Advancing Network Monitoring with Packet-Level Records and Selective Flow Aggregation

Ina Berenice Fink*, Ike Kunze*, Pascal Hein*, Jan Pennekamp*, Benjamin Standaert§, Klaus Wehrle*, and Jan Rüth*
*Communication and Distributed Systems, RWTH Aachen University, Germany
§Washington University in St. Louis, Missouri, United States
{fink, kunze, hein, pennekamp, wehrle, rueth}@comsys.rwth-aachen.de · b.g.standaert@wustl.edu

***Abstract*—Due to its superior efficiency, network operators frequently prefer flow monitoring over full packet captures. However, packet-level information is crucial for the timely and reliable detection, investigation, and mitigation of security incidents. Currently, no solution effectively balances these two contradicting approaches, forcing network operators to compromise between efficiency and accuracy. In this paper, we thus propose HybridMon, a hybrid solution that combines condensed packet-level monitoring with selective flow-based aggregation to strike a new balance between efficiency and accuracy. Operating on the data plane of P4-programmable switches, HybridMon enables fine-grained, practical, and flexible network monitoring at Tbps speeds. We validate the effectiveness of HybridMon through extensive evaluations using Internet backbone and university campus traffic traces, demonstrating its reliability and performance in network forensics and intrusion detection contexts. Our results show that HybridMon reliably monitors all flows while reducing the output bandwidth to 12 % to 20 % compared to packet monitoring when exporting standard features.**

***Index Terms*—Security Services, Control and Data Plane Programmability, Monitoring and Measurements**

I. Introduction

Network monitoring serves as a vital data source for a myriad of applications [1], including *network forensics* [2], [3] and *intrusion detection* [4]. However, their operation is challenged by growing network sizes, increasing the vulnerability surface [5], and performance demands of monitoring solutions [6]. Simultaneously, cyberattacks become more sophisticated with deployments scaling to state-level threats—visible in the Russian invasion of Ukraine [7]—and tailored to individual targets as became evident during the COVID-19 pandemic [8]. Therefore, today's network monitoring needs to fulfill two key requirements for effective network security: (i) the ability to handle high traffic volumes and (ii) the timely provision of comprehensive information for detailed analysis.

Two major monitoring approaches exist: packet and flow monitoring. Packet monitoring captures full packets, ensuring maximum accuracy; however, today's traffic volumes render it largely impractical due to exceptional storage and processing demands [9]. Flow monitoring poses an efficient alternative by storing flow statistics in structured records [10]. While it reduces the load on network and subsequent storage and processing components, it nowadays finds increasing use for intrusion detection [11]–[14]. However, relying on flow monitoring for network forensics and intrusion detection limits capabilities and effectiveness due to its significantly lower accuracy [9], [15]–[18], and delayed information availability [14]. Thus, network operators require new reliable approaches that better balance the trade-off between accuracy and performance. To this end, programmable networking devices provide new opportunities for implementing efficient and flexible monitoring [1]. Nonetheless, a broadly applicable solution which balances accuracy and efficiency is still missing.

To close this gap, we propose **HybridMon**, a hybrid approach between packet and flow monitoring that leverages programmable P4-switches. Instead of full packets or aggregated statistics, HybridMon exports a condensed and structured record with *selected packet-level information* for every packet. This approach minimizes overhead while preserving packet-level accuracy, allowing for timely and detailed analysis. Additionally, HybridMon incorporates multiple mechanisms to systematically reduce output further and alleviate the load on subsequent components without broadly decreasing accuracy. To this end, *selective aggregation* enables flow monitoring for less-relevant (user-defined) shares of traffic, while *fine-granular filtering* allows irrelevant traffic to be excluded upfront. Intrusion Detection Systems (IDSs) can also automatically feed back temporary filter rules, e.g., to exclude malicious high-volume traffic from monitoring once detected. Lastly, HybridMon offers high *deployability* as it runs on commercial-off-the-shelf hardware and can either replace existing switches in-line or be deployed off-path. It handles up to 3.2 Tbps of input per device and provides standardized output compatible with common analysis tools (e.g., nfdump [19]).

Contributions. Our main contributions in this paper are:

- We identify the limitations of traditional network monitoring, and derive concrete requirements to fill the gap.
- Addressing these requirements, our P4-based hybrid approach **HybridMon** efficiently exports customizable packet and flow information in the common IPFIX format.
- Our extensive evaluation of HybridMon demonstrates that operators can benefit from (i) accurate and custom output for security-related tasks, (ii) Tbps throughput and reduced output compared to full packet capture, and (iii) reliable operation, even with heterogeneous traffic patterns.

Open Science Statement. We open-source our implementation [20] under the GPLv3 license.

II. Traditional Network Monitoring and the Road Ahead

The operation of network security applications directly depends on the quality and quantity of their input [21], provided by monitoring systems. In this context, we first examine the trade-offs between packet- and flow-based monitoring. We then derive requirements for the design of new solutions and shortly examine related work.

A. Comparison of Packet and Flow Monitoring

We assess the content, use for attack detection, and performance of today's prevalent approaches in the following.

Content. While packet monitoring involves full packet capture, flow monitoring exports aggregated statistics of multiple packets from the same flow in the form of *flow records*, which are gathered by *flow collectors* [10]. As a result, flow records lead to the loss of individual packet characteristics. Furthermore, they typically include only information up to the transport layer. The structure of flow records and their detailed transmission are defined by flow export protocols such as NetFlow [22] or IPFIX [23] which serve as input for a broad range of general-purpose analysis tools [19], [24], [25].

Attack Detection. Packet monitoring offers maximum information gain and allows for arbitrary processing, including deep packet inspection (DPI), which facilitates flexible analysis of all (unencrypted) content up to the application layer [26]. Thus, packet monitoring provides the best start for attack detection. In turn, flow monitoring is well applicable for detecting numerically striking, e.g., volumetric, attacks, but is generally less effective at identifying more subtle threats, such as slow or semantic attacks [9], [15], including low-rate DoS attacks [27]. Recent research further indicates that the performance of ML-based intrusion detection can suffer from the coarse granularity of flow-based statistics [16], [17] while packet-level information, such as packet sizes or inter-arrival times, can significantly enhance the accuracy of novelty detection [18]. Lastly, aggregating packets into flows delays the forwarding of monitored data, potentially resulting in delays of several minutes in attack detection [14].

Performance. Depending on the network, packet monitoring is not always feasible [9] due to its storage and processing requirements. Additionally, resources may be wasted if packets lack usable payloads due to encryption. Flow monitoring substantially reduces the load on subsequent network, storage, and processing components compared to packet monitoring [10], but it requires additional processing power for flow metering and export. Furthermore, flow-based monitoring solutions struggle with short-lived flows [9] and high volumes of minimum-size packets [28], risking further information loss.

Given the shortcomings of packet and flow monitoring, network operators would significantly benefit from hybrid approaches that address their needs.

B. Toward Hybrid Network Monitoring

Based on the weaknesses of packet and flow monitoring, we derive precise requirements for hybrid solutions:

- **R1 Deployment:** Typically, network operators already have elaborate network infrastructures and analysis pipelines in use in which new monitoring systems should easily integrate. Providing output in a standardized and structured format, e.g., via IPFIX flow records, facilitates universal use and seamless processing at subsequent components.
- **R2 Output:** Accurate packet-level information is critical for attack detection and information above the transport layer can offer additional value [11], [29]. Therefore, solutions should encompass packet-level header fields and, if unencrypted, significant application layer information such as request types or error codes. Also, timely provision of the monitored information to security applications is vital for fast attack detection and reaction [14].
- **R3 Flexibility:** Reliable operation of the monitoring system and dependent security applications necessitates support of high traffic volumes and abnormal traffic patterns as well as optimized selection of monitored traffic to decrease load. Adjustment of the monitored traffic at runtime allows to proactively relieve all components from benign [30] and adverse traffic after detection [14].

Next, we discuss related work in light of **R1**-**R3**.

C. Related Work

Existing work primarily focuses on performant network telemetry [31]–[37] and traffic monitoring [6], [38]–[40]. Only a subset specifically addresses network security, focussing on enhanced flow export [41]–[43], or specific (integrated) security applications [14], [15], [44]–[47]. While we conclude from related work that custom monitoring of Tbps throughputs is possible with programmable network devices, we identify a lack of flexible approaches that provide standardized (**R1**) and fine-grained (**R2**) output, and can handle today's traffic volumes (**R3**). To bridge this gap, we propose HybridMon.

III. Packet-Level Monitoring with Selective Flow-Based Aggregation

This section introduces **HybridMon**, a performant hybrid solution for network monitoring. While hybrid monitoring could be achieved using existing software-based flow exporters like YAF [48], HybridMon leverages P4-programmable switches to achieve flexible monitoring at Tbps speeds. HybridMon specifically targets the Intel Tofino [49] switch to optimize compatibility and functionality with today's state-of-the-art hardware. Still, the design principles of HybridMon are not confined to Intel Tofino but adaptable to other P4 targets.

We cover the requirements from Sec. II-B as follows:

Deployment. Fulfilling **R1**, our monitoring relies on established flow export protocols, ensuring compatibility with common analysis tools. Running on a single switch, HybridMon is readily deployed in-line or off-path.

Output. In line with **R2**, HybridMon instantly exports flow records with a flow size of 1, i.e., one record for every packet, to deliver packet-level information while covering a comprehensive and customizable range of features. Thus, HybridMon facilitates rich, time-critical analysis.

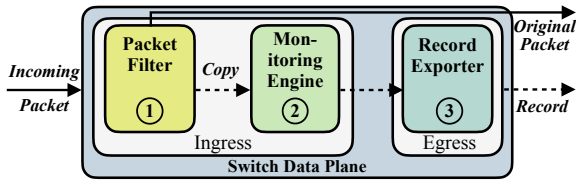

Fig. 1. HybridMon's three components: ① The packet filter extracts features, sorts out irrelevant traffic, and copies relevant packets. These copies are further processed by ② the monitoring engine before ③ the record exporter outputs flow records, either containing packet-level information or flow statistics.

Flexibility. Addressing **R3**, HybridMon includes a traffic filter and allows for the aggregation of selected (low-interest) traffic into flow statistics to reduce output. It further operates entirely on the switch's data plane, avoiding CPU-based bottlenecks and enabling reliable handling of diverse traffic patterns.

Design Overview. HybridMon comprises three (data-plane) components illustrated in Fig. 1. First, the ① Packet Filter identifies relevant traffic and copies the respective packets to the ② Monitoring Engine, which determines whether to subsample the corresponding flow. Then, the copied packets are passed to the ③ Record Exporter, which probabilistically generates a flow record for packets belonging to subsampled flows and a packet-level record for every other packet. The original packets remain unmodified and, if HybridMon is deployed in-line, are forwarded to their destination.

Overall, HybridMon enables detailed and customizable in-line and off-path network monitoring in high-speed networks, addressing the needs identified in Sec. II-B. In the remainder of this paper, we demonstrate its feasibility by presenting and evaluating an implementation for Tofino1.

IV. Open-Source Implementation of HybridMon

We implemented HybridMon in P4 for the Intel Tofino1 on top of basic Longest Prefix Match (LPM) routing functionality. We chose the IPFIX protocol [23] as it is a widely supported open standard, increasing deployability. In the following, we discuss our implemented subsampling and monitored features.

A. Subsampling Strategy

Heavy hitters are prime candidates for subsampling as packets of these flows typically exhibit similar characteristics, making their aggregation less critical regarding information loss while providing significant reduction potential. Thus, our monitoring engine implementation targets heavy hitters by leveraging PRECISION [50] with a 2-way associative flow table, split over 2 register arrays. We configured a flow table size of 65 536 entries, which is the maximum share of subsampled flows supported by our implementation on the used hardware. We further lowered the recirculation probabilities in our implementation by a factor of 6 compared to PRECISION, increasing the monitoring capacity.

B. Monitored Features

Our implementation currently supports 31 features, and our tests showed that we can easily employ 10 counters of 4 B each for applicable features of subsampled flows. The supported features include 27 standard IPFIX Information Elements (IEs) defined by IANA [51], including source and destination MAC and IP addresses, protocol ID, source and destination ports, number of octets, flow start time, IP time-to-live, fragmentation offset, ID, and flags, ICMP type and code, TCP flags, sequence number, and window size, and HTTP status code. Employing enterprise-specific IEs [51], we further introduced port-based detection of startTLS, and DNS over UDP, including DNS request and response code. These examples prove that HybridMon can easily export custom flow and packet-level features (cf. **R2**) above the transport layer while fully complying with the IPFIX protocol.

Our prototype demonstrates that HybridMon can be implemented on off-the-shelf switch hardware and supports standard flow export protocols, thereby satisfying **R1**. Subsequently, we evaluate our implementation with respect to **R2** and **R3**.

V. Evaluation of HybridMon

We provide an extensive evaluation to prove the feasibility of our design and ensure that the requirements established in Sec. II-A are met. We first investigate the data quality of HybridMon's output in Sec. V-A and how it copes with traffic from Internet backbone links in Sec. V-B, comparing it to the flow exporter YAF [48]. Then, we examine its performance and benefits in a university-specific use case in Sec. V-C.

A. Data Quality Assessment

We first examine the data quality of HybridMon to ensure that it is sound and can be used as input for reliable analysis (**R2**). For this purpose, we generated IPFIX records with and without HybridMon's heavy hitter-based subsampling, i.e., a mix of packet-level and flow records vs. packet-level records only. We compared the output to (i) the input and (ii) the traditional IPFIX-based flow records exported by YAF [48]. To highlight the applicability to production traffic, we used four real-world traffic traces as input, each randomly chosen from the publicly available datasets provided by CAIDA [52]–[55], containing extensive anonymized packet captures with real traffic recorded at high-speed backbone links.

To evaluate HybridMon on the datasets, we connected two workstations *W1* and *W2*, using 10 Gbps links to one 32-port Tofino-based switch running HybridMon. Then, we replayed each trace from W1 to the switch and collected its output, i.e., the generated IPFIX records, at W2. We conducted 10 runs for each of the four CAIDA traces to obtain significant results. Due to processing limitations of our workstations, we replayed the CAIDA traces at 100 Mbps and fed the same slowed-down traces into YAF to obtain comparable output. This adaption does not influence our evaluation results as it only affects the packets' timestamps but not the metadata, such as IP addresses or protocols (i.e., the targets of our evaluation).

1) Packet Coverage: We first evaluated the share of packets reported for every {protocol, direction, port}-tuple. To this end, we counted the packet numbers for each tuple in the original traces and compared them against the records generated by HybridMon and YAF. YAF covered 100 % of the

TABLE I
SHOWN ARE THE OUTPUT QUANTITIES AS A SHARE (IN %) OF THE INPUT QUANTITIES OF THE DIFFERENT APPROACHES AND HYBRIDMON'S SUBSAMPLING RATE, I.E., HOW MANY FLOWS WERE SUBSAMPLED. IPFIX RECORDS ARE RELATIVE TO THE INPUT PACKET COUNT.

Trace	Method	Packets	Bytes	IPFIX Records
CAIDA2011A	HybridMon	100.00	20.22	100.00
	sub. (7.21%)	51.42	10.40	51.42
	YAF	4.89	1.16	12.09
CAIDA2011B	HybridMon	100.00	17.61	100.00
	sub. (11.99%)	41.23	7.26	41.23
	YAF	2.78	0.70	8.52
CAIDA2015A	HybridMon	100.00	16.91	100.00
	sub. (8.15%)	46.25	7.80	46.24
	YAF	1.28	0.42	4.99
CAIDA2018A	HybridMon	100.00	12.07	100.00
	sub. (8.34%)	37.20	4.50	37.19
	YAF	3.57	0.43	7.12

packets of each tuple in its records. The same was the case for HybridMon *without* subsampling. With subsampling, we can only create flow records for incoming packets and not trigger record generation at an arbitrary time. Consequently, replacing a flow may result in losing its accumulated statistics since its last export. To mitigate the losses, we always generate a record for a flow's first and final packet (if, e.g., indicated through TCP flags). Thus, with subsampling, HybridMon reported 92 % of a tuple's packets in the worst case. However, for 99.9 % of the tuples, the packet coverage was still above 95 %, and 83 % had 100 % of their packets reported.

B. Monitoring Efficiency and Throughput

Next, we evaluated the resource efficiency and monitoring capacity of HybridMon (**R3**) using real hardware.

1) Resource Efficiency: We first evaluated HybridMon's output size in terms of packets, bytes, and records, again comparing it against the open-source tool YAF. To this end, we generated records including standard flow features (i.e., TCP flags and packet and byte counters), resulting in records of 50 B for HybridMon and 49 B for YAF. We then conducted 10 runs for each of the four CAIDA traces and averaged the numbers of exported packets, bytes, and records (cf. Tab. I).

YAF applies advanced software-based aggregation to *all* flows, creating low record numbers of up to 12 % of the input packets. YAF can also aggregate multiple records into one record packet, leading to even lower numbers of record packets than records and only 0.4 % to 1 % of the original bandwidth. In contrast, HybridMon's implementation on switch hardware does not enable such record aggregation. Thus, each record requires a packet and the number of output packets always equals the number of output records. Logically, no reduction of the output records and packets occurs with deactivated subsampling, where one packet-level record is generated for every monitored packet. Still, even this 1-to-1 mapping significantly reduces the bandwidth, as the condensed IPFIX records only need 12 % to 20 % of the original bandwidth. In turn, heavy hitter-based subsampling, which affected 7 % to 12 % of the flows, reduced the number of output records and record packets by around 48 % to 63 %, reducing the bandwidth to 4.5 % to 10.5 % of the original trace.

The results of this evaluation show that HybridMon provides lower output efficiency compared to traditional flow monitoring, which is an expected consequence of its higher granularity. However, its bandwidth is significantly reduced compared to packet monitoring, even without aggregation. Furthermore, we did not deploy any filter rules for our evaluation, which will additionally reduce the output size. Last, software-based solutions, e.g., deployed at additional middleboxes, could complement HybridMon by providing subsequent aggregation of its output and further reducing the load on the network.

2) Monitoring Capacity: Our Tofino1 switch has two independent pipelines, each accommodating 16 ports. The monitoring engine in our implementation leverages packets received on a single pipeline. Separate monitoring with both pipelines is feasible, e.g., to monitor different subnets, and theoretically allows for up to 3.2 Tbps for off-path monitoring and up to 1.6 Tbps for in-line monitoring. However, in practice, there are additional factors to consider, particularly the impact of recirculation when using subsampling, the deployment scenario, and different input patterns such as attack traffic. We evaluate the effects of these factors subsequently.

Impact of Recirculation. Our implementation employs heavy hitter-based subsampling using PRECISION [50], which probabilistically recirculates a portion of the input packets to add them to the flow register/heavy hitter list. This recirculation can reduce the amount of traffic HybridMon can handle to less than the switch's maximum capacity. To measure the impact of recirculation, we replayed each of the four CAIDA traces 10 times as done in Sec. V-A and counted the recirculated packets. Recirculation rates for these traces were consistent across runs and ranged from 0.6 to 1.1 %. We also measured the recirculation rate for the measurement traffic described in Sec. V-C, resulting in 2.8 %. These results show that recirculation has only minimal impact on the switch's capacity. With average traffic, Tofino1 with 3.2 Tbps line-rate can support around 3.16 Tbps. Even subject to unusual traffic, where the heavy hitter table is ineffective, recirculation only reduces the switch's maximum capacity to around 3.1 Tbps.

Impact of Deployment. The actual monitoring capacity further depends on the deployment scenario. To forward records directly to collectors, i.e., without loading the rest of the network, we need to occupy dedicated collector ports. Furthermore, in the case of short flows with small packets, e.g., caused by DNS traffic or SYN flooding, record packets might be more than twice the size of the original packet since subsampling cannot be applied and current data plane capabilities do not allow exporting multiple records per packet. To prevent record loss in this case, we need to account for twice as much output as input, dedicating at least twice as many ports to the collector ports as to the monitored traffic. Then, HybridMon provides a monitoring capacity of up to 1.03 Tbps for in-line and off-path monitoring when using both pipelines and considering recirculation. Thus, even in the worst case, HybridMon offers a significantly higher monitoring capacity than commercial hardware appliances for flow monitoring, which typically support up to 100 Gbps per device, e.g., the Flowmon Probe [56].

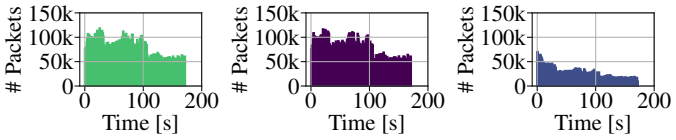

(a) University traffic (input) (b) HybridMon output (full) (c) HybridMon output (subsampled)

Fig. 2. Packet distributions (1 s bins) of input/output traffic originating from our university network's peering point.

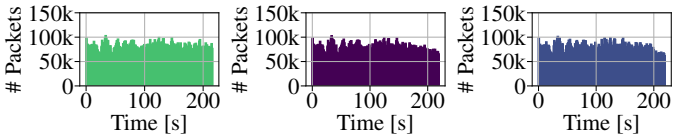

(a) Measurement traffic (input) (b) HybridMon output (full) (c) HybridMon output (subsampled)

Fig. 3. Packet distributions (1 s bins) of input/output traffic originating from our university's data center containing only special measurement traffic.

Impact of Volumetric Attacks. Just as other monitoring approaches, HybridMon is not immune to volumetric attacks. In particular, as described above, short flows with small packet sizes can amplify the monitoring output. However, HybridMon is more robust than common approaches that aggregate records on the control plane as their performance depends heavily on the successful aggregation of packets, which may fail in face of many simultaneous small flows and lead to record loss. In turn, HybridMon provides reliable monitoring as long as the data plane capacity is not exceeded (cf. Sec. V-B2). Additionally, detected attacks can be immediately blocked from monitoring through respective filter rules, e.g., installed by IDSs.

C. Use Case Evaluation: University Network Traces

HybridMon is designed to address the needs of network operators who operate a large backbone with heterogeneous network use regarding throughputs, users, and machines. Therefore, we verify the performance of HybridMon in our campus network (**R3**), and discuss how HybridMon's functional features benefit this use case (**R2**, **R3**).

1) Performance Evaluation: University networks must both cope with high traffic volumes and unusual traffic patterns due to ongoing research, e.g., caused by Internet measurements at our computer science department [57], resulting in outstanding shares of UDP traffic and small flows. Indeed, our network operators report that commercial flow export solutions, as partly deployed at our university's backbone, at times struggle with handling these loads and fail to generate records. Thus, we evaluated the performance of our HybridMon with real traffic traces covering two scenarios: standard and research traffic. Specifically, we were provided with truncated traffic traces captured at the peering point of RWTH Aachen University with its transit network, covering all non-internal traffic.

Regular Operation. The regular traffic capture covers almost 3 min of RWTH Aachen University's full traffic. We visualize its packet count distribution in Fig. 2a). Fig. 2b) details the output of HybridMon without subsampling, showing that the output distribution mirrors the input distribution. In contrast, Fig. 2c) shows the output when applying subsampling, resulting in significantly reduced packet counts.

Research Traffic. We provide the packet count distribution of the examined research traffic in (cf. Fig. 3a)). Using the respective capture as input for HybridMon, we observed that our subsampling strategy had little impact, resulting in similar input and output distributions for full and subsampled flow monitoring (cf. Fig. 3a) and Fig. 3c). However, all flows of the original trace were preserved, again covering over 99.9 % percent of their respective packets on average (cf. Fig. 3b).

Discussion. Similar to the attack traffic discussed in Sec. V-B2, the research traffic makes it hard to define permanent heavy hitters for subsampling due to its many small flows. Still, this challenging traffic is handled without effort or accuracy loss. Furthermore, our evaluation demonstrates that HybridMon effectively reduces the monitoring output compared to packet monitoring for regular traffic, i.e., the large majority of traffic.

2) Functional Benefits: As discussed in Sec. V-B2, HybridMon can handle challenging traffic as occurring in our university's network with up to 1.03 Tbps. According to our network operators, this throughput and the port density of Tofino1 are more than sufficient to cover the whole traffic of RWTH Aachen University (around 40 Gbps) with a single switch for the foreseeable future. Research traffic is also one example of traffic that regularly occupies high shares of monitoring and processing resources, although it is known and not of interest from a security perspective, and can be reasonably excluded through filtering (cf. Sec. III). Last, we added a port-based startTLS detection to HybridMon's feature list upon request of our university's network operators, who are particularly interested in observing such connections due to the numerous vulnerabilities of startTLS [58]. This highlights the capability and benefits of HybridMon for monitoring custom features to increase the visibility of critical characteristics.

Overall, our results emphasize that HybridMon can handle real-world traffic without any particular restrictions and provide use case specific features. We plan to thoroughly study HybridMon in a production environment at the university's uplink and replace commercial solutions if applicable.

VI. Conclusion

Network monitoring is key for detecting and investigating attacks, but we identify a lack of balanced solutions between packet monitoring and flow monitoring. Especially as broadly applied encryption (e.g., through QUIC) might eventually render full packet captures obsolete, we expect network operators to have an increasing demand for efficient monitoring of packet-level metadata. To fill this gap, we propose HybridMon, which combines customized IPFIX-based packet-level records, selective flow aggregation, and traffic filters. In particular, monitoring can be flexibly adapted at run-time, e.g., through IDS. Our open-source implementation is readily deployable inline or off-path and exports records at Tbps while providing compatibility with existing flow-based analysis tools and our evaluation confirms the practical feasibility of HybridMon for reliable and efficient network monitoring in the wild. In the

future, we will combine HybridMon with diverse analysis tools and IDS, to evaluate the impact of its increased monitoring granularity compared to traditional flow monitoring.

References

[1] D. Ding *et al.*, "Design and Development of Network Monitoring Strategies in P4-enabled Programmable Switches," in *NOMS*, IEEE, 2022.

[2] K. Kent *et al.*, "Guide to Integrating Forensic Techniques into Incident Response." NIST SP 800-86, 2006.

[3] R. Hunt and S. Zeadally, "Network Forensics: An Analysis of Techniques, Tools, and Trends," *Computer*, vol. 45, no. 12, 2012.

[4] B. Mukherjee *et al.*, "Network Intrusion Detection," *IEEE Network*, vol. 8, no. 3, 1994.

[5] S. J. Saidi *et al.*, "Exploring Network-Wide Flow Data With Flowyager," *IEEE Transactions on Network and Service Management*, vol. 17, no. 4, 2020.

[6] J. Sonchack *et al.*, "Scaling Hardware Accelerated Network Monitoring to Concurrent and Dynamic Queries With *Flow," in *ATC*, USENIX Association, 2018.

[7] European Parliament, "Russia's war on Ukraine: Timeline of cyber-attacks." https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2022)733549, 2022 (accessed January 5, 2024).

[8] H. S. Lallie *et al.*, "Cyber security in the age of COVID-19: A timeline and analysis of cyber-crime and cyber-attacks during the pandemic," *Computers & Security*, vol. 105, 2021.

[9] A. Sperotto *et al.*, "An Overview of IP Flow-Based Intrusion Detection," *IEEE Communications Surveys & Tutorials*, vol. 12, no. 3, 2010.

[10] R. Hofstede *et al.*, "Flow Monitoring Explained: From Packet Capture to Data Analysis With NetFlow and IPFIX," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 4, 2014.

[11] F. Erlacher and F. Dressler, "FIXIDS: A High-Speed Signature-based Flow Intrusion Detection System," in *NOMS*, IEEE, 2018.

[12] L. Hellemons *et al.*, "SSHCure: A Flow-Based SSH Intrusion Detection System," in *AIMS*, Springer, 2012.

[13] L. Dias *et al.*, "Go With the Flow: Clustering Dynamically-Defined NetFlow Features for Network Intrusion Detection with DynIDS," in *NCA*, IEEE, 2020.

[14] R. Hofstede *et al.*, "Towards Real-Time Intrusion Detection for NetFlow and IPFIX," in *CNSM*, IEEE, 2013.

[15] S. Panda *et al.*, "SmartWatch: Accurate Traffic Analysis and Flow-State Tracking for Intrusion Prevention Using SmartNICs," in *CoNEXT*, ACM, 2021.

[16] H. Clausen *et al.*, "Controlling Network Traffic Microstructures for Machine-Learning Model Probing," in *SecureComm*, Springer, 2021.

[17] M. A. Salahuddin *et al.*, "Time-based Anomaly Detection using Autoencoder," in *CNSM*, IEEE, 2020.

[18] K. Yang *et al.*, "Feature Extraction for Novelty Detection in Network Traffic." arXiv:2006.16993, 2020.

[19] P. Haag, "nfdump." https://github.com/phaag/nfdump, 2015.

[20] I. B. Fink *et al.*, "Prototype Implementation of HybridMon." https://github.com/COMSYS/HybridMon, 2025.

[21] L. F. Sikos, "Packet analysis for network forensics: A comprehensive survey," *Forensic Science International: Digital Investigation*, vol. 32, 2020.

[22] B. Claise, "Cisco Systems NetFlow Services Export Version 9." RFC 3954, 2004.

[23] B. Claise *et al.*, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information." RFC 7011, 2013.

[24] C. Gates *et al.*, "More Netflow Tools for Performance and Security," in *LISA*, USENIX Association, 2004.

[25] nTop, "nProbe." https://www.ntop.org/products/netflow/nprobe/, 2015.

[26] T. AbuHmed *et al.*, "Deep Packet Inspection for Intrusion Detection Systems: A Survey," *Journal of the Korean Institute of Communication Sciences*, vol. 24, no. 11, 2007.

[27] A. Kuzmanovic and E. W. Knightly, "Low-rate TCP-targeted denial of service attacks: the shrew vs. the mice and elephants," in *SIGCOMM*, ACM, 2003.

[28] J. R. Binkley and B. Massey, "Ourmon and Network Monitoring Performance," in *USENIX ATC*, USENIX Association, 2005.

[29] P. Velan and P. Čeleda, "Application-Aware Flow Monitoring," in *IM*, IEEE, 2019.

[30] J. Amann and R. Sommer, "Providing Dynamic Control to Passive Network Security Monitoring," in *RAID*, Springer, 2015.

[31] A. Gupta *et al.*, "Sonata: Query-Driven Streaming Network Telemetry," in *SIGCOMM*, ACM, 2018.

[32] Y. Zhou *et al.*, "Flow Event Telemetry on Programmable Data Plane," in *SIGCOMM*, ACM, 2020.

[33] C. Misa *et al.*, "Dynamic Scheduling of Approximate Telemetry Queries," in *NSDI*, USENIX Association, 2022.

[34] G. Vassoler *et al.*, "VERMONT: Towards an In-band Telemetry-Based Approach for Live Network Property Verification," in *NOMS*, IEEE, 2023.

[35] Z. Xu *et al.*, "Information-Sensitive In-Band Network Telemetry in P4-Based Programmable Data Plane," *IEEE/ACM Transactions on Networking*, 2024.

[36] K. Papadopoulos *et al.*, "Deterministic and Probabilistic P4-Enabled Lightweight In-Band Network Telemetry," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, 2023.

[37] H. N. Nguyen *et al.*, "A comprehensive p4-based monitoring framework for l4s leveraging in-band network telemetry," in *NOMS*, 2023.

[38] Z. Liu *et al.*, "One Sketch to Rule Them All: Rethinking Network Flow Monitoring with UnivMon," in *SIGCOMM*, ACM, 2016.

[39] O. Michel *et al.*, "Software Packet-Level Network Analytics at Cloud Scale," *IEEE Transactions on Network and Service Management*, vol. 18, no. 1, 2021.

[40] R. Doriguzzi-Corin *et al.*, "Introducing packet-level analysis in programmable data planes to advance network intrusion detection," *Computer Networks*, vol. 239, 2024.

[41] B. Guan and S.-H. Shen, "FlowSpy: An Efficient Network Monitoring Framework Using P4 in Software-Defined Networks," in *VTC*, IEEE, 2019.

[42] J. Sonchack *et al.*, "TurboFlow: Information Rich Flow Record Generation on Commodity Switches," in *EuroSys*, ACM, 2018.

[43] F. Erlacher *et al.*, "Improving Network Monitoring through Aggregation of HTTP/1.1 Dialogs in IPFIX," in *LCN*, IEEE, 2016.

[44] N. Gray *et al.*, "High Performance Network Metadata Extraction Using P4 for ML-based Intrusion Detection Systems," in *HPSR*, IEEE, 2021.

[45] D. Barradas *et al.*, "FlowLens: Enabling Efficient Flow Classification for ML-based Network Security Applications," in *NDSS*, Internet Society, 2021.

[46] G. Gori *et al.*, "GRAPH4: A Security Monitoring Architecture Based on Data Plane Anomaly Detection Metrics Calculated over Attack Graphs," *Future Internet*, vol. 15, no. 11, 2023.

[47] P. Velan, "EventFlow: Network flow aggregation based on user actions," in *NOMS*, IEEE, 2016.

[48] C. M. Inacio and B. Trammell, "YAF: Yet Another Flowmeter," in *LISA*, USENIX Association, 2010.

[49] Intel Corporation, "Intel® Tofino™ Programmable Ethernet Switch ASIC." https://www.intel.com/content/www/us/en/products/network-io/programmable-ethernet-switch/tofino-series.html, 2020.

[50] R. Ben-Basat *et al.*, "Efficient Measurement on Programmable Switches Using Probabilistic Recirculation," in *ICNP*, IEEE, 2018.

[51] IANA, "IP Flow Information Export (IPFIX) Entities." https://www.iana.org/assignments/ipfix/ipfix.xhtml, 2007.

[52] CAIDA, "The CAIDA UCSD Anonymized Internet Traces 2011 – DirA 20110607-235600 UTC." https://catalog.caida.org/details/dataset/passive_2011_pcap, 2011.

[53] CAIDA, "The CAIDA UCSD Anonymized Internet Traces 2011 – DirB 20151217-133400 UTC." https://catalog.caida.org/details/dataset/passive_2011_pcap, 2011.

[54] CAIDA, "The CAIDA UCSD Anonymized Internet Traces 2015 – DirA 20151217-133400 UTC." https://catalog.caida.org/details/dataset/passive_2015_pcap, 2015.

[55] CAIDA, "The CAIDA UCSD Anonymized Internet Traces 2018 – DirA 20180816-135200 UTC." https://catalog.caida.org/details/dataset/passive_2018_pcap, 2018.

[56] Progress Software Corporation, "Flowmon Probe." https://www.flowmon.com/en/products/appliances/probe, 2009.

[57] M. Dahlmanns *et al.*, "Easing the conscience with opc ua: An internet-wide study on insecure deployments," in *IMC*, Association for Computing Machinery, 2020.

[58] H. Böck, "Vulnerabilities show why STARTTLS should be avoided if possible." https://blog.apnic.net/2021/11/18/vulnerabilities-show-why-starttls-should-be-avoided-if-possible/, 2021 (accessed October 11, 2024).