Using Explicit (Host-to-Network) Flow Measurements for Network Tomography

Ike Kunze
kunze@comsys.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Constantin Sander
sander@comsys.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Alexander Ruhrmann
ruhrmann@comsys.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Niklas Rieken
rieken@oms.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Klaus Wehrle
wehrle@comsys.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

ABSTRACT

QUIC hides implicit protocol semantics commonly used for flow performance measurements with TCP. Addressing this challenge, Explicit (Host-to-Network) Flow Measurement (EFM) schemes explicitly expose measurable signals: QUIC's spin bit enables latency measurements, and RFC 9506 defines mechanisms for packet loss. While the measurement accuracy of these schemes has already been studied, their utility for network management is unexplored. In this work, we demonstrate that EFM can enable network tomography (NT), which traditionally uses active measurements. Comparing EFM-based to traditional NT through extensive network simulations, we find that raw EFM output challenges NT due to fluctuations induced by links outside the monitored network. However, path segmentation capabilities of the spin and Q bits can remedy these issues. Overall, our findings highlight the general potential of EFM for network management.

CCS CONCEPTS

• **Networks** → **Network monitoring**.

KEYWORDS

Flow Monitoring, Network Tomography

ACM Reference Format:

Ike Kunze, Constantin Sander, Alexander Ruhrmann, Niklas Rieken, and Klaus Wehrle. 2025. Using Explicit (Host-to-Network) Flow Measurements for Network Tomography. In *Applied Networking Research Workshop (ANRW 25), July 22, 2025, Madrid, Spain.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3744200.3744769

ANRW 25, July 22, 2025, Madrid, Spain

ACM ISBN 979-8-4007-2009-3/25/07
https://doi.org/10.1145/3744200.3744769

1 INTRODUCTION

Network operators often struggle to identify performance impairments, learning of issues only from customers or fellow operators [7, 36]. To proactively detect such problems, many works focus on monitoring packet loss and delay, two pervasive indicators of impairments in day-to-day network operations [16]. While TCP's implicit semantics enable passively deducing these metrics from production traffic [3, 41], QUIC hides its corresponding semantics, thus preventing such measurements [28]. To address this challenge, QUIC includes the optional *spin bit* for continuous latency estimations. Similarly, RFC 9506 [16] defines several packet loss estimation concepts, collectively labeling them and the spin bit as Explicit (Host-to-Network) Flow Measurements (EFMs). Prior work has examined the EFM measurement behavior [12, 17, 32] and has shown that the spin bit is used on the web [30]. However, the utility of EFM for subsequent network management tasks has yet to be explored.

In this paper, we demonstrate that EFM can fuel network management using the example of network tomography (NT) [15], which traditionally assesses the network state based on active measurements. Reusing standard NT methods, we comprehensively compare traditional, TCP-based, and EFM-based NT. We conduct our study with an ns-3 simulation [29] to have a reliable ground truth and flexibly evaluate a broad parameter space. While we discover that raw EFM output is ineffective due to effects outside the monitored network, spin and Q bit path segmentation capabilities can mitigate these issues. Consequently, path-segmented EFM-based NT accurately estimates the network conditions experienced by production traffic. These findings inform efforts to incorporate additional EFM schemes into QUIC [19] and illustrate the broader potential of using EFM in network management tasks. Overall, we make three main findings.

- Effects of external networks make raw EFM output unsuitable for direct use in network tomography
- Spin and Q bit path segmentation capabilities can remedy these issues, improving latency and packet loss estimation
- The explicit semantics of EFM outperform corresponding solutions based on TCP's implicit semantics

2 BACKGROUND & MOTIVATION

Network tomography (NT) [15] is a broad field with extensive research over the years [39, 43]. It typically uses carefully orchestrated *active* measurements that provide information on end-to-end paths to assess the network state. This section gives an overview of NT and presents selected EFM schemes in detail. We also define our research objectives.

2.1 Traditional Network Tomography

NT has a large design space and traditionally relies on active end-to-end probes within the monitored network. With these measurements, NT can deduce different metrics for different goals, and NT designs differ in how they model the metrics and derive them from the underlying measurements.

Metrics and goals. NT can, e.g., discover network topologies [43] or identify broken links [4]. Most works target performance metrics, such as latency [9], packet loss [33], or bandwidth [18]. Latency and loss are pervasive problems of network operation [16] that harm goodput and responsiveness, so we focus on these metrics in our work.

Modeling and deriving the metrics. Most NT approaches are *algebraic* [35], i.e., they interpret the metrics as static link properties and apply algebraic solutions to the end-to-end measurements. In particular, they generally combine all measurements and formulate the underlying problem of identifying n link characteristics $x \in \mathbb{R}^n$ based on m measurements $b \in \mathbb{R}^m$ as a system of linear equations $A \cdot x + \epsilon = b$. Each row of $A \in \mathbb{R}^{m \times n}$ defines which links are covered by a measurement, and ϵ represents measurement errors. The assumption is that link characteristics are additive, as are link delays, or that they can be expressed in additive form, e.g., via *log* for packet loss. Using at least n linearly independent measurements ($m \geq n$), the system of equations is solvable.

The need for EFM. NT commonly uses active measurements, which are insufficient for grasping the network state seen by production flows [6, 34, 38]. As one solution, related work has proposed using passive information. Arzani et al. [5], e.g., collect end-host TCP statistics in data centers, but their solution requires end-host control and is not transferable to more diverse networks. Still, using production traffic and studying transport layer semantics serves as our main motivation. In particular, EFM schemes can capture production traffic performance and might adequately fuel NT. Next, we give an overview of the EFM schemes used in this work.

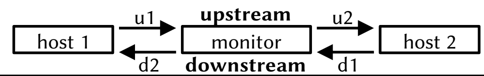

scheme	downstream	bidirectional
spin bit	(d2,u1,u2,d1)	(d2,u1),(u2,d1)
L bit	(d1,d2)	(u1,u2)
Q bit	(d1)	(u1)
R bit	(u1,u2,d1)	(d1,d2,u1)

Figure 1: Measurement paths (MPs) enabled by a unidirectional downstream monitor and additional MPs enabled by bidirectional visibility for the EFM schemes.

2.2 Explicit (Host-to-Network) Flow Measurement (EFM) Schemes

EFM schemes alleviate the loss of measurable information with QUIC and address general flow measurement challenges as the reliance on implicit semantics complicates estimations and risks inaccuracies due to ambiguous signals [1]. Explicitly measurable signals resolve this ambiguity and facilitate passive measurements. RFC 9506 [16] defines several latency and loss measurement schemes, and the spin bit is already an optional QUIC feature [27]. We focus on the spin bit and three schemes that were shown to provide reliable and timely packet loss measurements [32] — the L, Q, and R bits.

We next concisely present the schemes. Fig. 1 shows the measurement paths (MPs) they enable, i.e., the links *covered* by corresponding measurements of a uni- (downstream) or bidirectional (up- and downstream) monitor for one flow.

Spin bit. Spin bit clients flip the received spin bit value while servers reflect it, creating a square wave equal to the round-trip time (RTT). Unidirectional monitors can extract the RTT and cover the round-trip path. Bidirectionality enables distinguishing up- and downstream segments [28].

L bit. The L bit requires sender-side loss detection. The sender sets the L bit on one packet for each packet deemed lost. Monitors can assess the number of lost packets by counting packets with set L bits. The L bit covers the entire upstream MP; bidirectionality adds the downstream MP.

Q bit. The Q bit creates a square wave by alternately sending n packets with an unset/set Q bit. If n is known, unidirectional monitors determine the number of packets lost up to the monitor by comparing the tracked Q bit phase lengths to n. Bidirectionality adds the corresponding opposite MP.

R bit. The R bit uses the mechanism of and builds upon the Q bit. Instead of a fixed n, it reports the last received Q bit phase length. Unidirectional monitors can compare the observed R bit phase lengths to the original Q bit n to estimate packet loss from the original sender via the receiver to the opposite-side monitor. Bidirectionality adds the opposite MP.

Study objectives. Studying the applicability of EFM to diverse network management tasks, we identify NT as one task that likely benefits. In particular, the EFM schemes yield production traffic performance insight, thus complementing traditional active measurements, which struggle with fully grasping these conditions. In the following, we examine the efficacy of EFM-based NT, for which we first design a suitable NT system (Sec. 3) and describe our evaluation methodology (Sec. 4). We then evaluate the latency (Sec. 5) and packet loss (Sec. 6) estimation capabilities of EFM-based NT. In both settings, we compare its performance to traditional NT and NT fueled with corresponding TCP-based variants to quantify the benefits of explicit measurement information.

3 NETWORK TOMOGRAPHY DESIGN

To fairly compare the raw capabilities of the different NT variants, we design an NT system that can run on active *and* passive input, only using standard NT methods.

Measurements. Our active NT baseline uses UDP pings. Our passive NT variant can use TCP- and EFM-based measurements. For TCP, we use related approaches on latency (Dart [41]) and packet loss (RouteScout [3]) measurements. For EFM, we use the four schemes presented in Sec. 2.2. Each variant independently fuels one NT instance.

How to derive the network state? We choose an *algebraic* approach and derive mean latencies and packet loss ratios of the network links. We take the mean of all individual measurements for each flow and scheme, then determine the provided MPs, before finally attempting to derive the network state. For each scheme, we collect the measurements in a system of linear equations (cf. Sec. 2.1), adding entries to connectivity matrix $A \in \mathbb{R}^{m \times n}$ and measurement vector $b \in \mathbb{R}^m$ to obtain link characteristics $x \in \mathbb{R}^n$, resulting in

$$A \cdot x + \epsilon = b \Leftrightarrow ||A \cdot x + \epsilon - b|| = 0. \quad (1)$$

Measurement jitter ϵ is the main reason for estimation errors. Uniquely solving Eq. (1) is only possible if there are exactly as many linearly independent measurements m as there are links n ($m = n$). Thus, we often only approximate the actual solution. When using standard routing, most measurements only use a small share of the available links such that A is likely to be sparse. Hence, we use the LSQR algorithm [37] for solving Eq. (1) as it is efficient for sparse matrices. The resulting $\vec{x}$ yields estimates for *all* network link metrics.

Parameters. NT has a broad parameter space; fully exploring this space or optimizing performance is out of the scope of this work. Instead, we provide a first look at the general capabilities of EFM-based NT, choosing straightforward methods for flow and monitor selection. In particular, we perform pair-wise pings between all network nodes for our active NT variant. For EFM- and TCP-based NT, we monitor all production traffic flows on all available monitors.

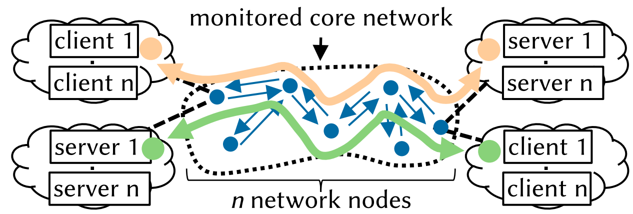

Figure 2: Each real node maps to one ns-3 node, each real link to two directional ns-3 links. We connect one client/server node to each network node, each holding one application for each of the n total network nodes.

Hence, both cases likely represent the best possible scenario. In practice, operators aim to minimize costs and would, e.g., carefully select between which nodes to perform pings or which flows to track. We leave exploring corresponding flow and monitor selection strategies for future work.

4 EVALUATION METHODOLOGY

We conduct our study with ns-3 [40] to obtain a reliable ground truth, simultaneously use all passive variants, and flexibly evaluate various parameters. We publish our code [29] for reproducibility. To ensure a fair assessment, we assume active probes and production traffic experience the same network conditions. We also assume knowledge of the network topology and forwarding state to reliably determine MPs.

Topologies and network setup. We aim to study representative Internet service provider (ISP) networks, but there are only few recent, usable, and publicly available topologies. Hence, we choose publicly available topologies of two European research networks, CESNET [13] and GÉANT [20], and a reference topology of the German Tier-1 ISP Telekom [24].

Our ns-3 networks use a dedicated network node for each node of the topology, as visualized in Fig. 2. Since we model networks as *directed* graphs, we map each real link to two directed links in ns-3, which we configure with a bandwidth of 10 Gbps. For GÉANT, we set the link latencies to the speed of light latency between the geographical locations such that latencies range from 0.2 ms to 9.5 ms with a mean latency of 1.8 ms. Due to their geographical density, we configure fixed latencies of 1 ms for each link of CESNET and Telekom.

Traffic setup. We generate payload traffic as follows: first, we add two end-hosts to each network node to represent foreign networks, one host representing distributed clients, the other representing servers as illustrated in Fig. 2. Each client establishes a TCP connection to one server such that we obtain pair-wise connections between all network nodes, including to the server of the same node, but with varying routes and end-to-end latencies. The flows use shortest path routing, and each flow transmits the median website size of 2022 [2], i.e., 2.3 MB of data, using TCP's Cubic congestion

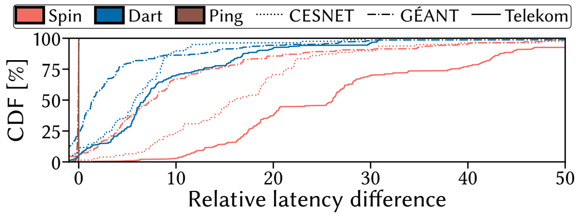

Figure 3: Relative latency difference between raw NT estimations and ground truth in different networks.

control. As access links often represent the overall bottleneck, we use dedicated 50 Mbps edge links with a fixed delay of 10 ms. Each ping flow transmits 40 probes of 40 B in 1 s intervals, i.e., the default frequency of Linux's *ping*.

Passive measurements. To compare the EFM schemes and assess their benefit compared to implicit semantics, we track all EFM schemes and TCP variants simultaneously. This approach would require significant modifications in real networks as no real transport protocol includes all schemes. In our ns-3 simulation, we easily accommodate these changes with *virtual measurement headers*: We tag each payload packet with metadata for *all* schemes so that we can apply each scheme concurrently. Monitors deployed on all network nodes use the virtual measurement headers as the source for their monitoring. We then collect all individual measurements and analyze the overall behavior via post-processing.

Latency and packet loss ground truth. Our simulation tracks latency and packet loss ground truths for all links. For packet loss, we count the total number of transmitted packets and packets dropped at any point from reception at the link endpoint until the successful forwarding. For latency, we measure the time from the reception on a node until forwarding to the next node. This design does not capture the metrics on the first hops, i.e., the links outside of the monitored network, reflecting a realistic scenario with limited visibility into external networks.

5 LATENCY TOMOGRAPHY

Assessing link latencies is a common use case [35], and we focus on the corresponding capabilities of EFM-based NT as the first part of our study, starting with a baseline case.

5.1 Baseline Scenario

In our baseline scenario, we assess the general latency estimation capabilities when fueling NT with the spin bit in comparison to active pings and Dart. Aiming to assess the behavior when having the richest information and to rule out that links are not covered by measurements, we place monitors on all network nodes and use full pair-wise payload traffic and pings. We perform 30 measurement repetitions and combine the corresponding results for our assessment.

Results. Fig. 3 shows a CDF of the relative difference between the latency estimations of the three latency tomography variants and the mean ground truth latency per link as determined over all actual per-packet delays in the corresponding simulation run. A value of 1, e.g., corresponds to a 100 % overestimation of the baseline. We cut the x-axis at -1 as smaller values correspond to negative link estimations.

As can be seen, NT using the spin bit and Dart significantly overestimates the real latencies. The reason is that our system of linear equations only models links of the monitored network and not the external edge links. However, these links have high base latencies and represent the overall bottlenecks where congestion with queuing delay arises. This leads to large fluctuations in the measured latencies, which the LSQR approach tries to map to the core links, thus causing the overestimations. For ping, this effect does not occur as the measurements only cover the monitored network such that ping-based NT achieves high accuracy.

Takeaway. *Raw spin bit measurements cannot be used for NT with the system of linear equations as the latency introduced by external networks causes significant overestimation.*

Next, we assess if and how specific characteristics of the spin bit can help to improve the estimations.

5.2 Spin Bit & Dart Path Segmentation

Raw spin bit and Dart measurements include foreign networks which prevents sensible estimations. However, both schemes enable focusing the measurements on specific links using path segmentation under certain conditions.

Path segmentation. Path segmentation requires bidirectional visibility of a flow by two monitors; Dart and the spin bit can then measure the half-RTT on each monitor. This is always possible for Dart as it needs bidirectional visibility anyway, but it might reduce the measurable flows for the spin bit. Once both monitors have determined the half-RTT in the same direction, the operator can deduce the bidirectional latency of the links between the monitors by subtracting the measurement of the monitor that covers fewer links from the other monitor. We apply the described methodology on all monitors and repeat our evaluation from Sec. 5.1.

Results. Fig. 4 compares the path-segmented Dart and spin bit results with the raw ping-based estimation. The path segmentation drastically improves the results such that the spin bit now closely matches ping for CESNET and Telekom, where both achieve almost perfect estimation. Both are more inaccurate for GÉANT, which has a more diverse latency profile, while Dart shows higher inaccuracies throughout all topologies. We attribute the difference between the spin bit and Dart to Dart's reliance on implicit TCP semantics, as the explicit spin bit semantics seem to enable more accurate estimations of the underlying flow performance.

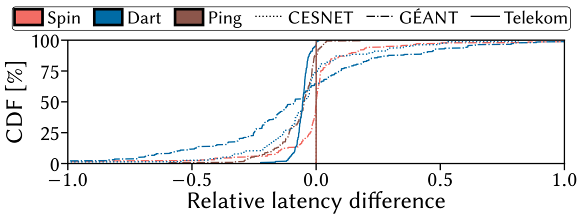

Figure 4: Path segmentation of the spin bit and Dart significantly improves the latency estimation.

Takeaway. *Spin bit path segmentation significantly improves the estimation performance and can reasonably fuel EFM-based NT. However, it requires bidirectional flow visibility and monitor cooperation. Finally, the explicit information of the spin bit outperforms Dart's implicit semantics.*

Next, we focus on loss tomography, which targets a more sporadic pattern that is hard to measure instantly.

6 LOSS TOMOGRAPHY

Compared to latency, packet loss has a more fluctuating behavior as it can, e.g., happen in bursts or only sporadically. Hence, measurements might easily miss or undersample these events. In this section, we study the performance of EFM-based NT regarding packet loss estimation. For this, we configure small random packet loss rates on some links and then repeat the general measurement setup from Sec. 5.

6.1 Baseline Scenario

Similar to Sec. 5, we start with a baseline where we use the raw output of the three EFM schemes (L, Q, R) and feed it into Eq. (1), reusing all parameters from Sec. 5.1. We compare the resulting estimations to NT based on ping and RouteScout. To focus on packet loss, we add random packet loss of 1 % to a single link, using 30 randomly generated scenarios overall. We measure each of these 30 scenarios once and combine the corresponding results for our assessment.

Results. Fig. 5 shows the packet loss estimation of the different approaches: The left plot shows the *relative* estimation error for links *with* packet loss, i.e., a value of 1 corresponds to an overestimation by 100 %. The right plot shows the *absolute* estimation error for links *without* packet loss, i.e., the real loss rate is 0 and the shown value is the estimation. We cut the x-axis of the left plot at −1.1 as values below −1 correspond to negative packet loss rates.

Looking at the left plot, we observe that most EFM schemes have a large estimation range with significant overestimation. However, the patterns for the L and R bits and RouteScout are similar. The Q bit tends to slightly underestimate the loss rate but shows a relatively stable performance across all scenarios. Our ping approach does not capture the packet loss as the 40 pings significantly undersample the shaped

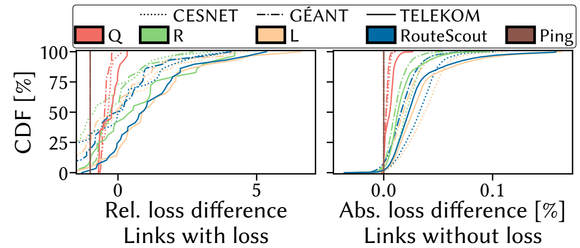

Figure 5: Loss difference between the NT estimations and ground truth.

packet loss rate [10]. In practice, there is no straightforward configuration regarding how many packets are needed as this number depends on the loss rate to be measured, and such measurements can entirely miss short-lived packet loss.

As we only configure packet loss on a single link, almost all links will not be subject to loss. Hence, we also inspect the absolute estimation error for those links in the right plot. The EFM schemes almost always overestimate the real packet loss rate. The L bit causes the largest overestimations as it assesses the end-to-end packet loss rate, which, again, includes the edge links, as was the case for latency. The Q bit provides the most accurate EFM results, although it is also affected by packet loss on the edge links. In contrast, the ping variant always correctly estimates a loss rate of 0.

Takeaway. *EFM-based NT can capture packet loss to some extent but is significantly affected by external links. Active measurements struggle even more as there is no straightforward way to configure how many probes are needed.*

As for latency, the EFM schemes suffer from external effects as the edge links will eventually cause packet loss for every flow. We again aim to equal out this external packet loss via path segmentation as we detail next.

6.2 Q Bit Path Segmentation

Of the available EFM (and TCP) schemes, only the Q and R bits enable path segmentation. However, the R bit relies on the Q bit and enables similar paths, such that we focus on the Q bit as a measurement scheme that can be used standalone.

Path segmentation. The Q bit path segmentation requires that a flow is observed by multiple monitors; each reports the number of observed Q bit markings. Operators can subtract the observed markings of a monitor later on the path from the ones of a monitor earlier on the path to deduce the packet loss between the observers. We apply this methodology to our results from Sec. 6.1 and repeat our evaluation, further broadening the studied packet loss spectrum.

Results. Fig. 6 visualizes the *relative* estimation error for links *with* packet loss in two scenarios: the left plot shows settings with a single faulty link but different loss rates, whereas the right plot has a fixed loss rate of 1 % but one to

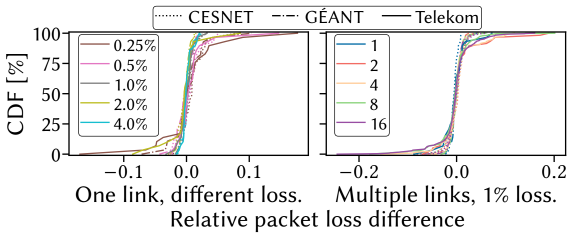

Figure 6: Path segmentation enables accurate Q bit estimations for various loss rates and impaired links.

16 randomly selected faulty links. As can be seen, path segmentation again significantly increases the loss estimation accuracy in all scenarios, with the estimation being within 20 % of the actual loss rate. Additionally, the segmented Q bit variant no longer falsely attributes any packet loss to links without actual packet loss (not shown), which increases the usefulness of such estimations as operators do not need to study false alarms. Upon closer inspection, we do observe differences, e.g., that the Telekom topology seems to be the most difficult as it has many of the highest estimation errors. Overall, however, the path-segmented Q bit works well.

Takeaway. *The path segmentation capabilities of the Q bit allow for very accurate packet loss estimation, irrespective of the packet loss rate or number of impaired links. It also has high practical usefulness with no false positives.*

In summary, our study shows that EFM-based NT can provide sensible results for latency and packet loss estimation. The path segmentation capabilities of the spin and Q bits are crucial in this process to filter external effects.

7 DISCUSSION

Our study shows that EFM can complement traditional active measurements for NT. In this section, we discuss implications of our study for bringing EFM-based NT to real networks and position our work in the context of related work.

EFM. The measurement behavior of the spin bit and the loss measurement EFM schemes has been studied in multiple works [12, 17, 32]. Additional work has examined the behavior and adoption of the spin bit on the web [30]. However, these works do not consider the applicability of EFM to specific network management tasks. For NT, we identify the spin and Q bits as the best choices as they enable high estimation accuracy with a total cost of only two bits.

NT methodology. Most related work on NT makes theoretical contributions (e.g., [8, 23]). Only few works look at practical angles of NT [21, 26]. We uncover practical challenges when using standard NT concepts for EFM-based NT. For example, EFM estimations can include links not modeled in our system of linear equations, significantly complicating latency and packet loss tomography. While path segmentation can alleviate such issues, it imposes additional requirements: flows need to be tracked by at least two monitors and with bidirectional visibility in case of the spin bit. Lowering these hurdles, we see worthwhile future work in approaches that can natively handle external links.

Gathering flow metrics. Standard flow monitoring, like IPFIX [14], does not provide latency or packet loss information. However, Dart [41], RouteScout [3], and the spin bit [31] map to high-speed hardware, and there are ongoing efforts to extend IPFIX for latency [22]. Hence, future work could focus on collecting the EFM output in a standard way.

Path information. Accurate path information is needed to map the flow metrics to path segments. Our simulation provides this information, but accessing it in real networks can be challenging. Concepts such as Magnifier [11] have already shown the feasibility of assessing assumptions on the current forwarding. Future work could extend these approaches to provide the required information for NT.

Flow selection. EFM-based NT must select flows at runtime, as tracking all flows in large networks is infeasible [25]. Ideally, the selected flows cover all links, and their MPs differ enough to identify individual links. However, picking a flow selection strategy is challenging as there are many options. For example, random sampling is sufficient in many related applications [3, 25, 41], but Internet traffic is also generally imbalanced [25]. Additionally, path segmentation requires that the same flows are tracked on multiple monitors, but *randomly* sampling the same flow at multiple locations is rare [42]. We argue that developing effective and efficient flow selection strategies is critical for future work.

8 CONCLUSION

Explicit flow measurements (EFMs) provide a crucial building block for network management, giving insight into key flow performance metrics. However, the utility of EFM schemes for network management tasks beyond plain measurements has never been studied before. In this paper, we have evaluated whether EFM can effectively fuel the network management task of network tomography (NT). In particular, traditional NT uses active probes to assess the network state, but such measurements do not necessarily accurately reflect the behavior of production traffic; EFM could provide the needed input. Our results from a broad ns-3 study suggest that raw EFM output struggles with measurement noise induced by foreign networks, failing to facilitate meaningful NT results. However, spin and Q bit path segmentation capabilities do enable high-accuracy latency and packet loss estimation, although at the cost of cooperation between monitors. In conclusion, EFM provides a sensible alternative input for NT, proving their general utility for network management.

ACKNOWLEDGMENTS

This work has been funded by the German Research Foundation DFG under Grant No. WE 2935/20-1 (LEGATO).

REFERENCES

[1] Mark Allman, Robert Beverly, and Brian Trammell. 2017. Principles for Measurability in Protocol Design. *ACM SIGCOMM Computer Communication Review* 47, 2 (2017), 2–12. doi:10.1145/3089262.3089264

[2] Web Almanac. 2022. Part IV Chapter 21 - Page Weight. https://almanac.httparchive.org/en/2022/page-weight Accessed: 2025-02-08.

[3] Maria Apostolaki, Ankit Singla, and Laurent Vanbever. 2021. Performance-Driven Internet Path Selection. In *Proceedings of the 2021 ACM SIGCOMM Symposium on SDN Research (SOSR)*. doi:10.1145/3482898.3483366

[4] Viviana Arrigoni, Novella Bartolini, Annalisa Massini, and Federico Trombetti. 2021. Failure Localization through Progressive Network Tomography. In *Proceedings of the 2021 IEEE International Conference on Computer Communications (INFOCOM)*. doi:10.1109/INFOCOM42981.2021.9488893

[5] Behnaz Arzani, Selim Ciraci, Luiz Chamon, Yibo Zhu, Hongqiang Liu, Jitu Padhye, Boon Thau Loo, and Geoff Outhred. 2018. 007: Democratically Finding the Cause of Packet Drops. In *Proceedings of the 2018 USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. https://www.usenix.org/conference/nsdi18/presentation/arzani

[6] François Aubry, David Lebrun, Stefano Vissicchio, Minh Thanh Khong, Yves Deville, and Olivier Bonaventure. 2016. SCMon: Leveraging Segment Routing to Improve Network Monitoring. In *Proceedings of the 2016 IEEE International Conference on Computer Communications (INFOCOM)*. doi:10.1109/INFOCOM.2016.7524410

[7] Ritwik Banerjee, Abbas Razaghpanah, Luis Chiang, Akassh Mishra, Vyas Sekar, Yejin Choi, and Phillipa Gill. 2015. Internet Outages, the Eyewitness Accounts: Analysis of the Outages Mailing List. In *Proceedings of the 2015 International Conference on Passive and Active Network Measurement (PAM)*. doi:10.1007/978-3-319-15509-8_16

[8] Novella Bartolini, Ting He, Viviana Arrigoni, Annalisa Massini, Federico Trombetti, and Hana Khamfroush. 2020. On Fundamental Bounds on Failure Identifiability by Boolean Network Tomography. *IEEE/ACM Transactions on Networking* 28, 2 (2020), 588–601. doi:10.1109/TNET.2020.2969523

[9] Yigal Bejerano and Rajeev Rastogi. 2003. Robust Monitoring of Link Delays and Faults in IP Networks. In *Proceedings of the 2003 IEEE International Conference on Computer Communications (INFOCOM)*. doi:10.1109/INFCOM.2003.1208666

[10] Gino Biondini. 2015. An Introduction to Rare Event Simulation and Importance Sampling. In *Handbook of Statistics*. Elsevier, 29–68. doi:10.1016/B978-0-444-63492-4.00002-2

[11] Tobias Bühler, Romain Jacob, Ingmar Poese, and Laurent Vanbever. 2023. Enhancing Global Network Monitoring with Magnifier. In *Proceedings of the 2023 USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. https://www.usenix.org/conference/nsdi23/presentation/buhler

[12] Fabio Bulgarella, Mauro Cociglio, Giuseppe Fioccola, Guido Marchetto, and Riccardo Sisto. 2019. Performance Measurements of QUIC Communications. In *Proceedings of the 2019 ACM/IRTF Applied Networking Research Workshop (ANRW)*. doi:10.1145/3340301.3341127

[13] CESNET. 2023. CESNET2/3: IP/MPLS Backbone Utilization. https://netreport.cesnet.cz/netreport/CESNET2_IP_MPLS_backbone_utilization/index.html

[14] Benoît Claise, Brian Trammell, and Paul Aitken. 2013. *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*. RFC 7011. IETF. doi:10.17487/RFC7011

[15] Mark Coates, Alfred O. Hero, Robert Nowak, and Bin Yu. 2002. Internet Tomography. *IEEE Signal Processing Magazine* 19, 3 (2002), 47–65. doi:10.1109/79.998081

[16] Mauro Cociglio, Alexandre Ferrieux, Giuseppe Fioccola, Igor Lubashev, Fabio Bulgarella, Massimo Nilo, Isabelle Hamchaoui, and Riccardo Sisto. 2023. *Explicit Host-to-Network Flow Measurements Techniques*. RFC 9506. IETF. doi:10.17487/RFC9506

[17] Piet De Vaere, Tobias Bühler, Mirja Kühlewind, and Brian Trammell. 2018. Three Bits Suffice: Explicit Support for Passive Measurement of Internet Latency in QUIC and TCP. In *Proceedings of the 2018 ACM Internet Measurement Conference (IMC)*. doi:10.1145/3278532.3278535

[18] Cuiying Feng, Jianwei An, Kui Wu, and Jianping Wang. 2022. Bound Inference and Reinforcement Learning-Based Path Construction in Bandwidth Tomography. *IEEE/ACM Transactions on Networking* 30, 2 (2022), 501–514. doi:10.1109/TNET.2021.3118006

[19] Alexane Ferrieux, Igor Lubashev, Giuseppe Fioccola, Isabelle Hamchaoui, Massimo Nilo, Fabio Bulgarella, and Mauro Cociglio. 2025. *Application of Explicit Measurement Techniques for QUIC Troubleshooting*. Internet-Draft. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-mdt-quic-explicit-measurements Work in progress.

[20] GÉANT. 2025. GN4-3N. https://network.geant.org/gn4-3n/ Accessed: 2025-02-08.

[21] Denisa Ghita, Can Karakus, Katerina Argyraki, and Patrick Thiran. 2011. Shifting Network Tomography Toward A Practical Goal. In *Proceedings of the 2011 ACM International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*. doi:10.1145/2079296.2079320

[22] Thomas Graf, Benoît Claise, and Alex Huang Feng. 2024. *Export of Delay Performance Metrics in IP Flow Information eXport (IPFIX)*. Internet-Draft. Internet Engineering Task Force. https://datatracker.ietf.org/doc/draft-ietf-opsawg-ipfix-on-path-telemetry Work in progress.

[23] Ting He, Athanasios Gkelias, Liang Ma, Kin K. Leung, Ananthram Swami, and Don Towsley. 2017. Robust and Efficient Monitor Placement for Network Tomography in Dynamic Networks. *IEEE/ACM Transactions on Networking* 25, 3 (2017), 1732–1745. doi:10.1109/TNET.2016.2642185

[24] Oliver Hohlfeld. 2008. *Statistical Error Model to Impair an H.264 Decoder*. Master's Thesis. Technical University of Darmstadt. http://downloads.ohohlfeld.com/theses/hohlfeld-masterthesis.pdf

[25] Thomas Holterbach, Edgar Costa Molero, Maria Apostolaki, Alberto Dainotti, Stefano Vissicchio, and Laurent Vanbever. 2019. Blink: Fast Connectivity Recovery Entirely in the Data Plane. In *Proceedings of the 2019 USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. https://www.usenix.org/system/files/nsdi19-holterbach.pdf

[26] Yiyi Huang, Nick Feamster, and Renata Teixeira. 2008. Practical Issues with Using Network Tomography for Fault Diagnosis. *ACM SIGCOMM Computer Communication Review* 38, 5 (2008), 53–58. doi:10.1145/1452335.1452343

[27] Jana Iyengar and Martin Thomson. 2021. *QUIC: A UDP-Based Multiplexed and Secure Transport*. RFC 9000. IETF. doi:10.17487/RFC9000

[28] Mirja Kühlewind and Brian Trammell. 2022. *Manageability of the QUIC Transport Protocol*. RFC 9312. IETF. doi:10.17487/RFC9312

[29] Ike Kunze. 2025. Using Explicit (Host-to-Network) Flow Measurements for Network Tomography. https://github.com/COMSYS/ns3-efm-study-overview

[30] Ike Kunze, Constantin Sander, and Klaus Wehrle. 2023. Does It Spin? On the Adoption and Use of QUIC's Spin Bit. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC)*. doi:10.1145/3618257.3624844

[31] Ike Kunze, Constantin Sander, Klaus Wehrle, and Jan Rüth. 2021. Tracking the QUIC Spin Bit on Tofino. In *Proceedings of the 2021 Workshop on the Evolution, Performance, and Interoperability of QUIC (EPIQ)*. doi:10.1145/3488660.3493804

[32] Ike Kunze, Klaus Wehrle, and Jan Rüth. 2021. L, Q, R, and T - Which Spin Bit Cousin Is Here to Stay?. In *Proceedings of the 2021 ACM/IRTF Applied Networking Research Workshop (ANRW)*. doi:10.1145/3472305.3472319

[33] Yichao Li and Ke Deng. 2022. Simultaneous Topology and Loss Tomography via a Modified Theme Dictionary Model. *IEEE Transactions on Signal Processing* 70 (2022), 4239–4251. doi:10.1109/TSP.2022.3191807

[34] Marc Linsner, Philip Eardley, Trevor Burbridge, and Frode Sorensen. 2015. *Large-Scale Broadband Measurement Use Cases*. RFC 7536. IETF. doi:10.17487/RFC7536

[35] Liang Ma, Ting He, Kin K. Leung, Ananthram Swami, and Don Towsley. 2013. Identifiability of Link Metrics Based on End-to-end Path Measurements. In *Proceedings of the 2013 ACM Internet Measurement Conference (IMC)*. doi:10.1145/2504730.2504738

[36] Edgar Costa Molero, Stefano Vissicchio, and Laurent Vanbever. 2022. FAst In-Network GraY Failure Detection for ISPs. In *Proceedings of the 2022 ACM SIGCOMM Conference*. doi:10.1145/3544216.3544242

[37] Christopher C. Paige and Michael A. Saunders. 1982. LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares. *ACM Trans. Math. Software* 8, 1 (1982), 43–71. doi:10.1145/355984.355989

[38] Cristel Pelsser, Luca Cittadini, Stefano Vissicchio, and Randy Bush. 2013. From Paris to Tokyo: On the Suitability of Ping to Measure Latency. In *Proceedings of the 2013 ACM Internet Measurement Conference (IMC)*. doi:10.1145/2504730.2504765

[39] Peng Qin, Bin Dai, Benxiong Huang, Guan Xu, and Kui Wu. 2014. A Survey on Network Tomography With Network Coding. *IEEE Communications Surveys & Tutorials* 16, 4 (2014), 1981–1995. doi:10.1109/COMST.2014.2320096

[40] George F. Riley and Thomas R. Henderson. 2010. The ns-3 Network Simulator. In *Modeling and Tools for Network Simulation*. Springer, 15–34. doi:10.1007/978-3-642-12331-3_2

[41] Satadal Sengupta, Hyojoon Kim, and Jennifer Rexford. 2022. Continuous In-Network Round-Trip Time Monitoring. In *Proceedings of the 2022 ACM SIGCOMM Conference*. doi:10.1145/3544216.3544222

[42] Olivier Tilmans, Tobias Bühler, Ingmar Poese, Stefano Vissicchio, and Laurent Vanbever. 2018. Stroboscope: Declarative Network Monitoring on a Budget. In *Proceedings of the 2018 USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. https://www.usenix.org/conference/nsdi18/presentation/tilmans

[43] Xian Zhang and Chris Phillips. 2012. A Survey on Selective Routing Topology Inference Through Active Probing. *IEEE Communications Surveys & Tutorials* 14, 4 (2012), 1129–1141. doi:10.1109/SURV.2011.081611.00040