Advancing Network Monitoring and Operation with In-band Network Telemetry and Data Plane Programmability

Jonatas A. Marques¹, Luciano Paschoal Gaspary¹

¹Instituto de Informática – Universidade Federal do Rio Grande do Sul (UFRGS) Caixa Postal 15.064 – 91.501-970 – Porto Alegre – RS – Brazil

{jonatas.marques,paschoal}@inf.ufrgs.br

Abstract. Modern communication networks operate under high expectations on performance and resilience mainly due to the continuous proliferation of nonelastic highly-distributed applications. In this context, closely monitoring the state, behavior, and performance of networking devices and their traffic as well as quickly troubleshooting problems as they arise is essential for the operation of network infrastructures. In this thesis, we make several contributions — based on in-band network telemetry and data plane programmability — that advance the discipline of network monitoring and operation. We formalize telemetry orchestration problems, prove their NP-Completeness, and propose polynomial computing time heuristic to efficiently solve real instances of these problems. We also design a system that combines in-band telemetry and in-network computation to enable the highly accurate and fine-grained detection and diagnosis of service-level objective violations. Finally, we introduce an approach that is able to recover from network link and node failures at data-plane timescales via policy-optimal paths. We also discuss opportunities and challenges for adapting this approach for other time-sensitive network management tasks.

Resumo. Redes de comunicação modernas operam sob altas expectativas de desempenho e resiliência. Isto, principalmente em função da continua proliferação de aplicações inelásticas altamente distribuídas. Neste contexto, torna-se essencial para a operação de uma infraestrutura de rede a monitoração aguçada do estado, comportamento, e desempenho do tráfego e dos dispositivos de rede assim como a resolução ágil de quaisquer problems os afetem. Nesta tese, faz-se diversas contribuições — alicerçadas em telemetria de rede in-band e programabilidade do plano — para a disciplina de monitoração e operação de redes. Formaliza-se problemas de orquestração de ações de telemetria, provase o pertencimento destes à classe de problemas NP-Completos, e propõe-se heurísticas polinomiais capazes de resolver instâncias reais destes problemas em tempo hábil. Projeta-se, também, um sistema que combina telemetria inband e computação in-network para possibilitar a detecção e diagnóstico de violações de service-level objectives de forma altamente acurada e precisa. Finalmente, introduz-se uma abordagem para recuperação de falhas de dispositivos e enlaces de rede que re-roteia tráfego através de caminhos ótimos (segundo políticas de encaminhamento) na escala de tempo do plano de dados. Também discute-se oportunidades e desafios relacionados a adaptação da abordagem proposta para a realização de outras tarefas de gerência de redes que são sensíveis ao fator tempo.

1. Introduction

Current networks operate with high expectations on performance (e.g., latency, bandwidth, availability), especially with the emergence and proliferation of new applications (e.g., algorithmic trading, telesurgery, and virtual reality video streaming) with architectures based on many interconnected components spread across multiple end-points [Balakrishnan 2021]. These applications and their users demand strict requirements, which to be met require defining clear goals for network performance, the so-called service-level objectives (SLOs), and troubleshooting problems that may prevent achieving such goals. Unfortunately, there is a myriad of problems that may impact the correct and efficient operation of a network, ranging from traffic congestions all the way to hardware failures. In this context, monitoring the state, behavior, and performance of networking devices and their traffic is essential for the operation of today's network infrastructures. Nevertheless, network monitoring is an inherently hard task, sometimes compared to searching for a needle in a haystack.

Existing tools and techniques are not engineered to monitor networks and help troubleshoot their problems with the level of detail and accuracy required nowadays. For example, and regarding the collection of metadata and statistics, traditional passive monitoring tools operate at coarse timescales (dozens of seconds and up) and, thus, lack adequate granularity to detect events such as short-lived traffic bursts (e.g., microbursts) that may be critical to modern applications. Active measurement techniques also do not provide sufficient time resolution; additionally, there is no guarantee that the network will route and prioritize probes in the same way as production packets. As a second example, and regarding troubleshooting SLO violations, approaches based on packet mirroring can help give visibility into the network to understand how packets are being processed and forwarded by devices. These approaches find their main challenge in keeping the monitoring overhead (i.e., required bandwidth and processing) under reasonable levels while still collecting fine grained data. Packet sampling, their common method for addressing this challenge, inherently leads these techniques to miss important events; deciding what and when to sample is hard.

Another challenge in meeting SLOs in modern times is the common dependency and delay in communication between the mechanisms that detect problems and the ones that find the solution to these problems. Consider, for example, the case of equipment failures and their impact on network availability. Existing solutions depend on some type of computation in the control plane at the time of failure and subsequent reconfiguration of forwarding tables. Computing the new forwarding entries for devices or, in broader terms, the solution to the problem can take considerable time. As the delay in reacting to failures leads to a significant number of packet drops, in the general case, the delay caused by the use of long control loops to solve problems can lead to substantial performance and financial loss. Ideally, the data plane should be able to react immediately at the time of failure. We note that the limitations and drawbacks presented by the existing monitoring tools and techniques result from the low level of flexibility in defining how packets are to be processed by the data plane in traditional networks. Even with Software-Defined Networking (SDN) and OpenFlow [McKeown et al. 2008], there is only support for standardized protocols; there is no freedom to define customized protocols and procedures.

2. Problem Statement

As a result to the presented scenario, the networking community has sought for more flexibility in the data plane, which recently culminated in the proposal of data plane programmability (DPP) [Bosshart et al. 2013]. DPP reshapes the SDN landscape by enabling network operators to reprogram forwarding devices in-field to deploy novel networking protocols, customize the network behavior, and consequently develop and support innovative services and applications. Protocols and packet processing procedures, in this new paradigm, are defined via domain-specific languages - e.g., P4 [Bosshart et al. 2014] - with support for abstractions to specify customized protocol headers, parsing logic, and match-action tables, for example. One interesting concept that gained traction with the introduction of programmable data planes is In-band Network Telemetry (INT) [Kim et al. 2015]. Within this concept, forwarding devices are programmed to annotate production packets with metadata regarding their state, behavior, and performance (such as port utilization, matched forwarding entries, and queuing delays). The annotated information is accumulated in a packet along its path and, at some point in the network, extracted and reported to analyzer servers. These servers piece together the received information to build an accurate and global view of the network, as observed by its traffic.

INT-based techniques have shown to produce monitoring data with an unprecedented level of accuracy and fine granularity [Kim et al. 2015]. That is because instead of relying on active probes, which may be subject to forwarding and routing behaviors different from those of the traffic of interest, the production packets themselves can be used to probe the network. Moreover, metadata collection can be made precisely during the instants when individual packets of interest are being processed at a device. As a consequence, INT makes it possible to detect and pinpoint network events that were previously imperceptible, such as microsecond congestions.

Although DPP brings greater flexibility to the development of monitoring mechanisms, to operate at line rate on high-speed links, data plane programs are constrained to a small time budget (dozens of nanoseconds) and a limited memory space (e.g., hundreds of megabits of SRAM and dozens of megabits of TCAM) [Bosshart et al. 2013]. Regarding INT, since it involves modifying production packets traversing the network, the amount of metadata that may be collected by a packet is limited by its original size and the network maximum transmission unit (MTU). Additionally, some of the performed actions may increase network load and impact performance. For example, embedding telemetry data into packets increases the load on network links, and generating report packets increases the load on forwarding devices and control channels. We argue that these constraints and factors need to be carefully considered in order to enable the full potential of data plane programmability to the discipline of network monitoring and operation.

3. Hypothesis, Research Questions, and Contributions

This thesis seeks to bridge the gap to materializing the new opportunities for network monitoring and operation brought up by Data Plane Programmability (DPP) and In-band Network Telemetry (INT). Our hypothesis is that INT along with DPP can be successfully applied to monitor and operate networks with per-packet granularity, practicable overheads, and at data-plane timescales. Figure 1 summarizes our work in three parts. In the first part, represented by INTO, we conduct a *theoretical exploration* of in-band network telemetry to provide an understanding of the tradeoffs between granularity and cost




Figure 1. Overview of this thesis with its main contributions.

for this monitoring approach. The second part, IntSight, consists of a *practical system* for efficiently detecting and diagnosing SLO violations via INT-based monitoring. In the third part, we design Felix to enable responding to network failures at data-plane speeds and describe how other network operation tasks can benefit from its *architectural design*. Diving into the specifics of these three parts, the contributions of our work, eight in total, are positioned as a result of the three research questions that we describe next.

Question 1. How can In-band Network Telemetry collection actions be orchestrated across devices in a network to maximize measurement quality while minimizing network and traffic overheads?

Our initial investigation into In-band Network Telemetry showed that when such concept is applied unsystematically to collect metadata from each one of the devices visited by each of the traffic packets traversing the network, substantial burden is placed on the traffic, network resources, and analysis servers. The overhead imposed by this burden can arrive at the point of halting the operation of the network or its effective monitoring. This can be due, for example, to the lack of the necessary space in traffic packets to collect all of the desired metadata, or of enough capacity on links and devices to transport the additional telemetry payload or for the analysis server to keep up with the telemetry reporting rate. With this observation, as our first contribution, we formalize what we call the In-band Network Telemetry Orchestration (INTO) problem by means of Integer Linear Programming. Our ultimate goal with INTO is to minimize monitoring overheads while still obtaining high-quality data. We formalize two variations of the INTO problem as mathematical programming models, each of them focusing on optimizing the usage of a specific resource: the packet processing capacity of devices and the bandwidth of links, respectively. We also prove that both variations of the orchestration problem are NP-Complete. Through an extensive evaluation using real network topologies, we confirmed that generating optimal solutions takes prohibitive amounts of time. As the second contribution, we address the scalability limitation of the mathematical programming models by designing heuristic algorithms. These algorithms are capable of computing high-quality solutions in polynomial computing time for the two variations of the INTO problem. Through our evaluation, we observe that the proposed heuristics are able to generate close to optimal solutions for all of the considered topologies under a second. We also evaluate the quality and costs associated with the proposed heuristics under different aspects and compare their results to identify what types of networks each heuristic is better suited to monitor.

In the following part of the thesis, we continue this study on INT by bringing DPP

into focus and answering the following intriguing question.

Question 2. Given the flexibility in packet processing provided by Data Plane Programmability, can monitoring data be pre-processed or consolidated by forwarding devices before being reported to the control plane to further reduce overhead and with no loss (or even improvement) to measurement quality?

One of the observations from our study of INT orchestration is that programmable data planes have capabilities that could allow for monitoring actions beyond raw metadata collection. In view of that, we introduce IntSight, a system that fully exploits the capabilities of network programmability [Bosshart et al. 2013, Bosshart et al. 2014] to monitor SLOs related to end-to-end delay and bandwidth guarantees. In a nutshell, IntSight discretizes time into sub-second, fixed-length slices called epochs. Network and flow-ofinterest traffic are monitored on a per-packet basis by the data plane to track their state, behavior, and performance (e.g., routing, contention, delay, packet drops, and provided bandwidth). Each production packet is instrumented to carry essential information (in a telemetry header), which has its values systematically updated (through in-network computation) as the packet moves towards the destination. Egress forwarding devices consolidate this information temporarily in memory. At the end of each epoch, the consolidated information kept for each flow of interest enables detecting and diagnosing SLO violation events, their causes, victims, and culprits. As a third contribution of the thesis, thus, we design and implement efficient data plane procedures for gradually computing path-wise metadata such as paths, contention points, end-to-end delays, and provided bandwidth. In our evaluation, we demonstrate the benefits (regarding functionality, performance, and resource footprint) of path-wise in-band network telemetry compared to state-of-the-art approaches, considering six representative network topologies. We observe that IntSight requires up to 10 times less header space and 4 times less memory space for monitoring tasks than the best state-of-the-art approach.

The positive answer to the second question - i.e., the success in pre-processing and consolidating monitoring data directly in the data plane - motivates our last question.

Question 3. Can part of the analysis and reaction logic (traditionally placed in the control plane) be offloaded to the data plane to enable detecting and reacting to network problems in shorter timescales?

We describe cases in which we find the answer to this question to be positive. First, as briefly mentioned, the consolidated information stored in forwarding devices under IntSight enables detecting and diagnosing SLO violation events. During the design of IntSight we observed that the way information is consolidated would enable the forwarding devices themselves to detect epochs in which SLO violations are present. As a result, our **fourth contribution** is an in-network, distributed, path-aware mechanism for monitoring network traffic capable of fine-grained and timely detection of SLO violations and other problems impacting performance (e.g., microbursts). This mechanism enables egress forwarding devices to carry out more selective reporting, only reporting information when it is useful for diagnosis. Control plane servers receive and analyze the generated reports to identify the culprit traffic disrupting network operation. In our evaluations, this approach to telemetry data reporting has shown to make judicious utilization of control plane bandwidth and analysis server resources without significant loss in accuracy and detail. IntSight generates up to two orders of magnitude fewer reports than the state-of-the-art approach to achieve the same diagnosis.

Following our work on IntSight, we present another step taken towards answering Question 3. We shift our focus to network equipment failures and investigate ways for more efficient and resilient rerouting to meet availability SLOs that arise in the context of SDN with programmable data planes. We propose Felix, a novel system that proactively computes forwarding tactics for the normal network state as well as failure scenarios in the control plane and programs these tactics in data plane devices along with a lightweight coordination protocol to immediately react to failures. In a sense, the control plane acts as a strategist devising recovery tactics to handle failures, while the data plane carries out these tactics accordingly when needed. This approach to the problem eliminates the need to wait for the control plane to compute and install new entries upon a failure while also enabling the use of the best alternate paths to bypass failures in general topologies. We make four main contributions with Felix. As the fifth contribution of the thesis, we devise a packet processing pipeline with customized match-action tables that forwards packets according to the current network state. The sixth contribution is the design of a lightweight protocol running on the fast path of switches to enable failure recovery coordination entirely in the data plane. The seventh contribution is the development of algorithms that compute and install alternate forwarding entries just in time to handle possibly-imminent failures. Through an extensive evaluation of Felix, we find that it considerably reduces reaction times while making sensible use of data plane in-device memory. When compared to existing SDN approaches, Felix presents average downtime that is several orders of magnitude lower, greatly reducing packet loss due to failures. The eighth, and final, contribution is the proposal of the Strategy-Tactic paradigm. This paradigm abstracts away Felix's elements to form an overall architecture that can be instantiated to perform other network operation tasks. We exemplify the generality of this paradigm by combining lessons learned from both Felix and IntSight to sketch IntReact, a system for SLO- and contention-aware rerouting at data-plane timescales.

4. Publications and Achievements

As part of the thesis, we have published four main papers [Cordeiro et al. 2017, Marques and Gaspary 2018, Marques et al. 2019, Marques et al. 2020] with a fifth paper [Marques et al. 2023] set to appear at the 2023 edition of the IEEE/IFIP Network Operations and Management Symposium (NOMS). One of our major achievements was receiving the Best Paper Award for our paper on INTO [Marques and Gaspary 2018] at the Brazilian Symposium on Computer Networks and Distributed Systems (SBRC). Another achievement was getting IntSight [Marques et al. 2020] published at ACM International Conference on Emerging Networking Experiments and Technologies (CoNEXT), considering its rigorous review process and small acceptance rate. There, we also earned the ACM reproducibility badges due to our artifacts being not only publicly available, but also functional and reusable.

In addition to our main papers, as part of active collaborations with our colleagues, we have published other six papers [Silva et al. 2018, Lapolli et al. 2019, Lapolli et al. 2019, Silva et al. 2020, Ilha et al. 2021, González et al. 2021, Dalmazo et al. 2021] (a seventh paper is set to appear at NOMS 2023 [Vassoler et al. 2023]) closely related to network programmability that were based on our knowledge in the theoretical, architectural design, and experimental fields obtained during the development of the thesis. A highlight for these additional publications is the Best Student Paper Award received at the IFIP/IEEE International Symposium on Integrated Network Management (IM) in 2019.

5. Conclusion

The work conducted throughout this thesis strongly suggests that our hypothesis that "Inband Network Telemetry (INT) along with Data Plane Programmability (DPP) can successfully be applied to monitor and operate networks with per-packet granularity as well as practicable overheads" is correct. We highlight that the contributions of this thesis touch both theoretical and practical aspects of network management. We mathematically formalized problems and proposed architectural paradigms as well as designed heuristic algorithms and monitoring and operation systems. Throughout, we give special focus to experimentation by considering real world scenarios. Our hope is that with this transversal investigation of the hypothesis, the research questions and their related challenges, we were able to not only tackle important problems, but also sensibly inform future work in the area.

Acknowledgments

This work was supported in part by CAPES – Brazil (Finance Code 1), CNPq – Brazil, RNP – Brazil, FAPESP – Brazil (#2020/05183-0), and CYTED – Spain (#519RT0580).

References

- Balakrishnan, H. (2021). Mind the app! SIGCOMM Lifetime Achievement Award (SIGCOMM'21 Keynote).
- Bosshart, P., Daly, D., Gibb, G., Izzard, M., McKeown, N., Rexford, J., Schlesinger, C., Talayco, D., Vahdat, A., Varghese, G., and Walker, D. (2014). P4: Programming protocol-independent packet processors. *SIGCOMM Comput. Commun. Rev.*, 44(3):87–95.
- Bosshart, P., Gibb, G., Kim, H.-S., Varghese, G., McKeown, N., Izzard, M., Mujica, F., and Horowitz, M. (2013). Forwarding metamorphosis: Fast programmable matchaction processing in hardware for sdn. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, pages 99–110, New York, NY, USA. ACM.
- Cordeiro, W. L. d. C., Marques, J. A., and Gaspary, L. P. (2017). Data plane programmability beyond openflow: Opportunities and challenges for network and service operations and management. *Journal of Network and Systems Management*, 25(4):784–818.
- Dalmazo, B. L., Marques, J. A., Costa, L. R., Bonfim, M. S., Carvalho, R. N., da Silva, A. S., Fernandes, S., Bordim, J. L., Alchieri, E., Schaeffer-Filho, A., Paschoal Gaspary, L., and Cordeiro, W. (2021). A systematic review on distributed denial of service attack defense mechanisms in programmable networks. *International Journal of Network Management*, 31(6):e2163.

- González, L. A. Q., Castanheira, L., Marques, J. A., Schaeffer-Filho, A., and Gaspary, L. P. (2021). Bungee: An adaptive pushback mechanism for ddos detection and mitigation in p4 data planes. In 2021 IFIP/IEEE International Symposium on Integrated Network Management (IM), pages 393–401.
- Ilha, A. d. S., Lapolli, A. C., Marques, J. A., and Gaspary, L. P. (2021). Euclid: A fully innetwork, p4-based approach for real-time ddos attack detection and mitigation. *IEEE Transactions on Network and Service Management*, 18(3):3121–3139.
- Kim, C., Sivaraman, A., Katta, N., Bas, A., Dixit, A., and Wobker, L. J. (2015). In-band network telemetry via programmable dataplanes. In *Proceedings of the 2015 ACM Symposium on SDN Research*, SOSR'15, New York, NY, USA. ACM.
- Lapolli, A. C., Marques, J. A., and Gaspary, L. P. (2019). Offloading real-time ddos attack detection to programmable data planes. In 2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM), pages 19–27.
- Marques, J., Levchenko, K., and Gaspary, L. (2020). Intsight: Diagnosing slo violations with in-band network telemetry. In *Proceedings of the 16th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '20, page 421–434, New York, NY, USA. Association for Computing Machinery.
- Marques, J. A. and Gaspary, L. (2018). Explorando estratégias de orquestração de telemetria em planos de dados programáveis. In Anais do XXXVI Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos, pages 1299–1312, Porto Alegre, RS, Brasil. SBC.
- Marques, J. A., Levchenko, K., and Gaspary, L. P. (2023). Responding to network failures at data-plane speeds with network programmability. In 2023 IEEE/IFIP Network Operations and Management Symposium (NOMS). To appear.
- Marques, J. A., Luizelli, M. C., da Costa Filho, R. I. T., and Gaspary, L. P. (2019). An optimization-based approach for efficient network monitoring using in-band network telemetry. *Journal of Internet Services and Applications*, 10(1):1–20.
- McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., and Turner, J. (2008). Openflow: Enabling innovation in campus networks. SIGCOMM Comput. Commun. Rev., 38(2):69–74.
- Silva, M. V., Marques, J. A., Gaspary, L., and Granville, L. Z. (2018). Identificação de fluxos elefantes em redes de ponto de troca de tráfego com suporte à programabilidade p4. In Anais do XXXVI Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos, pages 1131–1144, Porto Alegre, RS, Brasil. SBC.
- Silva, M. V. B. d., Marques, J. A., Gaspary, L. P., and Granville, L. Z. (2020). Identifying elephant flows using dynamic thresholds in programmable ixp networks. *Journal of Internet Services and Applications*, 11(1):1–12.
- Vassoler, G., Marques, J. A., and Gaspary, L. P. (2023). Vermont: Towards an in-band telemetry-based approach for live network property verification. In 2023 IEEE/IFIP Network Operations and Management Symposium (NOMS). To appear.