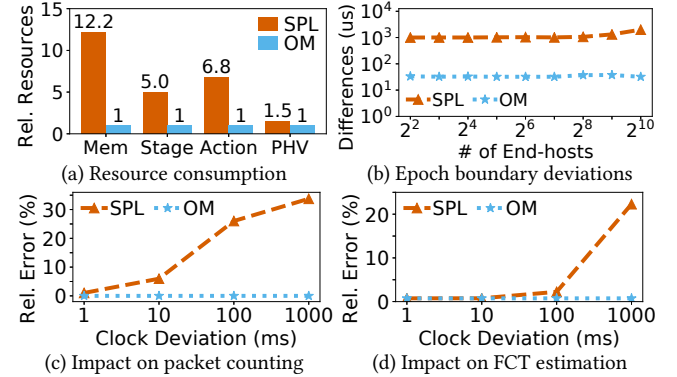


Figure 9: (Exp#4) Consistency.

CPU core. Figure 8(a) shows that epoch synchronization can be completed within 20 μ s, which is critical for epoch boundary deviations. Flow mapping takes more time, but it is invoked only at the startup. The most expensive operation is collective analysis, whose completion time is 180 ms for 10⁶ flows. We find that its major overhead comes from solving the linear system. Nevertheless, we can speed up the computation with more CPU cores [50, 66]. Figure 8(b) shows the impact of these operations on end-hosts by comparing the end-host throughput with and without these operations. All operations show limited throughput drop as we isolate them in distinct threads from packet processing (§7).

(Exp#4) Consistency. We evaluate epoch synchronization by comparing OmniMon with the open-source Speedlight (SPL) prototype [73], a synchronization solution for network telemetry. We adapt the SPL prototype, written in P4 BMv2 [55], to our Tofino switches. Figure 9(a) shows the switch resource usage, normalized to that of the full OmniMon version. SPL realizes the Chandy-Lamport algorithm [13] and incurs high overhead, while OmniMon is more resource-efficient as its epoch synchronization is done by end-hosts and the controller.

We examine the time deviations in per-epoch boundaries when the local clocks are loosely synchronized. We deploy multiple (up to 1,024) end-host instances in each server. We tune the local clocks of end-hosts to differ by 1 ms. Figure 9(b) shows that SPL suffers from a large deviation (nearly to the clock difference 1 ms) since it relies on only in-band packets for synchronization. If the traffic does not span all end-hosts, some end-hosts remain in the old epoch. OmniMon limits the deviation to 60 μ s even for 1,024 end-hosts, as it leverages the controller to bound the deviation and uses mTCP to bypass kernel space (§7).

We further evaluate the accuracy of epoch synchronization in a straggler case. We configure one straggler end-host whose local

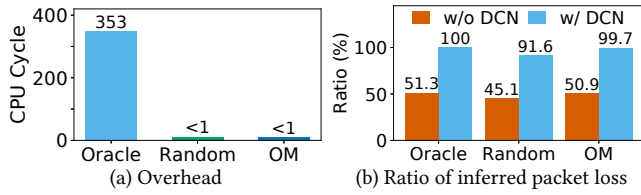


Figure 10: (Exp#5) Accountability.

clock differs from those of others by a range of 1 ms to 1 s. Figure 9(c) measures the average relative error of per-flow packet counts in this straggler end-host. For SPL, the error increases to 26% and 33% when the difference is 100 ms and 1 s, respectively, as the epoch boundary deviation maps many packets to a wrong epoch. OmniMon retains zero errors as it includes every packet in the same epoch across all end-hosts and switches. Figure 9(d) measures the average relative error of flow completion time. We focus on long-lived flows that exceed 1 s (which often imply important network events [14]). Although OmniMon measures coarse-grained timestamps in epochs (§6), its error is only 0.7%, as it bounds the epoch boundary deviation. For SPL, the error is above 22% when the clock deviation is 1 s.

(Exp#5) Accountability. Our default setup provisions sufficient slots to allocate each flow a unique flow mapping. Here, we consider an extreme case, in which we use 512 slots per switch (i.e., $\binom{512}{2} \approx 130\text{K}$ flow mappings). We set the epoch length to two minutes, so that an epoch now has 380K concurrent flows on average. We also set a packet drop rate of 0.5%. We compare OmniMon with two variants: (i) “Oracle”, which knows all lossy flows in advance and maps the lossy flows for minimum slot overlaps to have the maximum matrix rank; and (ii) “Random”, which randomly selects a slot for each flow. Figure 10(a) shows the number of CPU cycles for each invocation of the three schemes. OmniMon and Random consume less than one CPU cycle, while Oracle requires over 300 cycles per packet for maximum-rank mappings.

We also evaluate the fractions of per-switch, per-flow packet loss variables that can be solved via the four steps in §5.2. For each of the three schemes, we build the linear system using Steps 1 and 2, and consider two variants: (i) “w/o DCN”, in which we skip Step 3 and solve directly the linear system; and (ii) “w/ DCN”, in which we execute Step 3 that simplifies the linear system with data center network characteristics. Figure 10(b) shows that without Step 3, all schemes infer only around 50% of variables, as the number of lossy flows (i.e., n) exceeds the total number of available slots (i.e., 512 per switch). With the simplification of Step 3, Oracle infers all variables, while OmniMon infers 99.7% of variables. Here, OmniMon cannot achieve 100% of inference as we consider an extreme scenario, yet full accuracy is achievable in practical scenarios (Exp#8 in Appendix).

(Exp#6) Scalability. We compare OmniMon with the open-source SwitchPointer prototype [67] (SP) in scalability. We adapt the SP prototype, written in P4 BMv2, to our Tofino switches. Figure 11(a) shows the switch resource usage normalized to that of OmniMon. SP focuses on in-network visibility, while the resource-accuracy trade-off is not its focus. It achieves zero errors with perfect hashing, but incurs higher switch resource overhead for the complicated calculations. Note that perfect hashing needs to acquire all possible

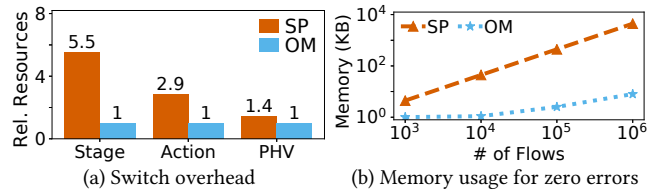


Figure 11: (Exp#6) Scalability.

flowkeys a priori to design perfect hash functions. This is feasible for IP addresses, but not for 5-tuples (with up to 2^{104} flows). Figure 11(b) compares the memory usage needed for zero-error guarantees. The memory usage of SP grows linearly with the number of flowkeys as each slot is associated with one flow, while that of OmniMon increases at a much slower rate, as K slots can accommodate $\binom{K}{d}$ flows (§5.2).

(Exp#7) End-to-end performance comparison. We compare OmniMon with the switch-controller architectures of Sonata [24] and Marple [52]. Our end-to-end performance comparison addresses 11 anomaly detection applications reported from Sonata [24]. The comparison includes three aspects: (i) per-switch resource usage; (ii) per-switch traffic sent to the controller; and (iii) per-application detection accuracy. We do not compare the end-host overhead, as Sonata and Marple do not involve end-hosts in their designs. As shown in Exp#1 and Exp#3, OmniMon has limited end-host overhead.

OmniMon tracks nine flow statistics to support the 11 applications. For Sonata, as in §2.2, we consider both “Sonata-FA (S-FA)” (with full accuracy) and “Sonata-AD (S-AD)” (with accuracy drop), in which every switch-side operator is configured with 2^{16} counters as in its open-source prototype. For Marple, we implement its key-value cache and evict old keys to handle hash collisions. We configure 2^{16} cache slots as in Marple [52]. For OmniMon, we employ a full configuration that includes all nine statistics (“OM-F”). We also employ the Sonata-based query engine (§7) to generate a configuration for each application (“OM-Q”). We set 10^5 flows at each epoch. Since Sonata, Marple, and OM-Q deploy each application individually, we average their results of the 11 applications (see Appendix for per-application results).

Figure 12(a) compares per-switch resource usage, normalized to that of OM-F. Sonata-AD sacrifices accuracy and consumes the least memory, stages, and actions. Sonata-FA and Marple require more resources for flowkey tracking and cache eviction. OM-F uses less switch memory usage than Sonata-FA and Marple, while tracking all nine statistics, due to its collective analysis (§3.2). Its usage of the other three types of resources is comparable to the single application deployment in Sonata-FA and Marple, since it performs only (simple) value update operations in switches. With the query engine, the resource usage of OM-Q drops further.

Figure 12(b) compares the per-switch traffic sent to the control plane. Sonata-FA and Marple preserve the accuracy by evicting colliding flows to the controller, but trigger much traffic as hash collisions frequently occur with numerous flows. Sonata-AD reduces per-switch, per-application traffic to 5 MB/s, but incurs accuracy drop. In contrast, OM-F triggers 180 KB/s of traffic without accuracy drop. The query engine further reduces the traffic (26 KB/s) as it tracks only the statistics of interest for each application.

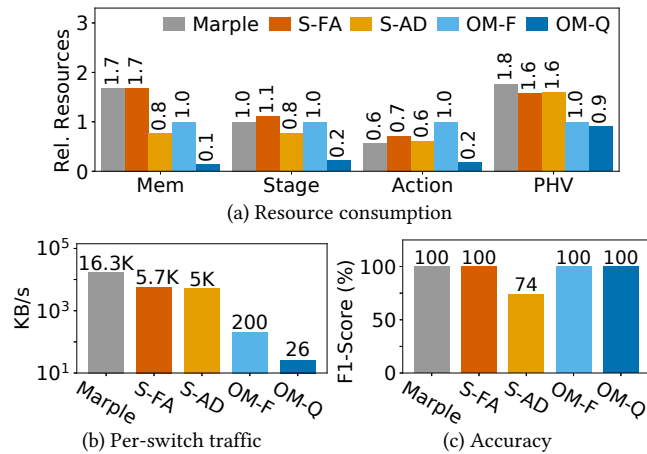


Figure 12: (Exp#7) End-to-end performance comparison.

Figure 12(c) compares the detection accuracy quantified by the F1-score. The F1-score of Sonata-AD is only around 74%, as it does not resolve the hash collisions in the data plane for resource efficiency. Marple and Sonata-FA improve the accuracy, but incur high resource overhead in both switches and the controller. OmniMon achieves both resource efficiency and full accuracy.

9 RELATED WORK

Resource-accuracy trade-off. Existing network telemetry proposals often make the resource-accuracy trade-off. End-host-based telemetry systems [1, 49] perform complete per-flow tracking, but have limited network-wide visibility and incur excessive memory usage. Switch-based approaches either relax accuracy [47] or restrict to specific queries [51]. Approximation algorithms [5, 26, 28, 29, 38, 43, 48, 62, 63, 65, 72, 75] trade accuracy for resource efficiency. Some systems mirror traffic to the controller, in a best-effort manner [59] or on pre-defined traffic patterns [25, 52, 69, 74, 76, 79], but inevitably miss the traffic that is not mirrored.

Recent hybrid architectures combine different entities for network telemetry [24, 32, 42, 52, 67]. To mitigate resource burdens, they rely on sampling [42] and/or event matching [24, 42, 52, 67] to focus on only the traffic of interest, and hence cannot achieve full accuracy. Also, they do not address consistency and accountability in entity coordination.

Consistency. HUYGENS [21] proposes software clock synchronization, but its machine-learning design is complicated in switch deployment. Swing State [44] ensures consistent data-plane states via state migration, but poses heavy state update costs. Synchronized Network Snapshot [73] realizes causal consistency (a form of weak consistency) in P4 switches. OmniMon achieves hybrid consistency and is more lightweight (Exp#4 in §8).

Accountability. Loss inference is well studied in network telemetry. Active probing [23, 68] sends probe packets to measure path losses, but only covers a subset of paths. Passive monitoring analyzes network traffic, but relies on domain knowledge [79] and/or employs approximation techniques [4, 22, 39, 61] to estimate loss rates. OmniMon exploits the complete statistics in end-hosts as the ground truth to form a linear system for exact packet loss inference.

10 CONCLUSION

OmniMon is a network telemetry architecture for large-scale data centers. Its design principle is to re-architect the network-wide collaboration and carefully coordinate telemetry operations among all entities. We show how OmniMon achieves both resource efficiency and full accuracy, with consistency and accountability guarantees. Experiments show the effectiveness of OmniMon over 11 state-of-the-art telemetry designs; more results are available in Appendix.

ACKNOWLEDGEMENTS

We thank the anonymous shepherd and reviewers for their valuable comments. The work was supported in part by National Key R&D Program of China (2019YFB1802600), National Natural Science Foundation of China (61802365), and “FANet: PCL Future Greater-Bay Area Network Facilities for Large-scale Experiments and Applications (No. LZC0019)”, and Research Grants Council of Hong Kong (GRF 14204017).

REFERENCES

- [1] O. Alipourfard, M. Moshref, Y. Zhou, T. Yang, and M. Yu. A Comparison of Performance and Accuracy of Measurement Algorithms in Software. In *Proc. of ACM SOSR*, 2018.
- [2] M. Alizadeh, A. Greenberg, D. A. Maltz, and J. Padhye. Data Center TCP (DCTCP). In *Proc. of ACM SIGCOMM*, 2010.
- [3] M. Alizadeh, N. Yadav, G. Varghese, T. Edsall, S. Dharmapurikar, R. Vaidyanathan, K. Chu, A. Fingerhut, V. T. Lam, F. Matus, and R. Pan. CONGA: Distributed Congestion-Aware Load Balancing for Datacenters. In *Proc. of ACM SIGCOMM*, 2014.
- [4] B. Arzani, S. Ciraci, L. Chamon, Y. Zhu, H. Liu, J. Padhye, B. T. Loo, and G. Outhred. 007: Democratically Finding The Cause of Packet Drops. In *Proc. of USENIX NSDI*, 2018.
- [5] R. Ben Basat, G. Einziger, R. Friedman, M. C. Luizelli, and E. Waisbard. Constant Time Updates in Hierarchical Heavy Hitters. In *Proc. of ACM SIGCOMM*, 2017.
- [6] T. Benson, A. Akella, and D. A. Maltz. Network Traffic Characteristics of Data Centers in the Wild. In *Proc. of ACM IMC*, 2010.
- [7] B. H. Bloom. Space/Time Trade-offs in Hash Coding with Allowable Errors. *Communications of ACM*, 13(7), 1970.
- [8] Bowden, Rhys Alistair and Roughan, Matthew and Bean, Nigel. Network Link Tomography and Compressive Sensing. In *Proc. of SIGMETRICS*, 2011.
- [9] T. C. Bressoud and F. B. Schneider. Hypervisor-based Fault Tolerance. In *Proc. of SOSP*, 1995.
- [10] Broadcom Tomahawk Ethernet Switch Series. <https://www.broadcom.cn/products/ethernet-connectivity/switching/strataxgs/bcm56960-series>.
- [11] Caida Anonymized Internet Traces 2018 Dataset. http://www.caida.org/data/passive/passive_dataset.xml.
- [12] J. D. Case, M. Fedor, M. L. Schoffstall, and J. Davin. Simple Network Management Protocol (SNMP), 1990.
- [13] K. M. Chandy and L. Lamport. Distributed Snapshots: Determining Global States of Distributed Systems. *ACM Transactions on Computer Systems*, 3(1):63–75, 1985.
- [14] A. Chen, Y. Jin, J. Cao, and L. E. Li. Tracking Long Duration Flows in Network Traffic. In *Proc. of INFOCOM*, 2010.
- [15] Cisco Tetration. <https://www.cisco.com/c/en/us/products/data-center-analytics/tetration-analytics>.
- [16] Data Plane Development Kit. <https://dpdk.org>.
- [17] M. Dobrescu, N. Egi, K. Argyraki, B.-g. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh, and S. Ratnasamy. RouteBricks: Exploiting Parallelism To Scale Software Routers. In *Proc. of ACM SOSP*, 2009.
- [18] G. W. Dunlap, S. T. King, S. Cinar, M. A. Basrai, and P. M. Chen. ReVirt: Enabling Intrusion Analysis through Virtual-Machine Logging and Replay. In *Proc. of OSDI*, 2002.
- [19] Eigen. <http://eigen.tuxfamily.org>.
- [20] S. K. Fayazbakhsh, L. Chiang, V. Sekar, M. Yu, and J. C. Mogul. Enforcing Network-Wide Policies in the Presence of Dynamic Middlebox Actions using FlowTags. In *Proc. of NSDI*, 2014.
- [21] Y. Geng, S. Liu, Z. Yin, A. Naik, B. Prabhakar, M. Rosenblum, and A. Vahdat. Exploiting a Natural Network Effect for Scalable, Fine-grained Clock Synchronization. In *Proc. of USENIX NSDI*, 2018.
- [22] D. Ghita, K. Argyraki, and P. Thiran. Network Tomography on Correlated Links. In *Proceedings of ACM IMC*, 2010.

- [23] C. Guo, L. Yuan, D. Xiang, Y. Dang, R. Huang, D. Maltz, Z. Liu, V. Wang, B. Pang, H. Chen, Z.-w. Lin, and V. Kurien. Pingmesh : A Large-Scale System for Data Center Network Latency Measurement and Analysis. *Proc. of ACM SIGCOMM*, 2015.
- [24] A. Gupta, R. Harrison, M. Canini, N. Feamster, J. Rexford, and W. Willinger. Sonata: Query-Driven Streaming Network Telemetry. In *Proc. of ACM SIGCOMM*, 2018.
- [25] N. Handigol, B. Heller, V. Jeyakumar, D. Mazières, and N. McKeown. I Know What Your Packet Did Last Hop: Using Packet Histories to Troubleshoot Networks. In *Proc. of USENIX NSDI*, 2014.
- [26] R. Harrison, Q. Cai, A. Gupta, and J. Rexford. Network-Wide Heavy Hitter Detection with Commodity Switches. In *Proc. of ACM SOSR*, 2018.
- [27] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. In *Proc. of ACM SIGCOMM*, 2013.
- [28] Q. Huang, X. Jin, P. P. C. Lee, R. Li, L. Tang, Y.-C. Chen, and G. Zhang. SketchVisor: Robust Network Measurement for Software Packet Processing. In *Proc. of ACM SIGCOMM*, 2017.
- [29] Q. Huang, P. P. C. Lee, and Y. Bao. SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference. In *Proc. of ACM SIGCOMM*, 2018.
- [30] P. Hunt, M. Konar, F. P. Junqueira, and B. Reed. ZooKeeper: Wait-free Coordination for Internet-scale Systems. In *Proc. of USENIX ATC*, pages 11--11, 2010.
- [31] E. Jeong, S. Wood, M. Jamshed, H. Jeong, S. Ihm, D. Han, and K. Park. mTCP: a Highly Scalable User-level TCP Stack for Multicore Systems. In *Proc. of USENIX NSDI*, 2014.
- [32] V. Jeyakumar, M. Alizadeh, Y. Geng, C. Kim, and D. Mazières. Millions of Little Minions: Using Packets for Low Latency Network Programming and Visibility. In *Proc. of ACM SIGCOMM*, 2014.
- [33] Khandelwal, Anurag and Agarwal, Rachit and Stoica, Ion. Confluo: Distributed Monitoring and Diagnosis Stack for High-Speed Networks. In *Proc. of USENIX NSDI*, 2019.
- [34] D. Kim, Y. Zhu, C. Kim, J. Lee, and S. Seshan. Generic External Memory for Switch Data Planes. In *Proc. of ACM HotNets*, 2018.
- [35] H. Kim, K. Claffy, M. Fomenkov, D. Barman, M. Faloutsos, and K. Lee. Internet Traffic Classification Demystified : Myths , Caveats , and the Best Practices. In *Proc. of CoNEXT*, 2008.
- [36] L. Lamport. Time, Clocks, and the Ordering of Events in a Distributed System. *Communications of the ACM*, 21(7):558--565, 1978.
- [37] L. Lamport. Paxos Made Simple. *ACM SIGACT News (Distributed Computing Column)*, pages 51--58, 2001.
- [38] Y. Li, R. Miao, C. Kim, and M. Yu. FlowRadar: A Better NetFlow for Data Centers. In *Proc. of USENIX NSDI*, 2016.
- [39] Y. Li, R. Miao, C. Kim, and M. Yu. LossRadar: Fast Detection of Lost Packets in Data Center Networks. In *Proc. of ACM CoNEXT*, 2016.
- [40] Y.-s. Lim, H.-c. Kim, J. Jeong, C.-k. Kim, T. T. Kwon, and Y. Choi. Internet Traffic Classification Demystified : On the Sources of the Discriminative Power. In *Proc. of CoNEXT*, 2010.
- [41] Linux Netfilter. <https://www.netfilter.org>.
- [42] X. Liu, M. Shirazipour, M. Yu, and Y. Zhang. MOZART: Temporal Coordination of Measurement. In *Proc. of ACM SOSR*, 2016.
- [43] Z. Liu, A. Manousis, G. Vorsanger, V. Sekar, and V. Braverman. One Sketch to Rule Them All: Rethinking Network Flow Monitoring with UnivMon. In *Proc. of ACM SIGCOMM*, 2016.
- [44] S. Luo, H. Yu, and L. Vanbever. Swing State: Consistent Updates for Stateful and Programmable Data Planes. In *Proc. of ACM SOSR*, 2017.
- [45] R. Miao, H. Zeng, C. Kim, J. Lee, and M. Yu. SilkRoad: Making Stateful Layer-4 Load Balancing Fast and Cheap Using Switching ASICs. In *Proc. of ACM SIGCOMM*, 2017.
- [46] A. W. Moore and D. Zuev. Internet Traffic Classification Using Bayesian Analysis Techniques. In *Proc. of SIGMETRICS*, 2005.
- [47] M. Moshref, M. Yu, R. Govindan, and A. Vahdat. DREAM: Dynamic Resource Allocation for Software-defined Measurement. In *Proc. of ACM SIGCOMM*, 2014.
- [48] M. Moshref, M. Yu, R. Govindan, and A. Vahdat. SCREAM: Sketch Resource Allocation for Software-defined Measurement. In *Proc. of ACM CoNEXT*, 2015.
- [49] M. Moshref, M. Yu, R. Govindan, and A. Vahdat. Trumpet: Timely and Precise Triggers in Data Centers. In *Proc. of ACM SIGCOMM*, 2016.
- [50] MUMPS. <http://mumps.enseeit.fr>.
- [51] S. Narayana, M. T. Arashloo, J. Rexford, and D. Walker. Compiling Path Queries. In *Proc. of USENIX NSDI*, 2016.
- [52] S. Narayana, A. Sivaraman, V. Nathan, P. Goyal, V. Arun, M. Alizadeh, V. Jeyakumar, and C. Kim. Language-Directed Hardware Design for Network Performance Monitoring. In *Proc. of ACM SIGCOMM*, 2017.
- [53] NetFlow. <https://www.ietf.org/rfc/rfc3954.txt>.
- [54] OpenvSwitch. <http://openvswitch.org>.
- [55] P4 Behavioral Model v2. <https://github.com/p4lang/behavioral-model>.
- [56] P4 Language. <https://p4.org>.
- [57] R. Pagh and F. F. Rodler. Cuckoo Hashing. *Journal of Algorithms*, 51(2):121--133, 2001.
- [58] PktGen. <https://pktgen-dpdk.readthedocs.io>.
- [59] J. Rasley, B. Stephens, C. Dixon, E. Rozner, W. Felter, K. Agarwal, J. Carter, and R. Fonseca. Planck: Millisecond-scale Monitoring and Control for Commodity Networks. In *Proc. of ACM SIGCOMM*, 2014.
- [60] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren. Inside the Social Network's (Datacenter) Network. *Proc. of ACM SIGCOMM*, 2015.
- [61] A. Roy, H. Zeng, J. Bagga, and A. C. Snoeren. Passive Realtime Datacenter Fault Detection and Localization. In *Proc. of USENIX NSDI*, 2017.
- [62] V. Sekar, M. K. Reiter, W. Willinger, H. Zhang, R. R. Kompella, and D. G. Andersen. cSAMP: A System for Network-Wide Flow Monitoring. In *Proc. of USENIX NSDI*, 2008.
- [63] V. Sekar, M. K. Reiter, and H. Zhang. Revisiting the Case for a Minimalist Approach for Network Flow Monitoring. In *Proc. of ACM IMC*, 2010.
- [64] sFlow. <http://www.sflow.org/>.
- [65] V. Sivaraman, S. Narayana, O. Rottenstreich, S. Muthukrishnan, and J. Rexford. Heavy-Hitter Detection Entirely in the Data Plane. In *Proc. of ACM SOSR*, 2017.
- [66] SuperLU. <http://crd-legacy.lbl.gov/~xiaoye/SuperLU>.
- [67] P. Tammana, R. Agarwal, and M. Lee. Distributed Network Monitoring and Debugging with SwitchPointer. In *Proc. of USENIX NSDI*, 2018.
- [68] C. Tan, Z. Jin, C. Guo, T. Zhang, H. Wu, K. Deng, D. Bi, and D. Xiang. NetBouncer: Active Device and Link Failure Localization in Data Center Networks. In *Proc. of USENIX NSDI*, 2019.
- [69] O. Tilmans, T. Bühler, I. Poesse, S. Vissicchio, and L. Vanbever. Stroboscope: Declarative Traffic Mirroring on a Budget. In *Proc. of USENIX NSDI*, 2018.
- [70] Tofino. <https://www.barefootnetworks.com/products/brief-tofino/>.
- [71] N. Williams, S. Zander, and G. Armitage. A Preliminary Performance Comparison of Five Machine Learning Algorithms for Practical IP Traffic Flow Classification. *SIGCOMM Comput. Commun. Rev.*, 36(5):5--16, 2006.
- [72] T. Yang, J. Jiang, P. Liu, Q. Huang, J. Gong, Y. Zhou, R. Miao, X. Li, and S. Uhlig. Elastic Sketch: Adaptive and Fast Network-wide Measurements. In *Proc. of ACM SIGCOMM*, 2018.
- [73] N. Yaseen, J. Sonchack, and V. Liu. Synchronized Network Snapshots. In *Proc. of ACM SIGCOMM*, 2018.
- [74] D. Yu, Y. Zhu, B. Arzani, R. Fonseca, T. Zhang, K. Deng, and L. Yuan. dShark: A General, Easy to Program and Scalable Framework for Analyzing In-network Packet Traces. In *Proc. of USENIX NSDI*, 2019.
- [75] M. Yu, L. Jose, and R. Miao. Software Defined Traffic Measurement with OpenSketch. In *Proc. of USENIX NSDI*, 2013.
- [76] L. Yuan, C.-N. Chuah, and P. Mohapatra. ProgME: Towards Programmable Network Measurement. In *Proc. of ACM SIGCOMM*, 2007.
- [77] ZeroMQ. <http://zeromq.org>.
- [78] Q. Zhang, V. Liu, H. Zeng, and A. Krishnamurthy. High-Resolution Measurement of Data Center Microbursts. In *Proc. of ACM IMC*, 2017.
- [79] Y. Zhu, N. Kang, J. Cao, A. Greenberg, G. Lu, R. Mahajan, D. Maltz, L. Yuan, M. Zhang, B. Y. Zhao, and H. Zheng. Packet-Level Telemetry in Large Datacenter Networks. In *Proc. of ACM SIGCOMM*, 2015.

APPENDICES

Appendices are supporting materials that have not been peer-reviewed.

Appendix A: Proofs for §4

We formally prove the properties of hybrid consistency in five steps.

Step 1. We start with a lemma regarding the epoch of each packet in Lemma 1.

LEMMA 1. *Every packet is included at the same epoch during its transmission.*

PROOF. Recall that the epoch of a packet is decided only once in the source end-host and embedded in the packet. The embedded epoch will not be modified during the packet transmission. Every entity along the packet transmission path will follow this decision and monitor the packet at the epoch that is being embedded. The lemma follows. \square

Step 2. We next examine the specific weak consistency model that OmniMon achieves. We first define the concept of *partial ordering*.

In our context, we say that two packets A and B exhibit a partial order “ A before B ” if and only if either of the following two cases happens: (i) A and B are sent by the same end-host and A is sent before B , or (ii) A and B are sent by different end-hosts, but B is sent after its source end-host receives A . Let $Epoch(P)$ denote the epoch at which some packet P resides. Lemma 2 proves that OmniMon always achieves a specific weak consistency model based on the *Clock Condition* [36].

LEMMA 2. *If two packets A and B exhibit a partial order “ A before B ”, then their embedded epochs must satisfy $Epoch(A) \leq Epoch(B)$ in all entities.*

PROOF. We examine the two cases under the definition of partial ordering. In Case (i), A and B are sent by the same end-host and A is sent before B . In this case, based on the local clock in the source end-host, we must have $Epoch(A) \leq Epoch(B)$. In Case (ii), A and B are sent by different end-hosts, but B is sent after its source end-host receives A . Suppose that B 's source end-host receives A at its local epoch (Epoch b), where $b \leq Epoch(B)$. If $Epoch(A) \leq b$, then clearly $Epoch(A) \leq Epoch(B)$; otherwise, B 's source end-host will update its local epoch to $Epoch(A)$, so $Epoch(B) \geq b = Epoch(A)$. \square

Step 3. Lemma 3 provides an upper bound of the time difference of an epoch among all end-hosts.

LEMMA 3. *Let Δ be the maximum delay between an end-host and the controller. The time difference of an epoch among all end-hosts is at most 2Δ .*

PROOF. We examine the time difference between the first end-host and the last end-host that enter the same epoch. For any Epoch x , suppose that the first end-host enters Epoch x at time t_s . Let t_l , t_c , and t_p be the times at which the last end-host receives the epoch notification from its local clock, the controller, and the first embedded packet, respectively. Then the time difference is

$$\min\{t_l, t_c, t_p\} - t_s \leq t_c - t_s.$$

Since the first end-host sends a message to the controller immediately at t_s , $t_c - t_s$ is at most the round-trip delay between an end-host and the controller, i.e., 2Δ . The lemma follows. \square

Step 4. With the bounded time difference of an epoch, Lemma 4 shows that all end-hosts reside at the same epoch (i.e., strong consistency) at most times.

LEMMA 4. *Let Δ be the maximum delay between an end-host and the controller. Let ℓ be the length of an epoch. If $\ell > 2\Delta$, OmniMon achieves strong consistency for a duration of no less than $\ell - 2\Delta$ at each epoch.*

PROOF. For any Epoch x , let t_1 and t_2 be the times at which the first and last end-hosts that enter Epoch x , respectively. Based on Lemma 3, we have $t_2 - t_1 \leq 2\Delta$. Also, let h be the first end-host that enters Epoch $x + 1$. Let t_3 and t_4 be the times at which h enters Epoch x and Epoch $x + 1$, respectively. By definition, we have $t_4 - t_3 = \ell$ and $t_1 \leq t_3 \leq t_2$. Now, we compute the following.

$$t_4 - t_2 = \ell + t_3 - t_2 \geq \ell + t_1 - t_2 \geq \ell - 2\Delta.$$

Note that $t_4 - t_2$ represents the duration during which all end-hosts are synchronized at Epoch x . The lemma follows. \square

Remark. If ℓ is sufficiently larger than Δ , OmniMon achieves strong consistency at most times.

Step 5. Finally, Lemma 5 bounds the number of out-band messages triggered per epoch in OmniMon.

LEMMA 5. *Suppose that OmniMon has u end-hosts. For each epoch, the controller and each end-host process $O(u)$ and $O(1)$ out-band messages, respectively, for epoch synchronization.*

PROOF. At each epoch, the controller receives at most u out-band messages, each from one end-host. However, the controller only relays the first received out-band message to all u end-hosts. Thus, the number of processed (including received and relayed) out-band messages in the controller is at most $2u$.

Also, at each epoch, an end-host sends at most one out-band message to the controller (only when its epoch update is triggered by the local clock). Also, it receives one message from the controller based on the above discussion. Thus, the overhead is $O(1)$. \square

Remark. We consider an extreme scenario with 10K end-hosts and an epoch length of 100 ms. Suppose that an out-band message is of size 64 bytes (including Ethernet, IP and TCP headers). This leads to at most $2 \times 10K \times (1000/100) * 64 = 12.8$ MB/s bandwidth, which is quite low in a data center network with 10K end-hosts.

Appendix B: Proofs for §5

We present proofs for the following two properties: (i) The rank of A has a tight upper bound $r(m - 1) + n$, and (ii) the number of mapped slots has an expected value equal to nd .

Rank of A . We prove the tight upper bound of the rank of matrix A . Recall that we construct a matrix M that comprises the rows of each M_t ($1 \leq t \leq m$), and the rank of M is r . We show that $r(m - 1) + n$ is the tight upper bound of the rank of A .

We first construct a large matrix

$$A^* = \begin{bmatrix} M & 0 & \cdots & 0 \\ 0 & M & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & M \\ I & I & \cdots & I \end{bmatrix}.$$

Obviously, A is a sub-matrix of A^* . Thus, the rank of A is tightly bounded by the rank of A^* . The remaining problem is to show that the rank of A^* is $r(m - 1) + n$.

Our proof builds on the definition of rank: the rank of a matrix is equal to the maximum number of linearly independent rows or linearly independent columns in the matrix. Here, we focus on the linearly independent rows (LIRs). Our idea is to first construct $r(m - 1) + n$ LIRs for A^* , and then show that any $r(m - 1) + n + 1$ rows in A^* must be linearly dependent.

LEMMA 6. *There exist $r(m - 1) + n$ LIRs in A^* .*

PROOF. We construct the LIRs as follows. Note that the rank of M equals r . Thus, we have r LIRs, each of which has size n . We can expand each of them with $(m - 1)n$ zeros, such that the expanded

row becomes one row of A^* . Specifically, for the i -th ($1 \leq i \leq m$) M in A^* , we can add $n(i-1)$ zeros before each LIR, and $n(m-i)$ zeros after it. Thus, we obtain rm expanded rows in total. We denote them by R_1, R_2, \dots, R_{rm} , each of which is an expanded row of size mn , such that $R_{r(i-1)+1}, R_{r(i-1)+2}, \dots, R_{r(i-1)+r}$ are expanded rows for the i -th M ($1 \leq i \leq m$).

Now we select the first $r(m-1)$ expanded rows, i.e., $R_1, R_2, \dots, R_{r(m-1)}$. We also select the last n rows in A^* , which are denoted by I_1, I_2, \dots, I_n . We show that this collection of $r(m-1) + n$ rows are linearly independent. Specifically, we consider the linear combination $L = \alpha_1 R_1 + \alpha_2 R_2 + \dots + \alpha_{r(m-1)} R_{r(m-1)} + \beta_1 I_1 + \beta_2 I_2 + \dots + \beta_n I_n$, which is a vector of length mn . We prove that L is a vector with all zero elements only if all coefficients α_i ($1 \leq i \leq r(m-1)$) and β_j ($1 \leq j \leq n$) are also zeros.

Our proof partitions L into m groups with n elements each. We first analyze the last group. If L is a zero vector, the j -th element ($1 \leq j \leq n$) in the last group is contributed only by I_j . Thus, β_j must be zero. Then for the i -th group ($1 \leq i \leq m-1$), only r rows $R_{r(i-1)+1}, R_{r(i-1)+2}, \dots, R_{r(i-1)+r}$ contribute non-zero values for the n elements in this group. Since $R_{r(i-1)+1}, R_{r(i-1)+2}, \dots, R_{r(i-1)+r}$ are linearly independent, we can deduce that $\alpha_{r(i-1)+1}, \alpha_{r(i-1)+2}, \dots, \alpha_{r(i-1)+r}$ must be all zeros to make the linear combination a zero vector. Thus, we prove that all coefficients are zeros. The lemma follows. \square

Next, we show that any $r(m-1) + n + 1$ rows in A^* must be linearly dependent.

LEMMA 7. *Any $r(m-1) + n + 1$ rows in A^* must be linearly dependent.*

PROOF. Consider an arbitrary collection that is composed of $r(m-1) + n + 1$ rows in A . If any M contributes more than r rows, these rows must be linearly dependent because the rank of M is r .

Thus, we focus on the case where each M contributes at most r LIRs. Without the loss of generality, suppose that the first $m-1$ matrices of M contribute r LIRs each, denoted by $R_1, R_2, \dots, R_{r(m-1)}$ as the above proof. The $n+1$ rows come from the last M and I_1, I_2, \dots, I_n . Obviously, at least one row is from the last M . For each row (denoted by P) in the last M , we can construct a linear combination $\alpha_1 R_1 + \alpha_2 R_2 + \dots + \alpha_{r(m-1)} R_{r(m-1)}$, such that $\alpha_1 R_1 + \alpha_2 R_2 + \dots + \alpha_{r(m-1)} R_{r(m-1)} + P$ is a vector of size mn and every group of n elements is the same as each other. Note that each I_i ($1 \leq i \leq n$) also has m groups of n elements and all groups are the same. Such $n+1$ rows must be linearly dependent since there are at most n linearly independent columns. Since the $n+1$ rows are either constructed by the linear combination of $\alpha_1 R_1 + \alpha_2 R_2 + \dots + \alpha_{r(m-1)} R_{r(m-1)} + P$ or directly from some I_i ($1 \leq i \leq n$), the $r(m-1) + n + 1$ rows are linear dependent. \square

The two lemmas lead to the following theorem.

THEOREM 1. *If M has a rank r , the rank of A has a tight upper bound $r(m-1) + n$, where n is the number of lossy flows and m is the number of switches traversed by any lossy flow.*

Number of mapped slots. We analyze the number of slots k mapped by n lossy flows, given K slots in total and d distinct slots

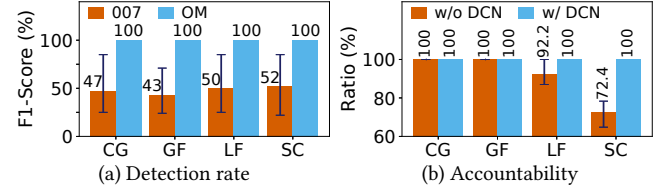


Figure 13: (Exp#8) Packet loss diagnosis.

per flow. For each slot, the probability that it is mapped by a particular flow is

$$p_1 = \frac{\binom{K-1}{d-1}}{\binom{K}{d}} = \frac{d}{K}.$$

With n lossy flows, the probability that it is touched by at least one flow is

$$p_n = 1 - (1 - p_1)^n \approx np_1 = \frac{nd}{K} \text{ with sufficiently large } K.$$

Since the K slots are independent, the number of touched slots follows a Binomial distribution with expectation $K \cdot p_n = nd$ and variation $K \cdot p_n(1 - p_n) = \frac{(nd)^2}{K}$.

Appendix C: Additional Experiments

Simulator. We build a simulator to evaluate OmniMon in large-scale deployment. The simulator forms an 8-ary Fat-Tree topology with 128 end-hosts and 80 switches. It uses a single thread to realize the packet forwarding and OmniMon operations of all end-hosts, switches, and the controller across epochs. It can also inject packet losses to mimic lossy environments. We use the same CAIDA trace [11] to drive our simulations as in tested experiments. In particular, we map the IP addresses in the trace to our end-hosts, such that the ratios of pod-local and aggregate-local flows are 15% and 70%, respectively. We also configure 8,192 slots in each switch (port) in per-switch (per-port) deployment.

(Exp#8) Packet loss diagnosis. We now use the simulator to show that OmniMon can detect four types of packet losses: (i) switch congestions (SG), (ii) gray failures (GF), (iii) switch crash (SC), and (iv) link failures (LF). We configure each type of packet losses as follows. We first randomly select eight switches (resp. links) as problematic switches (resp. links), implying that the failure probabilities for switches and links are 10% and 3%, respectively, in our 8-ray fat tree. In a switch congestion, we randomly drop packets at a rate of 1%; in a gray failure, we randomly sample 1% of flows to drop; for switch crashes and link failures, we drop all traffic in a faulty switch/link after a failure occurs. We assume that the failures occur in the forwarding module, so that flow values in each faulty switch remain available. We repeat the experiment 100 times for each type of packet losses. Note that our scenarios are even more severe than those reported in recent studies [4, 23, 60, 78].

We compare OmniMon with 007 [4], a diagnosis system that infers packet losses on a per-path basis. Since the results of 007 significantly vary across the 100 runs, we plot the error bars showing its maximum and minimum. Figure 13(a) compares their F1 scores (i.e., the harmonic mean of precision and recall). OmniMon achieves perfect detection for all failures, while the accuracy of 007 is quite low. The reason is that 007 builds on a flow voting mechanism for finding faulty switches or links. When the failure probability is

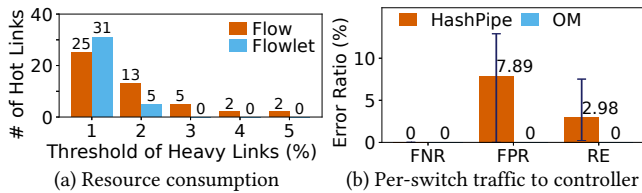


Figure 14: (Exp#9) Load balancing evaluation.

high, the presence of a large number of lossy flows can lead to many noisy votes. Figure 13(b) further shows how OmniMon achieves accountability by measuring the percentage of loss variables that can be solved based on the steps in §5.1. We consider two schemes: (i) skipping Step 3 and directly solving the original linear system (“w/o DCN”); (ii) taking into account the traffic characteristics of data center networks (DCNs) (i.e., executing Step 3) and solving the simplified linear system (“w/DCN”). Without the DCN characteristics, we can infer all loss variables except for switch failures (72%), since a switch failure produces much more lossy flows than the other three types of packet losses and makes the accountability condition fail. With the DCN characteristics, all loss variables can be solved, thereby achieving accountability.

(Exp#9) Load balancing evaluation. We study the performance evaluation effectiveness of OmniMon on the load balancing deployment in a data center network. We evaluate two load balancing algorithms: flow-based ECMP and flowlet-based CONGA [3]. For each algorithm, we feed our traces into the 8-ary Fat-Tree, and simulate the per-packet forwarding operations in each switch. OmniMon collects the actual per-flow bytes in each switch to construct a full view of the entire network.

Figure 14(a) counts the number of heavy links in both algorithms. Here, a link is considered to be heavy if the ratio of its traffic (in bytes) to the total traffic exceeds some threshold. We vary the threshold from 1% to 5% of the total traffic. The results show that CONGA introduces more heavy links than ECMP when the threshold is 1%. However, CONGA has fewer heavy links for larger thresholds (i.e., better load balancing). Figure 14(b) further identifies per-link heavy hitters that account for 10% of the link utilization. We measure the false negative rate (FNR), false positive rate (FPR), and per-flow relative error (RE) of OmniMon and HashPipe [65], a switch-based heavy-hitter detector. We plot the error bars showing its maximum and minimum. HashPipe incurs 8% of false positives and 3% of relative error since it only approximately records the top- k flows. In contrast, OmniMon achieves zero errors. This provides guidelines for administrators to fine-tune the deployment of a new load balancing algorithm.

Breakdown of resource consumptions. Tables 1 and 3 list the complete breakdown of switch overhead and the amount of traffic transferred to the controller, respectively, of each solution in Exp#2 and Exp#6. For Sonata, Marple, and OmniMon, the tables show the per-application resource consumptions in detail. We provide not only the results of the per-switch setup as in §8, but also those of the per-port setup.

Appendix D: Complete Entity Operations

Algorithms 2, 3, and 4 present the pseudo-code of the operations of end-hosts, switches, and the controller, respectively.

End-hosts. Algorithm 2 shows the operations of an end-host. Each end-host maintains: (i) a hash table (line 2), (ii) an array of egress slots (line 3), (iii) an array of ingress slots (line 4), (iv) a local epoch (line 5), (v) a set of switch indexes, which is pre-assigned by the controller and allocated to flows (line 6), and (vi) an array of counters for the usage of switch indexes (line 7).

When sending a packet (lines 8-19), an end-host first queries the flowkey in the hash table (line 9) to obtain both the host index and the switch index (line 10). Otherwise, if the hash table does not include the flowkey, the end-host chooses a new host index and a new switch index (line 12), and inserts the flowkey to the hash table (line 13) (the index selection (lines 37-42) has been elaborated in §3.2 and §5.2). It also increments the usage count of the switch index (line 14). With the host index, flow values in the egress slot will be updated (line 15). The end-host then embeds its current epoch, host index, and switch index into the packet (line 16) and emits the packet (line 17). If the flow terminates, it decrements the usage count (lines 18-19).

On the receipt of a packet (lines 20-27), the destination end-host extracts the source end-host (line 21), host index (line 22), and embedded epoch (line 23). It then updates the corresponding ingress slot (line 24). It updates its local epoch based on embedded packets (lines 25-27), local clock (lines 28-32), and controller messages (lines 33-36).

Whenever the local epoch is updated, an end-host not only performs synchronization, but also reports local results to the controller for collective analysis. It clears the hash table (i.e., removes terminated flows) and resets slot counters for the new epoch (lines 26, 30, and 35).

Switches. Algorithm 3 shows the switch operations. For each packet, a switch retrieves the switch index and epoch embedded in the packet (lines 2-3), as well as the set of slot indexes based on the switch index (line 4). It then updates the corresponding slots with the flow value in the packet (lines 5-6).

Controller. Algorithm 4 shows the controller operations. In collective analysis (lines 1-14), the controller first collects results from all end-hosts and switches (lines 2-3). For each flow, the controller first examines its flow values in the source and destination end-hosts via the host index (line 7). If the values in the two end-hosts are equal, the controller decomposes the values from each switch slot along the packet transmission path with the flow values in end-hosts (lines 8-10). Otherwise, if the flow values differ in the source and destination end-hosts (line 13), this flow is a lossy one (line 14). The controller performs exact loss inference for all lossy flows (lines 13-14). The controller also performs network-wide epoch synchronization (lines 15-19) (§4) and collaborative flow mapping (lines 20-21) (§5). For collaborative flow mapping, the controller hashes each switch index (indicating one possible mapping to slots), and returns a set of switch indexes to the end-host, such that the hash value of each switch index equals the end-host ID.

Appendix E: Generality of OmniMon

We provide more detailed elaboration on the generality of OmniMon.

Flowkeys. OmniMon supports various combinations of packet fields as flowkey definitions. Specifically, end-hosts can track the flowkeys that include all packet fields of interest (e.g., 5-tuple).

Solutions	Tofino resources																			
	Memory (MB)					Stage					Action					PHV (Byte)				
	Queries (Marple Sonata-FA Sonata-AD Omnimon-QE-Switch Omnimon-QE-Port)																			
TCP new connection	1.21	1.26	0.36	0.17	0.41	9	8	5	2	2	18	21	13	7	7	194	189	175	120	120
SSH brute force	2.92	2.47	1.85	0.17	0.69	12	11	7	2	2	24	30	25	7	7	244	201	218	116	116
SuperSpreader	2.33	2.41	0.67	0.02	0.02	10	11	5	1	1	21	26	14	2	2	235	198	162	112	112
Port scan	2.05	2.41	0.63	0.02	0.02	10	11	11	1	1	22	26	17	2	2	229	169	172	112	112
DDoS	2.33	2.41	1.23	0.02	0.02	10	11	8	1	1	20	26	17	2	2	233	201	205	112	112
SYN flood	3.49	2.51	1.39	0.45	1.94	9	12	6	3	3	31	35	37	15	15	272	223	258	120	120
Completed flows	2.36	2.49	0.84	0.31	0.76	10	12	11	2	2	26	34	45	11	11	235	223	255	120	120
Slowloris attack	2.62	2.49	1.97	0.31	2.43	11	11	7	2	4	24	34	37	11	11	277	232	247	120	120
DNS tunneling	1.27	2.6	0.8	0.17	1.25	6	12	5	2	2	19	34	22	7	7	201	231	205	120	120
Zorro attack	2.92	2.49	1.04	0.17	0.69	12	11	6	2	2	24	33	24	7	7	244	220	220	130	130
Reflection DNS	0.11	0.1	0.1	0.17	1.25	3	3	3	2	2	11	11	11	7	7	164	164	157	120	120
Total	23.61	23.64	10.88	1.98	9.48	102	35	74	20	22	240	133	262	78	78	2528	1981	2274	1302	1302
Sketch Solutions																				
FlowRadar	6.58					11					26					133				
UnivMon	6.4					12					67					128				
ElasticSketch	0.59					10					23					163				
SketchLearn	1.16					10					40					158				
OmniMon Solutions																				
Only packet count (per switch)	0.24					2					7					120				
Only packet count (per port)	1.25					2					7					120				
Full statistics (per switch)	1.49					9					39					130				
Full statistics (per port)	6.25					10					39					130				

Table 1: Complete switch resource overheads of solutions in Exp#2 and Exp#6.

Solutions	Traffic to Controller									
	Message (Msg/s)					Byte (KB/s)				
	Queries (Marple Sonata-FA Sonata-AD Omnimon-QE-Switch Omnimon-QE-Port)									
TCP new connection	171	17	16	10	10	10250.32	1.21	1.19	10	320
SSH brute force	10	0	0	10	10	25600.02	0	0	20	640
SuperSpreader	1555	3649	57	0	0	20578.89	3590.45	4.82	0	0
Port scan	816	1042	25	0	0	17971.58	1049.64	1.75	0	0
DDoS	1576	3528	94	0	0	20580.25	3843.88	150.45	0	0
SYN flood	536	23366	23366	10	10	30753.67	26172.16	26452.55	60	1920
Completed flows	72	1053	1053	10	10	20483.97	95.33	107.98	20	640
Slowloris attack	29153	29143	29143	10	10	1865.18	29025.58	29491.87	80	2560
DNS tunneling	315	474	474	10	10	10259.57	41.47	43.36	40	1280
Zorro attack	10	1	1	10	10	25600.05	0.03	0.08	20	640
Reflection DNS	315	305	305	10	10	19.57	23.26	23.26	40	1280
Total	34529	62398	54533	80	80	183963.07	63843.01	56277.31	290	9280
Sketch Solutions										
FlowRadar	10					16000				
UnivMon	10					42240				
ElasticSketch	10					3840				
SketchLearn	10					1320				
OmniMon Solutions										
Only packet count (per switch)	10					160				
Only packet count (per port)	10					1280				
Full statistics (per switch)	10					800				
Full statistics (per port)	10					6400				

Table 2: Complete triggered traffic of solutions in Exp#2 and Exp#6.

Queries	Recall (%)	Precision (%)	F1-Score (%)
TCP new connection	100	100	100
SSH brute force	100	100	100
SuperSpreader	45	75	56
Port scan	20	90	32
DDoS	54	92	68
SYN flood	99	40	57
Completed flows	100	68	81
Slowloris attack	33	82	47
DNS tunneling	100	100	100
Zorro attack	-	-	-
Reflection DNS	100	100	100
Mean	75	85	74

Table 3: Accuracy of Sonata-AD Queries in Exp#6.

During collective analysis, the controller can aggregate the flows for specific combinations of packet fields (e.g., source IP address). For example, end-hosts can track 5-tuple flowkeys, while the con-

troller can group the flows by the source IP address for per-host measurement.

Flow statistics. OmniMon supports different types of flow statistics, include the 344 statistics in [35, 40, 46, 71]. In particular, a flow statistic typically performs a specific aggregated operator over all packets of a flow. Currently, OmniMon supports various basic online analytical processing (OLAP) operators (e.g., COUNT, SUM, MAX, MIN), which are enabled by end-hosts and programmable switches. More complex operators (e.g., RANGE and AVG) can be computed by the basic operators offline.

However, there are two concerns that may potentially limit the support of general flow statistics. First, OmniMon cannot perform collective analysis on packet timestamps. Currently, it addresses this limitation by measuring timestamps at the granularity

Algorithm 2 End-host operations

```

1: function INITIALIZATION
2:   Hash table  $H = \emptyset$  ▷ §3
3:   Array of egress slots  $ES[h]$  for end-host  $h$  ▷ §3
4:   Array of ingress slots  $IS[h][e]$  for end-host  $h$  and epoch  $e$  ▷ §3
5:   Local epoch  $e_h = 0$  ▷ §4
6:   Set of switch indexes  $I$ , assigned by controller ▷ §5
7:   Array of switch index usage counters  $U[s]$  for switch index  $s$ 
   ▷ Egress packet processing, §3

8: function ON_SENDING_PACKET(packet  $p$  with flowkey  $k$ )
9:   if  $k \in H$  then
10:    (switchIndex, hostIndex) =  $H[k]$ 
11:   else
12:    (switchIndex, hostIndex) = GETINDEXES( $k, H, I$ )
13:    Insert  $H[k] = (\text{switchIndex}, \text{hostIndex})$ 
14:    Increment  $U[\text{switchIndex}]$  by 1
15:   Update egress slot  $ES[\text{hostIndex}]$  with  $p$ 
16:   Embed  $e_h$ , switchIndex and hostIndex in  $p$ 
17:   Emit  $p$ 
18:   if Flow  $k$  terminates then
19:     Decrement  $U[\text{switchIndex}]$  by 1
   ▷ Ingress packet processing, §3

20: function ON_RECEIVING_PACKET(packet  $p$ )
21:   Extract source end-host  $h$  based on source_ip
22:   Extract epoch  $e$  from  $p$ 
23:   Extract hostIndex from  $p$ 
24:   Update ingress slot  $IS[\text{hostIndex}][e]$  with  $p$ 
25:   if  $e > e_h$  then ▷ Epoch synchronization, §4
26:     Send  $H, ES,$  and  $IS$  to controller, then reset
27:      $e_h = e$ 
   ▷ Epoch synchronization, §4

28: function ON_NEW_EPOCH_BY_LOCAL_CLOCK(new epoch  $e$ )
29:   if  $e > e_h$  then
30:     Send  $H, ES,$  and  $IS$  to controller, then reset
31:      $e_h = e$ 
32:     Notify  $e$  to controller
   ▷ Epoch synchronization, §4

33: function ON_NEW_EPOCH_FROM_CONTROLLER(new epoch  $e$ )
34:   if  $e > e_h$  then
35:     Send  $H, ES,$  and  $IS$  to controller, then reset
36:      $e_h = e$ 
   ▷ Flow mapping, §3 and §5

37: function GETINDEXES(flowkey  $k$ , hash table  $H$ , switch index set  $I$ )
38:   hostIndex =  $H.\text{getPosition}(k)$ 
39:   if available switch index  $i \in I$  then
40:     switchIndex =  $i$ 
41:   else
42:     switchIndex =  $\arg \min_{i \in I} U[i]$ 
43:   return (switchIndex, hostIndex)

```

Algorithm 3 Switch operations

```

Data structures: Slots  $S[e][i]$  for  $i$ -th slot at epoch  $e$ 
1: function ON_RECEIVING_PACKET(packet  $p$ )
2:   Retrieve switch index switchIndex from  $p$ 
3:   Retrieve epoch  $e$ 
4:   Retrieve the set of slot indexes from switchIndex
5:   for all  $i \in$  set of slot indexes do
6:     Update  $S[e][i]$  with the flow value in  $p$ 

```

of epochs, as elaborated in §6.

The second limitation comes from the accountability guarantee, in which the linear system requires additive flow statistics. OmniMon addresses this by converting non-additive statistics into additive ones, followed by applying the linear system to the converted additive statistics to infer missing values due to packet loss. Specifically, OmniMon still sums up the non-additive values in switches. To compute the correct sum in switches, it leverages the correct flow statistics in end-hosts: whenever the statistics are updated, a source end-host embeds the (zero-error) difference between

Algorithm 4 Controller operations

```

Data structures: Epoch  $e_c$ , Set of all switch indexes  $I$  ▷ §3

1: function COLLECTIVE_ANALYSIS(epoch  $e$ ) ▷ §3
2:   Collect hash table  $H$ , egress slots  $IS$  and ingress slots  $ES$  from all end-hosts
3:   Collect switch slots  $S$  from all switches
4:   Set of lossy flows  $L_f = \emptyset$ 
5:   for all flow  $f$  in every  $H$  do
6:     Obtain (hostIndex, switchIndex) of  $f$  from  $H$ 
7:     if  $IS[\text{hostIndex}] == ES[\text{hostIndex}]$  then
8:       Obtain true flow values from  $ES[\text{switchIndex}]$ 
9:       for all switch  $s$  along  $f$ 's path do
10:        Subtract flow values of  $f$  from  $S[\text{switchIndex}]$  of  $s$ 
11:     else
12:        $L_f = L_f \cup \{f\}$ 
13:   if  $L_f \neq \emptyset$  then
14:     Perform packet loss inference (§5) ▷ §4

15: function ON_NEW_EPOCH(epoch  $e$ )
16:   if  $e > e_c$  then
17:      $e_c = e$ 
18:     for all host  $h$  do
19:       Send  $e_c$  to  $h$  ▷ §5

20: function FLOW_MAPPING(end-host  $h$ )
21:   return  $\{x \in I \mid \text{hash}(x) = h.\text{id}\}$ 

```

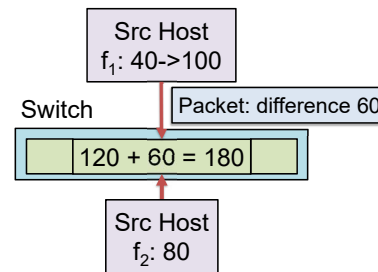


Figure 15: Example of non-additive statistics.

the old and new values of the statistics in the packet header. When a switch receives the packet, it extracts the difference and adds the difference to the corresponding slots. This extension incurs no additional overhead in switches, since a switch performs the same operations as for additive values. However, it requires extra fields in packets. Given that the updates of non-additive statistics are relatively less frequent (e.g., MAX and MIN), we expect that the extra overhead is acceptable.

Example. We illustrate how OmniMon addresses non-additive values via an example, in which it measures the maximum packet size of each flow. As shown in Figure 15, we assume that there are two flows f_1 and f_2 share a slot in the switch. At the beginning, the maximum packet size of f_1 and f_2 are 40 and 80, respectively. The slot hence records their sum 120. Now f_1 has a new packet with size 100. The source end-host of f_1 first updates the maximum packet size of f_1 , and then embeds the difference $100 - 40 = 60$ in the packet header. The switch adds up the difference so that the sum of maximum packet size becomes 180.