Poster: Precise Real-Time Monitoring of Time-Critical Flows

Kilian Holzinger[*], Henning Stubbe[*], Franz Biersack[*], Angela Gonzalez Mariño[‡], Abdoul Kane[‡], Francisco Fons Lluis[‡], Zhang Haigang[§], Thomas Wild[*], Andreas Herkersdorf[†], Georg Carle[†]
[*]firstname.lastname@tum.de [†]lastname@tum.de [‡]firstname.lastname@huawei.com [§]zanghaigang@huawei.com
[*†]Technical University of Munich, Germany [‡§]Huawei Technologies Düsseldorf GmbH, Germany

ABSTRACT

Ethernet is increasingly used in areas where time-critical and safety-relevant data are transported over the network along with best-effort flows, for example in intra vehicle networks or industrial networks. The resulting complex network architectures, time-sensitive networking configurations and system interactions are hard to foresee during the design phase. Therefore, it is hard to rule out any violations of flow specifications or timing and reliability requirements, especially in the presence of unpredictable failures.

In this work, the design of a flow-oriented network monitoring system for time-sensitive applications is presented. It continuously supervises relevant performance metrics with high precision and short detection delay. Moreover, it allows to check compliance with flow specifications in real-time. Initial evaluations using intra vehicle network traffic yield a high measurement precision.

ACM Reference Format:
Kilian Holzinger, Henning Stubbe, Franz Biersack, Angela Gonzalez Mariño, Abdoul Kane, Francisco Fons Lluis, Zhang Haigang, Thomas Wild, Andreas Herkersdorf, Georg Carle. 2021. Poster: Precise Real-Time Monitoring of Time-Critical Flows. In *The 17th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '21), December 7–10, 2021, Virtual Event, Germany*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3485983.3493356

1 INTRODUCTION

Ethernet is pervading more and more application areas which traditionally rely upon specialized field bus systems. This development can be seen, e.g., in intra vehicle networks (IVNs) of modern cars and industrial networks (INs) deployed in industrial control systems [6, 8]. One motivation for this shift are a series of optional Ethernet extensions, referred to as time-sensitive networking (TSN). They enable time deterministic and reliable communication which was the main reason for various legacy busses. Additionally, Ethernet solves the urge for higher bandwidths and also allows to use internet protocols along with specialized application-specific protocols. Thus, safety-critical flows are mixed with best-effort traffic for, e.g., a car entertainment system. The resulting heterogeneous requirements reflect in a complex hard- and software architecture. As a consequence, it is hard to guarantee that in all possible edge-cases no flow specification is violated. This challenge motivates a network monitoring system capable of detecting violations of flow and application requirements in time-critical networking scenarios.

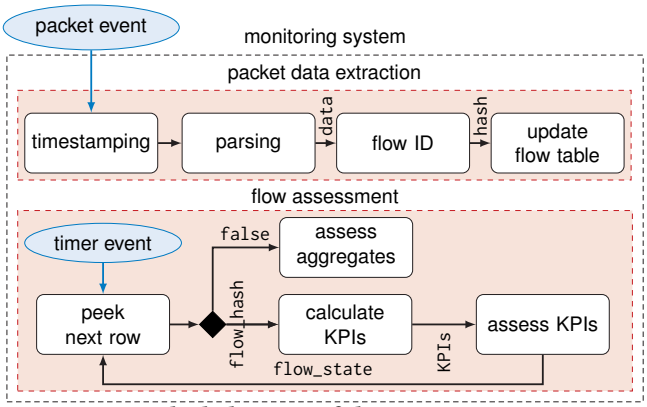

Figure 1: Block diagram of the monitoring system

Being able to observe anomalous behavior in real-time with short delay and high precision opens up the opportunity to use the gathered information subsequently for root cause analysis and mitigation steps. Oeldemann et al. [5] presented a FPGA-based load balancer to optimize resource utilization and improve application latency. Future work is to augment such an approach to consider behavior of network flows, or circumvent failures in the network.

In this work, we present a design of a monitoring system for time-critical flows and evaluate its accuracy and precision in an exemplary setting.

2 RELATED WORK

Time-critical flows are commonly found in IVNs, thus there are comprehensive reviews of the current state of the art and ongoing challenges available [6, 8]. Zeng et al. [8] see potential for additional safety of IVNs by monitoring. Performing flow-level monitoring in software is described in several publications already [2, 9], however without exploring time-sensitive scenarios. An approach to monitor IVN traffic is presented in [7]. The focus of that paper is on attack detection and evaluation is performed in a simulation.

3 DESIGN

A goal of the monitoring system design is to achieve short and predictable detection latency and easy integration into existing network nodes. It shall be scalable and support high flow counts and packet rates.

The monitoring system is deployed as a bump-in-the-wire and performs passive measurements in real-time, hence it does not

CoNEXT '21, December 7–10, 2021, Virtual Event, Germany

ACM ISBN 978-1-4503-9098-9/21/12.
https://doi.org/10.1145/3485983.3493356

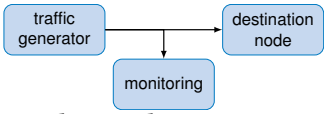

Figure 2: Three node measurement setup

interfere with network traffic. As a DPDK-based software implementation, it is portable to many architectures including SmartNICs. Its poll-mode drivers achieve very low and consistent latency. The program is separated into *packet data extraction* and *assessment* component, as depicted in the internal block diagram in fig. 1, which are interconnected with queues. The resulting packet processing pipeline gives the opportunity for scalability through data-parallelism as well as isolation of critical monitoring tasks from best-effort flows. Information is gathered in the *packet data extraction* step.

Packet data extraction. In the first step after a new *packet event*, a receive timestamp of the arriving packet is taken. Here, the capability of Intel© X550 network adapters which support hardware timestamping with a resolution of 12.5 ns are used. The *parsing* step extracts the involved headers from the received data. In many cases, these are: Ethernet, IP, and transport protocol headers. Also, information embedded in application protocols can be of interest. The monitoring architecture allows flexibility using a callback-based architecture and user-defined logic. Based on an extracted flow identifier, a concurrent hash table, holding necessary flow information, is updated. Both, uni- and bidirectional flows are supported.

Assessment. A *timer event* triggers the key performance indicator (KPI) calculation. All active flows are assessed based on flow information, potentially including historic information. In time-critical networks flows have timing requirements, often extended by a specified communication behavior. As an example, control traffic in INs and IVNs tends to be periodic, i.e., either the target packet inter-arrival-time is known or that the time between consecutive packet bursts is constant. The *assessment* functionality calculates a set of KPIs that capture the degree of alignment with the desired behavior. Based on threshold values, a classification is performed, e.g. to monitor upper limits of one-way-delay or packet loss. KPIs can be further aggregated on the basis of nodes, applications or network links to reason about root causes of a subpar network condition.

The monitoring can optionally use in-band information such as transmit timestamps embedded into encapsulation or application protocols. The Precision Time Protocol achieves sub-microsecond clock synchronization accuracy allowing for direct comparison of send and receive timestamps to calculate e.g., one-way-delays.

4 EVALUATION

Packet arrival timestamps are the basis of multiple relevant KPIs. Therefore, a crucial aspect of the initial evaluation is their precision and accuracy. The used setup is depicted in fig. 2 and resembles a real-world deployment of the solution: The traffic, in this case coming from the *traffic generator*, is provided to the *destination node*. The signal is split using passive fiber taps and hence is provided to the *destination* and the *monitoring* node. The *destination node* relies on MoonGen [1] and its precise packet sniffing capabilities [3].

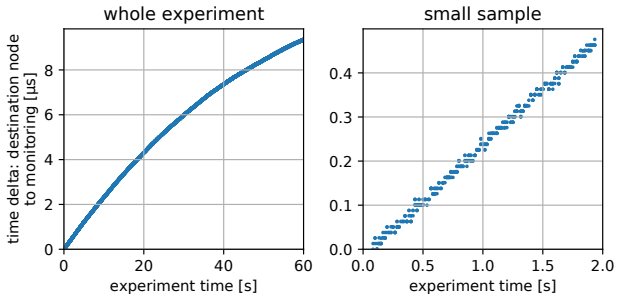

Figure 3: Difference of hardware timestamps between *destination node* and *monitoring*

The traffic produced at the *traffic generator* resembles traffic of an automotive camera system, like described in [4].

The plot in fig. 3 shows the difference of receive timestamps taken at the *monitoring* and *destination node* in an experiment of 60 s and a small subsample. Since the clocks are unsynchronized the time difference of the first packet is set to 0 µs. Hence, only relative changes over time are analyzed in this setup. In the plot, it can be seen that the resulting difference is dominated by systematic errors which can be attributed to clocks drifting slowly apart. The drift is non-linear. The detailed plot on the right shows that the precision is limited by the discrete clock to 12.5 ns. It can be concluded: packet arrivals on both nodes only differ by systematic measurement artifacts; they can be measured with very high precision. Thus, for the purpose of precise monitoring of time critical traffic this approach is viable.

5 CONCLUSION

Motivated by challenges in IVNs and INs, we present an architecture capable of monitoring time-critical flows. An initial DPDK-based implementation is evaluated. We find that the measurement instrument's timer resolution is sufficient to observe slight deviations from specified flow behavior.

There is ongoing work to make the architecture scalable and robust for higher packet rates. Triggering changes in the data-plane to effect e.g. packet forwarding based on detected flow impairments is a promising aspect of future work.

REFERENCES

[1] Emmerich et al. 2015. MoonGen: A Scriptable High-Speed Packet Generator. In *Internet Measurement Conference*.
[2] Emmerich et al. 2018. Efficient dynamic flow tracking for packet analyzers. In *IEEE 7th International Conference on Cloud Networking*.
[3] Gallenmüller et al. 2020. 5G QoS: Impact of Security Functions on Latency. In *IEEE/IFIP Network Operations and Management Symposium*.
[4] Migge et al. 2017. Insights on the Performance and Configuration of AVB and TSN in Automotive Ethernet Networks.
[5] Oeldemann et al. 2020. Inter-Server RSS: Extending Receive Side Scaling for Inter-Server Workload Distribution. In *28th Euromicro International Conference on Parallel, Distributed and Network-Based Processing*.
[6] Tuohy et al. 2015. Intra-Vehicle Networks: A Review. *IEEE Transactions on Intelligent Transportation Systems* (2015).
[7] Waszecki et al. 2017. Automotive electrical and electronic architecture security via distributed in-vehicle traffic monitoring. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (2017).
[8] Zeng et al. 2016. In-Vehicle Networks Outlook: Achievements and Challenges. *IEEE Communications Surveys Tutorials* (2016).
[9] Zhang et al. 2019. FloWatcher-DPDK: Lightweight Line-Rate Flow-Level Monitoring in Software. *IEEE Transactions on Network and Service Management* (2019).