

AN EXTENSIVE SURVEY OF IN-BAND NETWORK TELEMETRY: HISTORICAL INSIGHTS, CURRENT DEVELOPMENTS, AND SIGNIFICANT CHALLENGES

SAMI ABBAS NAGAR ADAM¹, AISHA HASSAN ABDALLA HASHIM^{1,2}, OTHMAN OMRAN KHALIFA^{1,3}

¹ Department of Electrical and Computer Engineering, International Islamic University Malaysia, Malaysia

²Dept. of Electrical and Electronic Engineering Science, University of Johannesburg, South Africa

³Libyan Center for Engineering Research and Information Technology, Bani Waleed, Libya

E-mail: elnajarsami@yahoo.com, aisha@iium.edu.my, khalifa@iium.edu.my

ABSTRACT

While In-Band Network Telemetry (INT) is revolutionizing network monitoring, the current literature lacks in-depth synthesis of the shortcomings and open questions associated with INT architecture. This paper provides critical review of INT technology based on systematic analysis of 80+ primary studies published from 2015 to 2026. There is a critical gap in understanding INT technology because of the absence of the general framework of evaluating the solutions with regard to visibility-overhead-security-ASIC capacity trade-off. The study finds out that 76% of INT solutions are concentrated on solving the problem of bandwidth, however, do not take into account issues like the effect of INT traffic on depletion of ASIC resources, TCAM/SDRAM bottleneck and security risks associated with INT. In this paper, a novel four-dimensional Visibility-Fidelity-Overhead-Security (VFOS) framework is developed and used to analyze the state of research in INT and demonstrate gaps such as absence of a standard approach in modeling heterogeneous telemetries and insufficient use of artificial intelligence. Contributions of this paper include (i) a critical taxonomy of distinction between measurement fidelity and operational sustainability, (ii) the first empirical proof that the scaling capacity of INT is limited by a non-linear scaling wall that depends on the network diameter, and (iii) an exciting future research agenda based on adaptive telemetry and hardware/software co-design approach.

Keywords: *In-Band Network Telemetry, Programmable Data Plane, Network Monitoring, Critical Analysis, Software-Defined Networking, P4 Programming.*

1. INTRODUCTION: THE TELEMETRY IMPERATIVE AND ITS DISCONTENT

1.1 The Growing Complexity of Network Management

The modern world of networks brings to light a paradox never seen before: despite the ever-growing complexity of the network infrastructure with regards to the increase in scale, variety and dynamism of the system, the methods of monitoring and managing the network remain firmly rooted in paradigms created in times of relatively static and monolithic structures [1]. Today's network management requires facing a host of challenges ranging from micro millisecond requirements for latency, to terabit per second throughput needs, to the use of multiple clouds crossing geopolitical borders and up to new applications such as autonomous cars and virtual reality [63]. With the assumption that network behavior may be deduced from aggregate statistics

and sampled data as a basis of modern approaches to network management [3], these techniques have become outdated and inadequate.

The primary deficiency in the conventional monitoring approaches is rooted in the use of what we refer to as the gap in observability and complexity: the growing discrepancy between the temporal and spatial resolution needed to manage the network efficiently and the abilities of the current infrastructure for monitoring [2]. Although active approaches based on Ping and Traceroute [11] are quite easy to implement, the main issue that arises here is the problem of probe-traffic divergence that results in the inability to obtain accurate measurements [12]. The passive monitoring schemes, such as NetFlow [4], sFlow [5] or IPFIX [6, 7], although measuring user traffic accurately, employ probabilistic sampling, which, due to its inherent nature, is unable to capture the temporal dynamics necessary to detect microbursts,

congestion incidents and failures [13]. The hybrid monitoring schemes, in turn, complicate the situation even further and, more often than not, fail to alleviate the basic trade-off between completeness of the measurement and additional efforts required to conduct it [2]. Figure 1 below shows the development and limitations of the monitoring paradigms. It shows the development from traditional active and passive monitoring methods towards the new In-Band Network Telemetry approach.




Figure 1: Evolution of Network Monitoring Paradigms

1.2 The Origin and Potential of In-Band Network Telemetry

The problem led to a new era in networking as In-Band Network Telemetry (INT), which emerged in 2015 as the result of collaboration between Intel, Dell, ARISTA, VMware, and Barefoot Networks [19], came into existence as an answer to these challenges. The main concept behind the INT is completely different from all previous approaches in that INT involves embedding the instruction to measure the network state right into the data packet, allowing each intermediate switch on the forwarding path to append to the packet header its metadata (switch identity, queue sizes, timestamps, ingress and egress ports). It leads to the following three breakthroughs:

First, per-packet, hop-by-hop visibility becomes possible without the sampling bias built into the probabilistic approaches. Packets become bearers of network intelligence with the capability of detecting anomalies on a microsecond scale, something not feasible with traditional aggregated statistics [13]. Secondly, in-line, in-network monitoring obviates the lag between collecting the data and the actual analysis that characterizes out-of-network approaches, making possible the closed loop approach in which network devices act upon congestion issues within the same timeline [39]. Thirdly, programmability facilitated by P4 [52] and

other domain-specific languages makes it possible for network engineers to specify the type of metadata to be gathered to implement flexible telemetry schemes.

The theoretical attractiveness of INT is indisputable and early proof of concepts – especially in data center spines and leaves as well as 5G/6G network slicing – showed INT's tremendous potential in network management applications [63]. But as will be made clear in this paper, there is a huge gap between the theory of INT and its practice with many unsolved problems in need of critical assessment [19].

1.3 The Critical Analysis Mandate: Importance of this Research

Although a plethora of studies concerning INT have been performed and published, and our comprehensive review revealed no less than 86 individual proposals, models, and implementations in between 2015 and 2026, there is a consistent and worrying characteristic feature of the body of knowledge: a great number of papers (as many as 83.7% of those included in our review) make an incremental progress in the development of INT architecture without putting a critical question to the underlying premises and limitations of the method [62,64,65]. Three problems underlined by the unthinking proliferation include:

Bandwidth-Efficiency Tradeoff Problem: The bulk of INT research is oriented towards optimization for bandwidth overhead by means of either compressing or otherwise reducing the amount of information sent through INT mechanism [66,68]. Our critical analysis, however, shows that the trade-off for efficiency improvement consists in increase in the computational complexity for switch data plane and may actually aggravate the very problem addressed through use of the technique [80].

Paradox of Security-Visibility Trade-off: By design, INT provides exposure of network topology and flow information inside packet headers while they pass through untrusted network domains [22]. While the visible nature of INT is appreciated by researchers, at the same time this characteristic introduces major privacy and security challenges. Our critical analysis has found that there is lack of focus on security concerns arising from exposing hop-by-hop information to adversaries in current state-of-the-art INT research [21].

Chasm between Heterogeneity and Interoperability: A majority of current INT solutions has been studied in controlled, homogeneous testbeds like data centers with identical hardware and uniform

management [59]. However, the future of network architecture will be characterized by heterogeneity due to different types of hardware, administrative domains, protocols, and hybrid public and private clouds and edge networks [61].

This paper attempts to bridge this gap by performing a critical synthesis that goes beyond mere cataloguing of INT techniques and focuses on evaluating the fundamental limits, potential side-effects and open challenges of these techniques. Our contributions include the following.

1. Critical Taxonomy: Our first key contribution is the development of a critical taxonomy framework for classification of the INT techniques not based on the benefits of each but rather on the underlying trade-offs in four areas: Visibility, Fidelity, Overhead and Security (VFOS) [19].

2. Limits of Scalability: By applying both theoretical analysis and empirical synthesis of the experimental results, we show the fundamental nature of the scaling problems related to non-linear increase of INT overhead depending on network diameter and number of flows [62].

3. Research Agenda: Based on our critical review, we identify and formulate an agenda for future research with the focus on adaptive, cognitive telemetry approaches, cross-domain collaboration and hardware-software co-design in consideration of switching ASIC limitations [80].

1.4 Structure of This Critical Analysis

This paper is structured to enable a critical assessment of its subject matter in an efficient way. First, Section 2 introduces the methodology for systematic review and critical analysis. Then, Section 3 provides a critical classification of INT architecture in terms of trade-offs associated with their evolution. Next, in Section 4, we analyze the uses of INTs and determine whether the uses reported in the literature constitute valid applications as opposed to merely desirable aspirations lacking any evidential support. Section 5 outlines the results of our critical analysis of INT limitations in terms of scalability issues, hardware limits, security issues, and model heterogeneity. Section 6 identifies and specifies the implications of our research for future studies in this domain. Section 7 describes the positioning of our paper against previous surveys and discusses the novelty of insights resulting from our critical analysis.

2. METHODOLOGY: SYSTEMATIC CRITICAL ANALYSIS

2.1. Literature Collection and Selection Criteria

The systematic critical analysis that we are undertaking has been based on the systematic study of INT literature, where our choice of literature has followed an extensive approach, which apart from the highly referenced and cited work, is expected to find other marginal and dissident viewpoints that will go against the conventional wisdom about the same. Searches were carried out within the IEEE Xplore database, ACM Digital Library, Google Scholar database, and Arxiv, and for that purpose, we used keywords such as "In-Band Network Telemetry," "INT," "In-Situ Operations Administration and Maintenance," "IOAM," "Programmable Data Plane Telemetry," and "P4 Telemetry" [52,56].

We were able to retrieve 347 publications, which underwent a three-phase literature screening process. The first phase excluded the duplicates and articles written in other languages than English to arrive at 298 articles. Phase 2 consisted of abstract and conclusion screening where we only included articles relevant to the architecture of implementation of the INT or application to filter out publications not contributing anything substantial about INT or were only about other aspects of telemetry [3]. From this process, we were able to filter out 146 publications. Phase 3 involved setting a minimum criterion for quality based on the type of venue of the publication and its level of citation to end up with 86 publications. The method for literature review and selection in our study is outlined in Figure 2.




Figure 2: Systematic Literature Review and Critical Analysis Methodology

2.2 Framework for Critical Analysis

Our analysis methodology is quite different from traditional survey methodologies since we have constructed a critical evaluation framework around four interrogation dimensions:

Architectural Assumptions: We look at the assumptions about the structure of each of the INT proposals in terms of network architecture, network traffic, hardware capabilities, and administrative boundaries, and we analyze whether such assumptions hold up in today's and future's networking landscape [19].

Experimental Evidence: For each proposal that asserts its performance advantages, we evaluate their experimental methodology and testing environment, dividing into those that have been validated in operational networks and simulated/emulated environments [60, 61].

Trade-Offs: For each of the solutions, we try to determine the trade-offs that they have, especially the triangular trade-off between telemetry resolution, bandwidth overhead, and computational complexity [62, 64, 68].

Unexplored Limitations: We try to identify the limitations discussed by the authors, as well as unexplored limitations that can be revealed through careful study of experimental parameters [65].

3. CRITICAL TAXONOMY OF INT ARCHITECTURES

3.1 History of Evolution: From Concept to Fragmentation

As evidenced by the systematic analysis we performed, the field of INT demonstrates evolutionary tendencies in line with the general evolution of programmable networking [1]. Indeed, its conceptualization (2015-2017) has been marked by a singular vision of the technology as a general-purpose substrate for telemetry providing exceptional visibility based on the latest possibilities offered by programmable data planes [52]. That era has generated seminal research into the creation of header formats, data plane implementations, and deployment architecture of INT [54], which had been mostly conducted by P4 Language Consortium and IOAM working group of the IETF [21,22]. Figure 3 provides a timeline of the evolution of INT. Key milestones of the evolution of In-Band Network Telemetry from its initial emergence in 2015 to current implementation.

Nonetheless, the following period (2018-2021) has been characterized by an increasing fragmentation of the INT field. As a result of dealing with the limitations such as ASIC processing constraints, packet size limitation, and bandwidth overhead, researchers and developers started to create specialized versions of INTs tailored for particular situations [13,14,15]. The diversity of approaches

like PINT [66], Sel-INT [13], PBT [23], iFIT [24] shows not an evolutionary trend but a tendency toward diverging incompatible design spaces [19]. Our observations show that while producing valuable information, such approach negatively impacts standardizing INTs that could operate together.

3.2 The VFOS Model: Describing INT Trade-offs

To comprehensively evaluate INT architecture design space, we define the Visibility-Fidelity-Overhead-Security (VFOS) model, an architectural characterization capturing the key trade-offs that INT systems have to make [19]:

Visibility (V): Level of visibility that an INT solution provides of network state information, in terms of spatial scope (which devices/elements are being measured) and temporal resolution (frequency of measurements) [13]. The level of visibility can be from sporadic (stochastic sampling, representative flows) to comprehensive (measurement of every single packet, every single hop) [66].

Fidelity (F): Accuracy and completeness of telemetry data compared to ground-truth. Fidelity consists of several metrics of precision of measurement (latency measurement accuracy, queue size estimates accuracy), completeness of collected metrics (collection of all important state variables), and consistency of measurements taken at multiple nodes [39].

Overhead (O): Cost of telemetry, in terms of bandwidth (extra bytes per packet), computation cost (cycling in processing metadata encoding and decoding) and storage costs (buffering and logging) [62]. The overhead is evident both in the data-plane (utilization of switch ASIC) and in the control plane (overhead incurred by collectors and analyzers) [80].

Security (S): This refers to the robustness of the system against potential attacks and the security of data [22]. Security aspects relate to data confidentiality (no unauthorized access to metadata in telemetry), integrity (metadata is not tampered with), and availability (robustness against telemetry-based denial of service) [21].

Figure 4 presents the VFOS framework as well as the multidimensional space that reflects the compromise between INT solutions. Four-dimensional framework for INT implementation characterization and fundamental trade-offs. Every dimension is defined in terms of the spectrum of INT solutions.




Figure 4: VFOS Critical Analysis Framework

Based on the VFOS framework, no INT solution is able to excel on all the four dimensions at the same time; each approach is a certain configuration of compromises depending on certain network, hardware and operational assumptions [19].

3.3 Passive and Active Approaches Reconsidered

The classification of methods into Passive Network Telemetry (PNT) and Active Network Telemetry (ANT), introduced by traditional measurement papers [2] has become outdated for programmable dataplanes [27]. It becomes clear during our discussion, that there are many aspects which the classical division of methods into passive and active approaches does not take into account [78].

Passive INT solutions are those when the telemetry is gathered within a normal data packets used in network. They are called so because no additional probes are sent. The most popular solutions of this type are P4-INT and IOAM [22]. The potential value of passive INT methods is obvious, and we will try to identify some shortcomings of them.

First of all, it should be noted that, being passive, INT method observes only the areas where the traffic goes through [13]. And sometimes there are blind spots due to non-uniform packet distribution. Metadata attachment also adds overhead to production packets, which may cause packet fragmentation for small MTU in certain situations [55] as well as influence forwarding efficiency in other ways [62]. Third, passive INT is flow-based and relies on current flows, so it is unable to perform proactive network probing and characterization when flows do not exist [78].

On the other hand, active INT involves injection of special probe packets which provide telemetry information along specified routes [78]. Although this provides stable coverage and regular intervals of measurements, we identify several negative aspects in our research [60]. The process of probe packet injection is resource-intensive and uses bandwidth, and in some cases may distort the network behavior that is under test—a typical example of observer effect [2]. In addition, probe

packets may follow different paths than production traffic, since the paths are determined by ECMP hashing or other mechanisms [61]. Figure 5 illustrates the difference between passive and active INT approaches.

Based on our VFOS analysis of the PNT/ANT domain, there emerges an interesting finding: studies that make claims to the comparison of these two techniques always tend to favor one technique by ignoring the inherent incommensurability of the two paradigms [19,78]. In our opinion, the very concept of PNT/ANT can be regarded more as a rhetorical tool, because the real challenge lies not in identifying whether telemetry is passive or active but in addressing how the aforementioned trade-offs are resolved.

3.4 Comparative Analysis of INT Implementations

In Table 1 below, we present a comparative analysis of some INT implementations, based on the VFOS analysis, enhanced by assessing the degree of deployment readiness in a qualitative manner. It should be stressed that the ratings are derived on the basis of the synthesized results obtained along with our perception of the problems associated with these results.

Critical Observations:

1. Overhead/Visibility Trade-off Is Not Linear [62]: The relationship between overhead decrease and visibility loss is not linear. For example, the probabilistic approach used by PINT [66] decreases overhead to 1-16 bits per packet but eliminates determinism which makes INT useful for microsecond-level anomaly detection.
2. No Implementation Is Fully Secured [22]: None of the implementations offers complete security guarantees, although it's obvious how dangerous can be exposing topology and flow information through packet headers [21]. As a result, "security" column in Table 1 is purely hypothetical because all the solutions are based on some trust assumptions.
3. Deployment Readiness Does Not Account for Hardware Constraints [80]: The values in deployment readiness column are calculated under assumption that Tofino programmable switches are available [68]. The deployment landscape in case of legacy hardware or FPGA solutions is quite immature, thus reducing deployment flexibility of INT.
4. Scalability Barrier [62]: In all the implementations, we found a clear trend—INT performance decreases non-linearly with network

diameter and number of flows. This effect is due to one inherent property—each node adds metadata to each packet, what creates scalability barrier no optimization can bypass [59].

4. CRITICAL EVALUATION OF PERFORMANCE TELEMETRY USING INT

4.1 One-Way Delay (OWD) Measurement

There is much research work discussing the measurement capabilities of one-way delay through INT-enabled performance telemetry [28,29]. Our critical evaluation shows that while the measurement capabilities do exist, their practical worth is taken for granted more than proven [19].

One-Way Delay (OWD) Measurement [28]: While being able to accurately measure one-way latency on a per-hop basis with microsecond precision does have its merits when it comes to the identification of network issues [29], we find that there are at least three major problems to overcome when it comes to OWD-INT measurement:

Timestamping of packets [28] and thus measurement of latency still poses a technical challenge in practice. While the microsecond precision is claimed, in reality, even though it may be theoretically accurate, GPS or PTP timestamps are often employed which makes the OWD-INT less accurate in practical applications than in lab environments [19].

Two, in addition to being path-specific and not flow-specific, delay data acquired through OWD-INT technology is also incomplete because of ECMP and multipath forwarding, where separate packets in the same flow will take different paths but the telemetry metadata collected by an individual packet is not going to tell the whole story about its flow [15].

Three, and this is the most important point, OWD measurement on a microsecond scale might be an interesting and useful technology, which is only useful in special use cases, including latency-critical tasks such as high-frequency trading or autonomous car guidance; however, for an overwhelming majority of enterprises and cloud workloads, a millisecond scale is enough [75]. Therefore, striving for microsecond accuracy is an optimization task that might not always be necessary.

Queue Depth Monitoring [30]: Although queue depths are intuitive indicators of the presence of a potential traffic congestion problem, they face the following challenge: by the moment when information regarding queue depth for a specific

packet becomes available, the queue depth value changes. Despite being an intrinsic feature of any monitoring technique [3], the INT enthusiasts exaggerate real time aspect of queue depth observation [39]. Indeed, the benefit of the technique is in offline processing of data and trend detection [75].

Two Way Delay and Quality of Service [31, 32]: Combination of QoS and TWD monitoring in INT context looks like a more relevant application [31], since such metrics directly correspond to SLA's and user experience [32]. Nevertheless, the analysis shows that two-way delay estimation based on INT suffers from all problems of one-way delay estimation, plus the additional problem of feedback loops [66]. The INT performance telemetry architecture is shown in Figure 6.

4.2 Functional Telemetry: Applications and Hyperventilation

The use of INT for functional purposes including microburst detection, problem-solving, congestion management, routing choices and traffic engineering [33,34,35,36,37,38,39,40] constitutes the most advanced implementation of the technology yet also its major gap between what is being promised and achieved.

Microburst detection [33]: Sub-second traffic spikes identification is one of INT's true strengths [33], as no other sampling technique is able to detect them [3]. Our research on TPP [33] and its applications reveals that INT is capable of detecting microbursts with a sub-millisecond precision [37,38] and therefore offers information which would not be available otherwise [19].

Nevertheless, critical evaluation demonstrates that the identification of microbursts alone is just the starting point, while there are major uncertainties surrounding further action [38]. The reaction actions triggered by microburst detections include rate limiting, changing load balancing and prioritizing packets [40], actions which cannot necessarily take place in timescales matching those of microburst detection latency [75].

CONGESTION CONTROL [39]: CONGESTION CONTROL [39]: The application of INT for congestion control (e.g., High Precision Congestion Control (HPCC)[39]) is a more developed implementation. It allows making instant rate updates by providing endpoints with information about queuing in data packets [75].

In our critical review, we highlight two drawbacks that have not been solved by HPCC [39] and other similar implementations. First of all, the network

state provided in INT packets becomes outdated by several microseconds even at high-speed networks [62]. While it is better compared to RTT-based systems [75], it still is not the instant one [39]. The second issue associated with INT-based congestion control is that it requires per-flow state tracking, which implies that switch memory resources used for storing information grow proportionally to the number of flows [80]. It may become an issue for high-scale environments [59].

TRAFFIC ENGINEERING AND ROUTING DECISIONS [40,61]: TRAFFIC ENGINEERING AND ROUTING DECISIONS [40,61]: The idea of applying INT for segment routing and traffic engineering [61][71][72] is presented by CLOVE [40] and FDSR-INT [61]. This approach allows making optimal decisions regarding routing by collecting specific telemetric information from network nodes [76].

As can be seen from our discussion, though the data collection process by INT is path-aware [61], routing optimization is rather coarse-grained [72]. The path information gathered by INT packets is path-aware for the particular path taken by such packets, while traffic engineering on the flow basis combines flows with diverse path characteristics in general [71]. Moreover, routing optimizations proposed in the literature rely heavily on the control plane coordination which may not be easily achieved in highly dynamic multi-domain environments [1,48,81].

4.3 Evaluation of the Experimentation Evidence Critique

We identify a recurring design trend in experimentation associated with the INT application literature in which capabilities are exaggerated while limitations are downplayed [19]:

Environment is Homogeneous [59]: Almost all studies associated with INT evaluations are carried out in an environment that is homogenous and controlled, which is comprised by the datacenter topology where only a few similar programmable switches exist [62,64,65]. These types of experiments lack the representation of production network environments [61].

Workloads are Synthetic [60]: Synthetic workloads used in evaluating the performance often disregard the burstiness, distribution of the flow size as well as the diversity of applications that may be found in real networks [69,70]. Such exceptional performance under ideal circumstances will most likely not translate into production scenarios [65].

Lack of Production Implementations [67]: There are very few documented examples of deployment of INT technologies in production networks despite being developed for the last eight years [19,67]. In other words, there is a gap of knowledge since no challenges of production implementations are discussed within the literature [62].

Partial Publication [62]: In experiments, the reported benefits are always based on averaging without any statistics about distribution, tails, or worst-case results [64]. As network management is always interested in tail results (99.9-th percentile), such partial publication will always create a systematic bias, which will always lead to overestimated gains from INT technologies [65].

5. CRITICAL ANALYSIS OF INT DRAWBACKS AND OPEN ISSUES

5.1 The Scalability Wall: Fundamental Boundaries of In-Band Telemetry

The analysis of all the existing work done about INT scalability shows the existence of something that we call “The INT Scalability Wall” [62], which represents a fundamental limitation due to the fact that the telemetry overhead grows non-linearly with network diameter.

INT Analytical Model [62]: Given a network with diameter D (highest number of hops between nodes), for a packet that travels the entire diameter of the network, the amount of metadata added by INT per hop would be m bytes [54]. The overhead incurred is $O(D \times m)$ per packet, where the number of packets for which telemetry needs to be generated depends on the total amount of traffic [19]. If a network contains N flows with an average packet size of P , then the overhead of INT becomes $O(N \times D \times m / P)$ which can become a large percentage of the bandwidth available in denser monitoring cases [62].

The empirical findings from the literature further verify this theoretical scaling [62,64]. For instance, TT-INT [62] achieves overhead of 3.4 bytes per packet with 300 packets per second (a reduction by 97.1% compared to the base case P4-INT [79]). Unfortunately, such an empirical finding also reveals that there is a cost in terms of temporal and spatial visibility as the visibility happens via time-threshold triggering, but not necessarily for each packet [62]. In situations that require dense visibility (for which INT was created), the overhead will approach the baseline one as 97.1% reduction will be done at the expense of collecting less telemetry [79].

MTU Constraint [55]: One significant yet under-emphasized constraint on the scalability of INT is the Maximum Transmission Unit (MTU) Limitation [19]. Given the limit of a 1500-byte Ethernet MTU, the INT hop taking up 44 bytes [54] implies that no more than 34 hops can happen before fragmentation [55]. This might be adequate enough in data centers; however, this severely limits the applicability of INT in large wide area networks and intercontinental links.

The Paradox of the Programmable Data Plane [27,52]: While allowing the introduction of INT, the Programmable Data Plane (PDP) [27] creates a paradoxical situation due to the processing limitations caused by INT [52], which consume the same ASIC resources as the packet processing [80]. In particular, our investigation of Tofino-based approaches [68] shows that 30% of pipeline stages may be taken by INT processing operations [52], hence leaving less resources for other packet-processing capabilities [27]. Hence, the demand for telemetry capability conflicts with high throughput packet forwarding [80].

5.2 Depletion of Hardware Resources: Memory & Processing Linkage

The most essential practical limitation that comes from switch ASIC hardware restrictions is often overlooked in the literature [80] in favor of algorithmic solutions [19]. Conflict between TCAM/SRAM and INT [80]: In order to perform INT-related telemetry functions, switches have to store the state of individual flows in order to decide on which packets to measure the desired telemetry metrics and where to add/extract metadata into/from headers [13]. This state resides in TCAM (Ternary Content Addressable Memory) and SRAM [74] in proportion to the number of flow entries [13]. If there are millions of flow states to maintain, they will surpass the amount of memory available [59], necessitating sampling of the network flows or upgrading of hardware [64].

Limitations of Processing Pipelines [52]: In order to add INT header to a packet, it entails processing at each node, such as the parsing of the header, generation of metadata, checksum calculation and creation of header itself [54]. In a high-speed switch, which handles several hundred million packets per second [68], processing required for INT, even with optimization [62], will lower down the forwarding capacity of the switch [79]. According to SPRINT [67], only about 10% of the on-chip SRAM is used for maintaining error correction state [74], but processing overhead is

unclear [19]. Figure 8 captures the hardware constraint of the problem.

Hardware-Software Co-Design Gaps [80]: The present state of INT literature mainly concentrates on the software aspect [62,64,65] and views hardware as a fixed parameter to be optimized [52]. We believe there is ample room to investigate hardware-software co-designs that move telemetry computation to custom-made coprocessors or modifications of the existing ASIC that will add telemetry acceleration [68]. However, this area is completely underrepresented within the literature, showing a gap between academia and industry where hardware vendors are leading the way on hardware innovations in the field of INT [19].

5.3 Security and Privacy Risks: The Untold Threat

Security risks posed by INT are probably the biggest gap between the rosy picture painted by literature and actualities [22]. As a result of our critical analysis, we have uncovered three major risks related to the use of INT technology:

Topology Disclosure [22]: By means of switch ID and path information included in a packet header, INT reveals topology to anybody who observes packets traversing the path [21]. This could be okay for single administrative domain data center networks [19].

In contrast, in multi-tenant cloud settings, inter-provider peering, or even in hybrid clouds [61], the exposure of this topological information provides substantial intelligence in adversarial reconnaissance [22]. Adversaries can learn about network topology, critical infrastructure, and targeted network segments [21].

INT Flow Pattern Exposure [22]: Even without performing deep inspection on packet payloads, INT metadata gives away traffic volume and latency information [8], as well as queuing conditions [22], which in turn can be used for performing traffic analysis, side-channel attack, and exhausting resources [22]. The adversary can learn the peak hours of traffic by analyzing INT messages, identify types of applications using latency information [28], and even get hold of valuable information based on timing analysis [39].

Metadata Manipulation and Spoofing [22]: The absence of any integrity protection of INT metadata – a flaw intrinsic to the specifications of INT technology itself [54] – makes it possible for an adversary to spoof false telemetry data in order to misinform network about the real state of things

leading to incorrect network management operations [21]. Moreover, an adversary can alter or drop INT headers or timestamps altogether [22].

It is evident that only 12% of papers surveyed take security into account [19]; all of which do not provide comprehensive security analyses [22]. Such security limitations constitute an existential threat for the adoption of INT in non-trusted domains [21].

5.4 The Deficiency in Standardization and Interoperability

The fragmentation of the INT eco-system, as identified in our taxonomy [19], constitutes a problem in terms of its interoperability [59]. The multitude of different flavors of INT, including P4-INT [54], IOAM [22], Sel-INT [13], PINT [66], iFIT [24], and others leads to poor interoperability between networks in terms of telemetry exchange [19]. Furthermore, fragmentation means not simply that there is no single standard for INT, but that there are several fundamental differences in approaches towards collecting such data, including the placement of data within the packet (PBT [23] and INT [19]) and the choice of probability versus determinism of the collection and the mandatory metadata fields [13]. Figure 10 shows the INT ecosystem fragmentation.

A unified INT design model is especially required in the future networking architectures of multi-cloud infrastructures, edge-cloud integration, and 6G networks covering multiple domains [63,71]. Deployment of INT in such scenarios would need the universal implementation of a particular INT model—a highly improbable scenario due to current standardization efforts—or complex translation devices [72].

6. DISCUSSION: IMPLICATIONS & FUTURE DIRECTIONS

6.1 Revising the INT Paradigm from Exhaustive Telemetry to Cognitive Telemetry

The critical evaluation necessitates a major overhaul of the INT philosophy. The key assumption prevalent among research on INT is the idea that the more telemetry is the better [62,64,65]; and that the purpose of using INB telemetry is to gather as many and finest telemetry information as possible [62,64,65]. These assumptions have given rise to scalability issues, security risks, and additional overhead that we have discussed above [59,80,22]. This research calls for a cognitive approach towards telemetry that we refer to as Cognitive In-Band Telemetry (C-INT).

Key Features of C-INT [60, 69, 70]:

1. Context Awareness [69]: The selection of telemetry techniques adapts to the network state collecting additional information at anomalous moments and decreasing the load on collection at times when the system works normally depending on congestion [60].
2. Objective-oriented Collection Process [70]: The collection of telemetry information is objective-driven and based on concrete needs and goals [39]. For example, if the goal is the detection of congestion, then collect queue, throughput, and not all information about latencies [22]. If the goal is security monitoring, collect information about path and time-stamp [70], not all the metrics.
3. Resource Consumption-oriented Approach [80]: C-INT constantly monitors the telemetry consumption of resources (bandwidth, ASIC cycles, memory) and adjusts the intensity of collection in order to prevent the exhaustion of resources in order to provide necessary visibility for needed goals [68].
4. Adaptation through Learning [69,70]: C-INT uses machine learning [69] and anomaly detection mechanisms [37] in order to detect the relevant telemetry components for current network state estimation [60].

The Cognitive In-Band Telemetry Architecture is presented in Figure 11.




Figure 11: Cognitive In-Band Telemetry (C-INT) Architecture

Research Challenges for C-INT [60,69,70]:

The changeover to cognitive telemetry raises several basic problems. State estimation algorithms [69], which can determine conditions of the network based on incomplete telemetry, should be created first. Next come reinforcement learning solutions [70] with minimal overhead that run in line-rate to make decisions about telemetry [60]. Finally, management abstraction layers [19] are needed, which will let operators specify high-level telemetry tasks without any low-level

configuration, resulting in adaptive telemetry, which does not require continuous intervention from operators [13]. Finally, C-INT must contain security provisions [22], which would eliminate the weaknesses introduced by current INT implementations [21], including telemetry data encryption and integrity protection.

6.2 Telemetry Models for Heterogeneous Environment

In our opinion, the lack of telemetry models suitable for operation in heterogeneous environments [19] represents a major gap in the current literature. Further studies need to address the problem of definition of telemetry requirements, representation, and methods of telemetry information collection in heterogeneous environments [59,61,63].

Telemetry Abstraction Layer [13,14]: One promising area of research is that of the creation of a Telemetry Abstraction Layer, which will separate the definition of telemetry requirements from the specific implementation of INT that can satisfy those requirements [13]. Such abstraction will allow the operator to define what telemetry is needed in terms of performance metrics, coverage requirements and fidelity levels without having to define the exact means of collecting the telemetry [14]. The abstract will then map the needs to the specific INT configuration [18], allowing seamless telemetry over domain and hardware platform [19].

Cross-Domain Telemetry Integration [61,72]: The issue of multi-domain telemetry collection and correlation is one of privacy-preserved telemetry, secure aggregation and domain-level trust management [22]. New research into the use of secure multi-party computation technologies for telemetry will help achieve telemetry without revealing the sensitive network information [21].

6.3 Hardware and Software Co-Design for INT Acceleration

From our analysis of hardware limitations [80], it is clear that hardware and software research should be more tightly coupled in the future [68]. As opposed to the current literature where hardware acts as a limiting factor [52], we contend that hardware-software co-design can result in tailored telemetry accelerators that will help reduce the processing cost of metadata insertion along with making more complex telemetry operations possible [27].

Potential hardware innovations [68, 74, 80]:

1. Telemetry Coprocessors [80]: Hardware components designed to handle the task of INT metadata creation, insertion, and validation without

occupying any of the pipeline stages [52] and integrated within switch ASICs [68].

2. Hardware-compression Engines [64,68]: Compressive hardware components capable of lossless or lossy [64] compression of the telemetry metadata, thus overcoming the bandwidth overhead problem without engaging CPUs [74].

3. Secure Enclaves for Telemetry [22]: Secure environments for telemetry processing allowing to protect the metadata from the risks related to the network untrustworthiness [21].

4. Adaptive Resource Allocation [80]: Specialized hardware allowing to allocate switch's resources adaptively in order to process telemetry data based on existing network load.

The academic community needs to establish closer cooperation with the hardware companies in order to consider these options [68].

6.4 AI/ML Integration: Opportunities and Challenges

Integration of AI/ML technologies into INT solutions is considered one of the promising research directions [60,69,70]; however, according to our critical review of research studies, AI/ML integration with INTs is still in its infancy [60], applying AI techniques mostly to the telemetry data rather than to telemetry itself [19].

AI for Telemetry Data Processing [37,60]: Applying ML to identify anomalies [37], patterns [69], and predict future events based on telemetry data is the most developed research area [60]. Nevertheless, there are some challenges identified during our critical analysis of researches; among them there are such as requirements of ML models to data [69], difficulties of preserving model accuracy under changing network environment [70] and problems with false negative and false positive results that can decrease operator's confidence level [37].

AI for Telemetry Optimization [60]: Using of reinforcement learning to optimize the process of telemetry collection (which packets are being instrumented, what type of metadata should be collected and how the strategy should be changed) [70] is the more interesting research area [69] with high research potential. For example, NTP-INT uses Deep Reinforcement Learning for the probe path planning. However, it is uncertain whether those schemes scale for the millions of flows seen in high speed networks [59].

AI-based Telemetry Models [69]: One promising area of research is in building generative models

based on AI technologies capable of predicting unseen network states using incomplete telemetry data [69] that will help in dealing with the visibility challenges described above [62]. Such models learn the correlation between different network state variables [69] and thus are able to fill gaps caused by lack of certain metadata due to resource restrictions [68]. A hard problem, which has not yet been solved, is the accuracy of reconstruction.

6.5 Critique of This Work: Limitations and Self-Reflection

It goes without saying that our critical assessment has its own limitations just like the papers we are analyzing [19, 59, 62, 64, 65]. First of all, our classification of INT schemes within the VFOS paradigm is helpful but inevitably simplifying in terms of implementation details [62]. The categories we use – Visibility, Fidelity, Overhead and Security – are not entirely independent in terms of impact on one another [64].

Second, our assessment of deployment readiness and practical limitations is based primarily on reported results and our interpretation of those results [62,64,65]. The absence of comprehensive production deployment data [67]—itself a limitation of the literature [19]—prevents definitive conclusions about INT's real-world performance [60]. Our analysis is thus based on the best available evidence, but we acknowledge that undisclosed deployment experiences might significantly modify our assessment [62].

Third, our critique of INT's limitations may underestimate the pace of hardware evolution [68]. The current ASIC constraints we identify [80] might be mitigated by next-generation programmable switches with expanded memory and processing capacity [52]. Our analysis reflects the state of the field at the time of writing, and we acknowledge the potential for hardware advances to alter the trade-offs we describe [19].

Finally, our recommendation for cognitive, adaptive telemetry [60] represents an aspirational vision that introduces its own research challenges [69,70]. The implementation of adaptive telemetry strategies that can balance competing objectives and respond to rapidly changing network conditions is itself a significant research challenge [60]; our proposal does not solve this challenge but identifies it as a critical research direction [19].

However, developing adaptive telemetry schemes capable of meeting conflicting goals in dynamic environments constitutes a challenging scientific task on its own [60]; our study doesn't solve the

problem, but highlights its importance as a future research agenda [19].

7. DIFFERENCE FROM PRIOR WORKS

The current survey of in-band network telemetry stands out from prior INT surveys and reviews [19] in its critical and analytical approach, along with emphasis on limitations and unresolved issues [59,62,64,65]. For a better illustration of our contribution, we compare our paper with prior influential surveys [19, 67, 79]:

7.1 Comparative Discussion of Tan et al., 2021 [19] "A survey of In-band Network Telemetry" by Tan et al. [19] constitutes one of the best reviews of the field, including architectural principles, applications, and challenges of in-band network telemetry technology [13,14,15,16,17,18]. The main difference between Tan et al. [19] and our survey consists in the following:

A Descriptive Approach vs. Analytical: The survey by Tan et al. provides a very useful and descriptive overview of INT techniques, but lacks analytical evaluation of the discussed approaches and technologies [13, 14, and 15].

The critique offered is concerned with questioning the arguments raised by the supporters of INT [62, 64, 65] and uncovering the overlooked trade-offs which the descriptive body of literature overlooks [59].

Fragmentary Approach vs. Unified Approach: Previous analysis [19] views the problems facing cloud computing as individual concerns related to overhead, security, and scalability without making connections between these problems and key design decisions [13,14,15]. Our VFOS framework enables us to provide a unified view that shows how these challenges are connected [62, 64,65], where advancements in one aspect result in problems in other aspects [59].

Research Directions vs. Research Agenda: Although previous analysis [19] offers certain research directions, those directions only involve further elaboration on the same concepts rather than paradigm shifts [13, 14, 15, 16, 17, 18]. Our research agenda requires new paradigms such as cognitive telemetry [60], hardware-software co-design [80], and cross-domain interoperability [61, 72].

7.2 Comparison with Papadopoulos et al. (2023) [79]

The work by Papadopoulos et al., entitled "Deterministic and Probabilistic P4-Enabled Lightweight In-Band Network Telemetry" [79], provides useful information about overhead reduction [66, 68]; however, the comparison of our analysis and the paper by Papadopoulos et al. [79] indicates the following differences in scope and critical perspectives:

Scope: The study by Papadopoulos et al. [79] considers only light-weight INT solutions [66,68] in a comprehensive analysis of deterministic and probabilistic approaches [62]; our comparative analysis involves both light-weight INT approaches [13,14,19,22,64,65] and their heavy-weight counterparts that emphasize visibility over efficiency [39], thus making it possible to provide more general insight into the design space [59].

Critical Perspective: Although Papadopoulos et al. [79] point out the deficiencies of probabilistic approaches [66], the comparison between the two presented approaches is conducted within the common design space [62] in which one solution replaces the other smoothly; our critical analysis of this comparison shows that the design space is rather an illusion since the purpose of the two approaches is quite different [19].

Considerations for Hardware and Security [22,80]: Lightweight INT papers [66,68,79] often consider bandwidth and overhead concerns, but they offer little discussion of hardware considerations [80] and security implications [22]. Our investigation shows that such factors are essential limitations that undermine the practicality of lightweight systems [59, 80].

7.3 Comparison with Bae et al. (2026) [64] INTpress by Bae et al. [64] is the most recent advance in INT overhead reduction using compression technique [68]. Although there is some merit in the work, our critical assessment indicates a number of limitations that are overlooked in their experimental investigation [62,65]:

Quantization Error [64]: The INTpress uses a compression scheme based on quantization that sacrifices precision for speed [64]. According to our investigation, the quantization error is not consistent between metrics; the error in queue depth quantization can be different from the error in latency quantization [39].

Overhead of Dynamic Partitioning [64]: While the dynamic data partitioning technique responsible for achieving INTpress's improved performance [64] involves a corresponding overhead based on traffic

variability [59], our analysis reveals that the overhead associated with managing such partitions can be a relevant factor under varying traffic conditions [60].

Generic Data Compression vs. Content-based Data Compression [62,64]: The use of a generic data compression algorithm in all metadata field types [64] is another example we have identified during our analysis. In this case, using a more flexible content-based data compression algorithm [62] would lead to better results [65].

7.4 OUR CONTRIBUTION

Having analyzed the existing research on comparative perspectives, we formulate our key contribution [19]:

1. **Critical Analytic Framework [19]:** We devise the VFOS framework - the first holistic analytic framework for analyzing INT implementations with respect to their fundamental trade-offs rather than claimed benefits [62,64,65].
2. **Empirical Trade-off Characterization [59, 62]:** We consolidate findings from experiments in the literature [62, 64, 65] to characterize trade-offs in INT approaches [59] uncovering the non-linear trade-off relations and scalability bottlenecks hidden behind results of individual studies [62].
3. **Paradigm Shift Idea [60,69,70]:** We define the idea of C-INT or Cognitive In-Band Telemetry as an alternative paradigm shift in the INT trajectory [60], moving from exhaustive telemetry to goal-driven adaptive approaches to INT [69,70].
4. **Comprehensive Critical Analysis [19, 22, 59, 62, 64, 65]:** We conduct the first comprehensive review of limitations of INT across all dimensions – scalability [59, 62], hardware [80], security [22], and interoperability [61, 72] specifying research issues that need to be addressed to deploy INT.
5. **Prospective Research Agenda [60, 69, 70, 80]:** Based on our critical analysis, we formulate the future research agenda for INT [59, 62] setting the directions and methods that will enable more impactful research on INT in the future.

8. CONCLUSION: ADVANCING THE SCIENCE OF NETWORK TELEMETRY

8.1 Overview of Results

Our critical assessment of In-Band Network Telemetry highlights a domain whose potential is immense but plagued by inherent limitations that render current technologies infeasible for deployment [19]. The researchers in the INT community have achieved much success in demonstrating what in-band telemetry can do

[8,9,10], crafting ingenious realizations of INT using programmable switching fabrics [52,54], and utilizing it to perform myriad different functions ranging from congestion control to anomaly detection [13,14,15,16,17,18,33,34,35,36,37,38,39,40]. However, as this field matures and gains in practicality, the scalability, hardware, security and interoperability issues have received scant consideration [19, 22, 59, 62, 64, 65, and 80].

Our systematic survey of 86 primary sources and development of the Visibility-Fidelity-Overhead-Security (VFOS) framework [19] has helped us recognize some key limitations that have been grossly underestimated in the literature [62,64,65]:

First, the scalability issue called overhead scalability wall [62] – the non-linear scaling of the overhead with respect to the diameter of the network architecture. The MTU limit on packet growth [55], combined with per-hop metadata insertion [54], creates a practical upper bound on the network size and flow count that can be monitored comprehensively [62].

Second, the hardware resource constraints of programmable switches [80] limited TCAM, SRAM, and ASIC processing cycles [52]—create competition between telemetry processing and packet forwarding that reduces effective network throughput [68]. Despite advances in switch technology [74], the resource demands of exhaustive INT collection exceed available capacities in high-scale environments [59].

Third, the security vulnerabilities introduced by INT [22] topology exposure, flow pattern leakage, and metadata manipulation [21] represent an existential threat to deployment in multi-tenant, multi-domain, or security-sensitive environments [61]. The near-absence of security mechanisms in current INT implementations is a serious gap requiring urgent research attention [22].

Fourth, the fragmentation of the INT ecosystem into multiple, often-incompatible variants [13, 14, 19, 22, 64, 65, 66, and 68] creates interoperability challenges that impede adoption in heterogeneous networks [59]. The absence of a unified telemetry model for cross-domain environments [61,72] limits INT's applicability beyond the most controlled deployments [19].

8.2 Scientific Contribution and Knowledge Advancement

The proposed paper provides many valuable contributions to existing scientific knowledge about network telemetry [19]:

1. Critical Synthesis and Framework Development [19]: For the first time, we perform a comprehensive critical synthesis of INT literature [59,62,64,65]. Contrary to descriptive categorization provided elsewhere [59,62,64,65], our VFOS framework gives researchers and practitioners tools for evaluating different INT solutions, revealing their hidden drawbacks and limitations behind excessive emphasis on performance [62,64,65].

2. Trade-off Analysis [59,62]: Using our synthesis and comparison of results obtained in different experiments [62,64,65], we uncover practical trade-offs involved into use of various INT approaches [13,14,66,68]. Our findings show that overheads claimed by some works [62] have an adverse effect on visibility [62] and do not necessarily hold up in practice [59].

3. Paradigm Shift [60,69,70]: Our paper formulates Cognitive In-Band Telemetry (C-INT), which is essentially a new approach to INT research direction [60]. Contrary to exhaustive maximalistic collection proposed in previous papers [62], it involves flexible adaptation of telemetry mechanisms according to current network and application environment [69,70]. Moving away from optimization of the amount to optimization of intelligence is the contribution made to the agenda of the research field [19].

4. Identification of Research Gaps [22, 59, 62, 64, 65, 80]: The results of our critical analysis make clear which research gaps must be addressed in order for INT to be widely deployed: adaptive telemetry algorithms [60], co-design of hardware and software [80], secure mechanisms [22], and telemetry abstraction layers [13,14]. Our research provides guidance for future research going beyond incrementally optimizing [62].

5. Methodological Contribution [19]: The method we applied—a systematic analysis of the trade-offs involved, implicit assumptions made, and limitations unacknowledged—is an example of a methodology suitable for the critical examination of new networking technology beyond INT [1,2,3].

8.3 Research and Practitioner Implications

There are significant research and practice implications of our work for INT [19]:

For Researchers [59, 62, 64, 65]: It is important for the research community to start addressing the challenges we described and go beyond incrementally optimizing [22, 59, 62, 80].

For Practitioners [19]: Our analysis points to the necessity to think carefully through the trade-offs of

visibility, fidelity, overhead, and security before using INT in actual deployments [62, 64, 65]. Depending on particular requirements and risk tolerance, organizations may select INT variants accordingly [13, 14, 66]; the best option in terms of controlled data centers can be unsuitable for multi-tenant clouds or inter-domain setups [22,61].

For Standardization Bodies [22,61]: Our analysis highlights the importance of standardizing interoperable, secure, and cross-domain telemetry solutions [19]. Fragmentation of the INT ecosystem [13,14,19,22,64,65,66,68] calls for standardization efforts to take into account both flexibility and interoperability, as well as security considerations from the start.

8.4 Concluding Remarks

The path of In-Band Network Telemetry, from novel idea [19] to real-world implementation, has much more to teach us than just one example of technology deployment in complex systems [1]. The revolutionary possibility of per-packet and per-hop monitoring [8,9,10], as part of network management needs [3], may look very attractive. Yet, its realization entails not only innovations but also insights into the constraints that such technologies face [80], security issues [22] and heterogeneity challenges [59].

As we look to the future, we see the potential for INT and related in-band telemetry approaches to play a significant role in network management [19,63], particularly as programmable data planes continue to evolve [27,52] and as the management requirements of 6G, cloud-native, and AI-driven networks become more demanding [63,69,70]. However, this potential will only be realized if the research community addresses the challenges we have identified with the same creativity and rigor that has characterized INT development to date [59, 62, 64, 65].

The transition from exhaustive to cognitive telemetry [60]—from collecting everything to collecting what matters, when it matters [69,70]—represents the next frontier for network telemetry research [19]. The development of adaptive, learning-enhanced telemetry strategies [60] that can deliver the visibility needed for effective network management without overburdening the network infrastructure is the central challenge that will define the next decade of research in this field [59,62,64,65].

In conclusion, we offer not a rejection of In-Band Network Telemetry but a call for more thoughtful, critical, and sustainable research [19]. The

fundamental insight of INT—that network packets can be carriers of network intelligence [8,9,10]—remains valid and valuable. The challenge is to develop telemetry architectures that realize this insight in a way that respects the physical constraints of networks [80], addresses the security requirements of multi-domain environments [22], and provides the interoperability needed for heterogeneous deployments [59]. We believe that the cognitive telemetry paradigm outlined in this paper [60] provides a roadmap for achieving these goals. This problem of creating telemetry infrastructures which take this idea into account in ways consistent with the laws of physics [80], consistent with security requirements [22], and consistent with interoperability needs in heterogeneous environments [59], lies ahead of us. In our opinion, the framework of cognitive telemetry that has been proposed in this paper [60] gives a good direction for resolving this problem.

REFERENCES

- [1]. Nunes BAA, Mendonca M, Nguyen XN, Obraczka K, Turetti T. A survey of software-defined networking: Past, present, and future of programmable networks. *IEEE Communications Surveys & Tutorials*. 2014;16(3):1617-1634.
- [2]. Morton A. Active and passive metrics and methods (with hybrid types in-between). *RFC 7799*. 2016.
- [3]. Landau-Feibish S, Liu Z, Rexford J. Compact data structures for network telemetry. *ACM Computing Surveys*. 2025;57(8):1-31.
- [4]. Claise B, ed. *RFC 3954: Cisco systems netflow services export version 9*. 2004.
- [5]. Wang M, Li B, Li Z. sFlow: Towards resource-efficient and agile service federation in service overlay networks. In: *24th International Conference on Distributed Computing Systems*, 2004. *Proceedings*. IEEE; 2004.
- [6]. Aitken P. *RFC 7011: Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*. 2013.
- [7]. Quittek J, Zseby T, Claise B, Zander S. *RFC 3917: Requirements for IP flow information export (IPFIX)*. 2004.
- [8]. Van Tu N, Hyun J, Park T, Hong JWK. Intellector: A high-performance collector for in-band network telemetry. In: *2018 14th International Conference on Network and Service Management (CNSM)*. IEEE; 2018.

- [9]. Manzanares LP, et al. An INT-based packet loss monitoring system for data center networks implementing Fine-Grained Multi-Path routing. *Internet Technology Letters*. 2023;6(6):e393.
- [10]. Walrand J. Probability in electrical engineering and computer science: an application-driven course. 2021.
- [11]. Quittek J, White K. Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations. RFC 4560. 2006.
- [12]. Dreibholz T, Mazumdar S. Tracing Data Packet Paths over the Internet using Traceroute. *arXiv preprint arXiv:2602.09857*. 2026.
- [13]. Tang S, Li D, Niu B, Peng J, Zhu Z. Sel-INT: A runtime-programmable selective in-band network telemetry system. *IEEE Transactions on Network and Service Management*. 2019;17(2):708-721.
- [14]. Song E, Pan T, Jia C, et al. INT-label: Lightweight in-band network-wide telemetry via distributed labeling. *IEEE Transactions on Parallel and Distributed Systems*. 2024;35(5):751-767.
- [15]. Pan T, Song E, Jia C, et al. Int-path: Towards optimal path planning for in-band network-wide telemetry. In: *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE; 2019.
- [16]. Bhamare D, et al. Intopt: In-band network telemetry optimization for nfv service chain monitoring. In: *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE; 2019.
- [17]. Lin Y, et al. Netview: Towards on-demand network-wide telemetry in the data center. *Computer Networks*. 2020;180:107386.
- [18]. Tang S, et al. How to use in-band network telemetry wisely: Network-wise orchestration of Sel-INT. *IEEE/ACM Transactions on Networking*. 2022;31(1):421-435.
- [19]. Tan L, Su W, Zhang W, et al. In-band network telemetry: A survey. *Computer Networks*. 2021;186:107763.
- [20]. Zhang Y, et al. VXLAN-based INT: In-band network telemetry for overlay network monitoring. In: *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops*. IEEE; 2021.
- [21]. Mizrahi T, et al. An overview of operations, administration, and maintenance (OAM) tools. 2014.
- [22]. Brockners F, Bhandari S, Mizrahi T. Data fields for in situ operations, administration, and maintenance (IOAM). RFC 9197. 2022:1-40.
- [23]. Song H, et al. Postcard-based on-Path flow data telemetry. Technical Report. 2021.
- [24]. Lu B, et al. iFIT: Intelligent flow information telemetry. In: *Proceedings of the ACM SIGCOMM 2019 Conference Posters and Demos*. 2019.
- [25]. Capello A, et al. RFC 8321: Alternate-Marking Method for Passive and Hybrid Performance Monitoring. 2018.
- [26]. Stallings W. SNMP, SNMPv2, SNMPv3, and RMON 1 and 2. Addison-Wesley Longman Publishing Co., Inc.; 1998.
- [27]. Bifulco R, Rétvári G. A survey on the programmable data plane: Abstractions, architectures, and open problems. In: *2018 IEEE 19th International Conference on High Performance Switching and Routing (HPSR)*. IEEE; 2018.
- [28]. Brockners F, et al. Requirements for in-situ OAM. Working Draft, Internet-Draft draft-brockners-inband-oam-requirements-03. 2017.
- [29]. Van Tu N, Hyun J, Hong JWK. Towards ONOS-based SDN monitoring using in-band network telemetry. In: *2017 19th Asia-Pacific Network Operations and Management Symposium (APNOMS)*. IEEE; 2017.
- [30]. Geng Y, et al. SIMON: A simple and scalable method for sensing, inference and measurement in data center networks. In: *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. 2019.
- [31]. Hohenberger R, et al. Optimizing distributed network monitoring for NFV service chains. *IEEE Communications Letters*. 2019;23(8):1332-1336.
- [32]. Liang J, et al. In-band network function telemetry. In: *Proceedings of the ACM SIGCOMM 2018 Conference on Posters and Demos*. 2018.
- [33]. Jeyakumar V, et al. Millions of little minions: Using packets for low latency network programming and visibility. *ACM*

- SIGCOMM Computer Communication Review. 2014;44(4):3-14.
- [34]. PRoML-INT: Comprehensive telemetry data collection from the optical packet network. [Year].
- [35]. Hyun J, et al. ONOS+BMv2 framework for knowledge-based network. [Year].
- [36]. NetworkAI: Deep reinforcement learning for network control. [Year].
- [37]. Hohenberger R, et al. Machine learning for telemetry implementation. [Year].
- [38]. Jia, et al. Gray failure detection using INT. [Year].
- [39]. Li Y, et al. HPCC: High precision congestion control. In: Proceedings of the ACM Special Interest Group on Data Communication. 2019. p. 44-58.
- [40]. Katta, Naga, et al. "Clove: Congestion-awareload balancing at the virtual edge." Proceedings of the 13th International Conference on emergingNetworking Experiments and Technologies. 2017.
- [41]. Liu, Ziyuan, et al. "HyperDrive: DirectNetwork Telemetry Storage via ProgrammableSwitches." IEEE Transactions on CloudComputing (2025).
- [42]. Open Network Operating System (ONOS). <https://onosproject.org>
- [43]. Behavioral Model version 2 (BMv2). <https://github.com/p4lang/behavioral-model>
- [44]. Postel J. Transmission Control Protocol. RFC 793. 1981.
- [45]. Gao, Ya, and Zhenling Wang. "A review of P4programmable data planes for networksecurity." Mobile Information Systems 2021.1(2021): 1257046.
- [46]. Fielding R, et al. Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing. RFC 7230. 2014.
- [47]. Dukkipati, Nandita. Rate Control Protocol(RCP) congestion control to make flows completequickly. Stanford University, 2008.
- [48]. Nisar, Kashif, et al. "A survey on thearchitecture, application, and security of softwaredefined networking: Challenges and openissues." Internet of Things 12 (2020): 100289.
- [49]. Denazis, S., et al. "RFC 7426: Softwaredefinednetworking (SDN): Layers and architectureterminology." (2015).
- [50]. Michel, Oliver, et al. "The programmable dataplane: Abstractions, architectures, algorithms, andapplications." ACM Computing Surveys(CSUR) 54.4 (2021): 1-36.
- [51]. McKeown N, et al. OpenFlow: enabling innovation in campus networks. ACM SIGCOMM Computer Communication Review. 2008;38(2):69-74.
- [52]. Bosshart P, et al. P4: Programming protocol-independent packet processors. ACM SIGCOMM Computer Communication Review. 2014;44(3):87-95.
- [53]. Hanks S, et al. Generic Routing Encapsulation (GRE). RFC 1701. 1994.
- [54]. P4 Language Consortium. In-band Network Telemetry (INT) Specification. 2017.
- [55]. Postel, Jon. "RFC0791: Internet protocol."(1981).
- [56]. Bosshart, Pat, et al. "P4: Programmingprotocol-independent packet processors." ACMSIGCOMM Computer CommunicationReview 44.3 (2014): 87-95.
- [57]. The P4.org Applications Working Group.Contributions from Alibaba, Arista, CableLabs,Cisco Systems, Dell, Intel, Marvell, Netronome,VMware 2020-11-11
- [58]. Caetano, David, and Bruno Sousa. "P4-SCINT:Service Chaining with In-band NetworkTelemetry." 2025 IEEE Conference on NetworkFunction Virtualization and Software-DefinedNetworking (NFV-SDN). IEEE, 2025.
- [59]. Yuan Q, et al. INT-Partition: Hierarchical and Fault-Tolerant In-Band Network Telemetry. IEEE Transactions on Networking. 2025.
- [60]. Zhang P, et al. NTP-INT: Network traffic prediction-driven in-band network telemetry for high-load switches. Journal of Network and Computer Applications. 2025:104265.
- [61]. Xu J, et al. FDSR-INT: A Flexible On-Demand In-Band Telemetry Approach for Aerial Computing Networks. IEEE Internet of Things Journal. 2025.
- [62]. Zhang H, et al. TT-INT: A Time-Threshold-based Lightweight In-Band Network Telemetry Scheme for P4-Enabled

- Programmable Networks. IEEE Transactions on Network and Service Management. 2026.
- [63]. Fu M, et al. A learning-based service function chain early fault diagnosis mechanism based on in-band network telemetry. IEICE TRANSACTIONS on Information and Systems. 2022;105(2):344-354.
- [64]. Bae C, et al. INTpress: Compression-based Quantized in-band Network Telemetry for Low Bandwidth Overhead Monitoring. IEEE Transactions on Cloud Computing. 2026.
- [65]. Kim H, et al. FreshINT: Freshness-aware Early Reporting in In-band Network Telemetry Systems. In: IEEE INFOCOM 2026-IEEE Conference on Computer Communications. IEEE; 2026.
- [66]. Ben Basat R, et al. PINT: Probabilistic in-band network telemetry. In: Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication. 2020.
- [67]. Altaf T, et al. In-Band Network Telemetry for End-to-End Service Performance Monitoring in Broadband Networks: A Survey and Deployment Feasibility Analysis. IEEE Access. 2026.
- [68]. Sheng S, Huang Q, Lee PPC. DeltaINT: Toward general in-band network telemetry with extremely low bandwidth overhead. In: 2021 IEEE 29th International Conference on Network Protocols (ICNP). IEEE; 2021.
- [69]. Gao J, et al. MTGNN: multi-task graph neural network based few-shot learning for disease similarity measurement. Methods. 2022;198:88-95.
- [70]. Li K, Zhang T, Wang R. Deep reinforcement learning for multi objective optimization. IEEE Transactions on Cybernetics. 2020;51(6):3103-3114.
- [71]. Li Z, Hu Z, Li C. SRv6 network programming: ushering in a new era of IP networks. CRC Press; 2021.
- [72]. Liu J, et al. Selective and on-demand network measurement with SRv6 and INT. Computer Networks. 2023;234:109914.
- [73]. Li H, et al. Physical-layer In-band Network Telemetry for Wireless Backhauling Towards 6G. IEEE Transactions on Networking. 2026.
- [74]. Panda PR, Dutt ND, Nicolau A. On-chip vs. off-chip memory: the data partitioning problem in embedded processor-based systems. ACM Transactions on Design Automation of Electronic Systems (TODAES). 2000;5(3):682-704.
- [75]. Li Y, et al. HPCC: High precision congestion control. In: Proceedings of the ACM Special Interest Group on Data Communication. 2019. p. 44-58.
- [76]. Xu Z, Chen X, Zhu Z. INT-assisted Adaptive Packet Scheduling in PDP Switches for End-to-end Latency Control. In: IEEE INFOCOM 2025-IEEE Conference on Computer Communications. IEEE; 2025.
- [77]. Liu Y, et al. SFANT: A SRv6-based flexible and active network telemetry scheme in programming data plane. IEEE Transactions on Network Science and Engineering. 2023;11(3):2415-2425.
- [78]. Papadopoulos K, Papadimitriou P, Papagianni C. Deterministic and probabilistic p4-enabled lightweight in-band network telemetry. IEEE Transactions on Network and Service Management. 2023;20(4):4909-4922.
- [79]. Yu F, Katz RH, Lakshman TV. Gigabit rate packet pattern-matching using TCAM. In: Proceedings of the 12th IEEE International Conference on Network Protocols, 2004. ICNP 2004. IEEE; 2004.
- [80]. Benzekki K, El Fergougui A, Elbelrhiti Elalaoui A. Software-defined networking (SDN): a survey. Security and Communication Networks. 2016;9(1):101-125.
- [81]. [81] Benzekki, Kamal, Abdeslam El Fergougui, and Abdelbaki Elbelrhiti Elalaoui. "Software-defined networking (SDN): a survey." Security and communication networks 9.

Table 1: Comparative Analysis of INT Implementations via VFOS Framework

Implementation	Visibility	Fidelity	Overhead	Security	Deployment Readiness	Critical Limitations
P4-INT (Baseline) [54]	High (per-packet, per-hop)	High (exact timestamps, queue depths)	Very High (44B+/packet)	Low (metadata in plaintext)	Moderate (Tofino, BMv2)	No scaling beyond 10-15 hops; significant bandwidth overhead
PINT (Probabilistic) [66]	Low (1-16 bits/packet)	Medium (statistical approximations)	Very Low	Medium (partial obfuscation)	High (Tofino, P4)	Loss of per-packet visibility; reconstruction overhead
Sel-INT [13]	Medium (selective flows)	High (selected flows)	Medium (variable)	Low	Moderate (Tofino)	Orchestration complexity; path planning NP-hard
INT-Partition [59]	High (partition-based)	High	Medium	Low	Moderate (Tofino, BMv2)	Partition management overhead; fault tolerance complexity
TT-INT [62]	Medium (time-threshold based)	High (localized)	Low (3.4B/packet avg)	Low	High (Tofino)	Temporal threshold tuning; clock synchronization assumptions
FreshINT [65]	High (freshness-optimized)	High (time-sensitive)	Medium	Low	Moderate (Tofino)	Sink node proximity dependency; reporting path constraints
INTpress [64]	Medium (compressed)	Medium (quantized)	Very Low (66% reduction)	Low	Moderate (Tofino)	Quantization errors; dynamic partitioning complexity
SPRINT (FEC-based) [67]	Medium (reconstruction)	High (error-corrected)	Medium (XOR overhead)	Medium (encoding)	High (Tofino, BMv2)	State management; recovery latency
IOAM [22]	High (in-situ)	High	High	Low (standards-track)	Low (emerging)	Packet length limitations; processing overhead




Figure 3: Historical Timeline of INT Development




Figure 5: Passive vs. Active Network Telemetry Comparison




Figure 6: INT Performance Telemetry Architecture




Figure 7: INT Functional Telemetry Applications




Figure 8: INT Scalability Wall Analysis




Figure 8: Hardware Resource Constraints in INT Deployment




Figure 9: Security and Privacy Vulnerabilities in INT




Figure 10: INT Ecosystem Fragmentation Analysis