

Δ ELTA: Differential Energy-Efficiency, Latency, and Timing Analysis for Real-Time Networks

Stefan Reif
Friedrich-Alexander University
Erlangen-Nürnberg
reif@cs.fau.de

Andreas Schmidt
Saarland Informatics Campus
andreas.schmidt@cs.uni-saarland.de

Timo Hönig
Friedrich-Alexander University
Erlangen-Nürnberg
thoenig@cs.fau.de

Thorsten Herfet
Saarland Informatics Campus
herfet@cs.uni-saarland.de

Wolfgang Schröder-Preikschat
Friedrich-Alexander University
Erlangen-Nürnberg
wosch@cs.fau.de

ABSTRACT

The continuously increasing degree of automation in many areas (e.g. manufacturing engineering, public infrastructure) lead to the construction of cyber-physical systems and cyber-physical networks. To both, time and energy are the most critical operating resources. Considering for instance the Tactile Internet specification, end-to-end latencies in these systems must be below 1ms, which means that both communication and system latencies are in the same order of magnitude and must be predictably low. As control loops are commonly handled over different variants of network infrastructure (e.g. mobile and fibre links) particular attention must be paid to the design of reliable, yet fast and energy-efficient data-transmission channels that are robust towards unexpected transmission failures. As design goals are often conflicting (e.g. high performance vs. low energy), it is necessary to analyze and investigate trade-offs with regards to design decisions during the construction of cyber-physical networks.

In this paper, we present Δ elta, an approach towards a tool-supported construction process for cyber-physical networks. Δ elta extends the previously presented X-LAP tool by new analysis features, but keeps the original measurements facilities unchanged. Δ elta jointly analyzes and correlates the runtime behavior (i.e. performance, latency) and energy demand of individual system components. It provides an automated analysis with precise thread-local time interpolation, control-flow extraction, and examination of latency criticality. We further demonstrate the applicability of Δ elta with an evaluation of a prototypical implementation.

CCS CONCEPTS

•**Networks** \rightarrow **Cyber-physical networks**; Network reliability;
•**Computer systems organization** \rightarrow *Real-time systems*; *Embedded and cyber-physical systems*; Dependable and fault-tolerant systems and networks;

KEYWORDS

Performance Evaluation, Simulation and Modelling Tools of Real-Time Networks (automotive, aerospace, multimedia, etc.), Networked Embedded Systems and Sensors, Cyber-Physical Systems, Internet of Things

ACM Reference format:

Stefan Reif, Andreas Schmidt, Timo Hönig, Thorsten Herfet, and Wolfgang Schröder-Preikschat. 2018. Δ ELTA: Differential Energy-Efficiency, Latency, and Timing Analysis for Real-Time Networks. In *Proceedings of RTN'2018 The 16th International Workshop on Real-Time Networks, Barcelona, Spain, July 2018 (ECRTS-RTN'18)*, 6 pages.

1 INTRODUCTION

With the increasing degree of automation in domains such as manufacturing engineering, mobility and logistics, as well as public infrastructure, great efforts are being made to close the gap between cyber (or digital) and physical processes, commonly described as *cyber-physical systems* (CPS). The communication between these CPSs is becoming particularly important, as control loops are handled over different variants of network infrastructure (e.g. fibre links, LTE/5G) to make measurements (i.e. sensor data) available from a variety of systems and implement the necessary, distributed coordination. These communicating systems are commonly referred to as *cyber-physical networks* (CPN) and represent a special division of *real-time networks*. CPNs share the requirements of the underlying cyber-physical systems which implies that such networks must ensure fault-tolerance or resilience, provide real-time characteristics, allow the execution of distributed tasks, and incorporate self-sufficient systems which are battery driven and demand low-power operations with predictable energy footprints.

CPNs not only require novel solutions on all layers of a communication stack (i.e. transport, network, operating system and hardware), they also demand for cross-layer integrations to fulfill these requirements [12]. The work on solutions for the arising challenges requires tools that analyze the timing behavior as well as the energy demand at system level. Correspondingly, such development tools support the construction process of CPNs and allow the evaluation of system changes, validate guarantees (i.e. as to time and energy demand), and guide the process of decision making for further improvements of the system design. This paper presents

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ECRTS-RTN'18, Barcelona, Spain

© 2018 Copyright held by the owner/author(s).

and evaluates an approach towards tool-based, automated timing and energy demand analysis of CPNs. The approach improves and extends X-LAP [17], a timing-analysis tool specifically designed for CPNs. We demonstrate the validity of the proposed approach by analyzing the *Predictably Reliable Real-time Transport (PRRT)* protocol, which is a transport layer protocol for CPNs. PRRT respects application constraints, such as maximum latency and tolerance to lost messages, and parametrizes its internal mechanisms for congestion and error control to fulfill these.

The contribution of this paper is threefold:

- We present an approach which jointly captures, analyzes, and correlates the runtime behavior (i.e. performance, latency) and energy demand of CPNs
- We propose an automated analysis which provides precise thread-local time interpolation, control-flow extraction, and analysis of latency criticality.
- We extend our analysis to compare multiple sets of traces, which automatically extracts the impact of code changes or hardware configurations (i.e. varying processor speeds).

The remainder of this paper is structured as follows: Section 2 gives background information as well as related work. Section 3 presents details on the implementation, while Section 4 shows the evaluation results. Section 5 concludes the paper and gives directions for further research and implementation.

2 BACKGROUND AND RELATED WORK

In CPSs and CPNs in particular, we see that the requirements that are posed [12] cannot be fulfilled by analyzing the network or operating system independently. However, the networking and operating systems domains have been using completely different sets of tools to evaluate system performance and validate adherence to design goals. Typically, this process involves abstracting away the other domain [5, 19] and simplifying it by characterising representative loads (from a system perspective) or representative task execution times (from a network perspective). With the precision and timescale required for CPSs [2], these abstractions are no longer sufficient and it is required to properly co-design and analyze both domains at the same time [3, 4].

Communication and operating system latencies used to be in different orders of magnitude (communication: 10ms to 10s; system: 10us to 1ms), if we for instance consider applications running on the Internet and covering large physical distances. In contrast, CPN applications have stricter time constraints (cf. the Tactile Internet [6]), e.g. a total end-to-end latency of 1ms, so that the two domains tend to contribute to the overall latency more evenly. For the networking domain, this means that the communication distances are strictly limited (below 300km to achieve 1ms propagation delay with speed of light) and queueing delays must be kept as close to zero as possible, using pacing and congestion control. The latter is already motivated in previous work on data center networking [1], which does not deal with real-time requirements, but still requires that the delays are predictable and low [4]. Considering that data centers and clouds are going to play an import role in offloading tasks for upcoming CPN applications, these approaches will also be part of network stacks tailored to the needs of CPNs.

Another approach to timing and energy-demand analysis is off-line static analysis, before run-time. The goal of such an analysis is typically to derive a *worst-case execution time* (WCET) [22] and *worst-case travel time* (WCTT) [14] for network protocols. Delta and X-LAP, in contrast, monitor the behavior at run-time. In consequence, Delta is inherently restricted to the actually observed behavior, and it cannot guarantee that the collected samples cover worst-case scenarios. However, this approach is still viable for soft real-time systems that are often too complex for exhaustive static analysis. Besides, the information gathered with X-LAP and Delta can be used to verify timing information gained by static analysis. An example for a tool-based approach for the construction of real-time networks based on static analysis is RT-Appia [18]. This framework composes protocol components to real-time network stacks. Thus, the protocol stack is tailored to the application demand in order to improve WCET analyzability.

2.1 X-LAP

X-LAP¹ [17] is a timing analysis tool specifically designed for CPNs. In these networks, the network protocol stacks are often tightly coupled with operating system in order to fulfill real-time requirements with high reliability. X-LAP therefore works cross-layer, by profiling the performance on transport, operating system and network layer. Furthermore, it works on multiple nodes and allows to correlate the timing information taken on these systems, without explicit clock synchronization. X-LAP consists of two components: 1) a set of C-level calls to be injected into the code for capturing high resolution clock- and cycle-stamps with minimal overhead; 2) a set of analysis procedures to determine the causes of latency and jitter. Sender and receiver systems that use these C-level calls are going to produce a table with packet traces after termination, which gives the clock- and cycle stamps together with the packet sequence number. Afterwards, the processing chain of X-LAP interpolates timestamps—we only take them when necessary, because cycle-stamps lead to a lower overhead induced by the profiling—and calculates the durations of certain processing steps, as well as the overall execution time by different threads. Consequently, X-LAP is able to compute exact processing durations for fine-grained processing steps of every packet, through the entire CPN—providing insights that other (single-layer) tools, e.g. *wireshark*² for networking or *valgrind*³ for systems, cannot deliver.

2.2 Predictably Reliable Real-time Transport

We use X-LAP for evaluating the *Predictably Reliable Real-time Transport (PRRT)*⁴ protocol [8]. PRRT is a protocol that aims to overcome the shortcomings of existing transport layer protocols, such as TCP, UDP and QUIC [7], in particular with respect to time- and resilience-awareness. In contrast to these solutions, PRRT provides a *partially reliable ordered datagram stream with bounded latency*, in contrast to TCP (*fully reliable ordered byte stream*), UDP (*datagram stream*) and QUIC (*multiplexed fully reliable ordered byte streams with reduced number of round-trips*), which all do not provide bounds on timing.

¹<http://xlap.larn.systems>

²<https://www.wireshark.org/>

³<http://valgrind.org/>

⁴<http://prrt.larn.systems>

PRRT achieves this by a) letting the application state its requirements on maximum latency and acceptable residual error rate and b) incorporating this into the protocol's operation, together with measurements of the channel's propagation delay, loss rate and bottleneck bandwidth. Therefore, it is necessary that code segments across multiple threads within PRRT take a predictable amount of time and face only minimal jitter. For achieving predictable delay on the network layer, PRRT uses multiple mechanisms: First, it uses *adaptive hybrid error correction* [13] to choose an optimal configuration for providing proactive (FEC) and reactive (ARQ) error control, based on the application constraints and channel parameters. An optimal coding scheme ensures that after the application's maximum tolerable latency, at most its maximum tolerable residual loss is present. Second, it uses packet pacing to adjust its effective sending rate to the bottleneck bandwidth, effectively avoiding queuing delays when packets are sent out in bursts. Finally, pacing together with measurements on the propagation delay are used to compute the congestion window over the given path and avoid further losses due to competitions for bandwidth with other applications.

In general, evaluations using X-LAP can be executed with other transport layer protocols as well, as long as the source code is available and can be changed to inject clock- as well as cyclestamping calls. Furthermore, the protocol must provide some notion to identify packets, e.g. a sequence number, so that traces for individual pieces of data can be generated.

3 IMPLEMENTATION

The implementation of Delta bases on X-LAP [17], and reuses its high-precision, low-interference time-stamping architecture. However, it significantly extends the off-line analysis component, providing additional insights and increasing automation.

The analysis approaches we present in this paper augment the timing analysis facilities of X-LAP by a *differential* approach that *compares* the timing of multiple experiments. We extract information that are particularly useful for the design, implementation, and evaluation of real-time networks.

The analysis approaches work *retroactively*, in the sense that they derive information from observing a real-time network stack in action. Such information is necessarily imprecise because it is restricted to the actually observed behavior. However, we show in the following, that we can derive valuable information with little design-time and run-time costs. *Proactive* approaches that *predict* the behavior, for instance by static analysis, are beyond the scope of this paper.

3.1 Control Flow Reconstruction

When searching for and mitigating the interferences between different processing steps in a transport stack, the control flow graph is an essential tool to foster the analysis. As control flows can vary between different protocol versions or even minimal code changes, it is crucial for a cross-layer analysis to extract this piece of information from the captured data. Obviously, this cannot fully replace the expertise of a developer and a thorough analysis of the source code, but in particular with concurrent systems, such an empirical approach can greatly support the development process.

The reconstruction works as follows: We compute a *happens-before* relation [10] $A \rightarrow^+ B$ of events A and B if, for every packet, A has a lower time-stamp than B . If $A \rightarrow^+ B$ and $B \rightarrow^+ A$, A and B are concurrent. Concurrency can occur if events happen in different threads, or in interrupt handlers, such as timers.

Since *happens-before* is transitive, we derive a *happens-directly-before* relation $A \rightarrow B$ of events A and B if $A \rightarrow^+ B$ and no event E exists with $A \rightarrow^+ E \rightarrow^+ B$. This relation reconstructs control flows, because it reveals the sequence of events in the protocol stack. A representative sub-graph is depicted in Figure 1, showing the flow packets follow through a specific version of the PRRT stack.

3.2 Latency-Criticality Analysis

The analysis steps in the following sections rely on the information whether the duration of a specific code segment impacts the end-to-end latency. We therefore need to quantify the relation between code segments and the end-to-end latency. Thereto, we define the *latency criticality* of a code segment $\langle A, B \rangle$, as the Pearson correlation coefficient between the duration of $\langle A, B \rangle$ and the *end-to-end* (E2E) latency, over all packets in a trace. This value describes whether the E2E latency depends on the duration of a code segment. Thereby, we utilize the control flow reconstruction to automatically detect whether A and B refer to an actual code segment in the protocol stack.

The latency-criticality analysis measures how individual code segments influence the end-to-end delay. Thus, it gives a hint on which parts of the protocol are worth considering for further optimization, to fulfill the requirements of CPNs.

3.3 Timing-Predictability Analysis

A key property of CPNs, as well as real-time networks, is timing predictability. Therefore, our analysis identifies protocol parts that exhibit unpredictable behavior.

The reasons for non-reproducible timing behavior are many-fold. For instance, many network protocols incorporate concurrency due to timers, task synchronization or thread interaction. Besides, hardware and software-related interference such as cache misses or OS noise [21] influence the timing behavior of code segments. Furthermore, the timing of specific code segments might depend on current channel properties (e.g. packet loss or propagation delay).

In order to identify code parts where the timing behavior is not sufficiently predictable, we run the same exact version of the code multiple times and trace each packet using X-LAP. Afterwards, we reconstruct the control flow from the timestamps to identify the code segments of interest.

For each code segment, we apply the *k-sample Anderson-Darling-Test* [20]. This stochastic test is a well-established method to check whether multiple samples are derived from the same distribution. While, technically, the durations are drawn from the same actual runtime-distribution of code segments, the actual data might show significant differences between consecutive runs.

The stochastic test only classifies timing behavior as "different" when it is certain, given a configurable significance parameter. When the test is not certain enough, it classifies the timing distributions as "similar". However, this test can give false-negatives when timing is actually unpredictable but, incidentally, behaves the

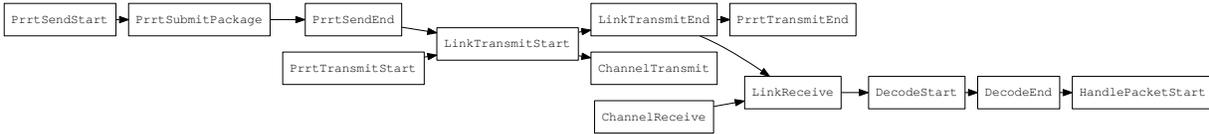


Figure 1: A control flow graph reconstructed from the *happens-directly-before* relation

same at all observed experiment runs. However, if the test result is “different”, then the two experiment runs have certainly shown different timing behavior.

The actual importance of the unpredictability results depends on the latency criticality of the code segment. At some protocol parts, unpredictability can be expected, for instance, regarding timers that cause concurrent code execution. Therefore, the timing unpredictability results should be analyzed jointly with the results of the latency-criticality analysis.

3.4 Modification Tracking

When changing parts of the implementation of protocols, the timing behavior of (at the first glance) independent protocol parts can change. The reason is interference that is sometimes hidden: For instance, cache effects and contention on hardware buses can cause non-trivial interference between seemingly independent protocol parts. Besides, if the timing of network packets changes slightly, the entire protocol behavior might adapt.

Our analysis reveals such timing interference. We execute different protocol versions and, for each version, capture packet traces with X-LAP. Then, we apply the Anderson-Darling-Test similar to the timing predictability evaluation described in Section 3.3. This test reveals code segments where the timing behavior has changed. We can also identify control flow changes using the automated control flow reconstruction described in Section 3.1.

3.5 Energy-Efficiency Analysis

Another use of differential analysis for real-time networks is energy-efficiency optimization. Modern hardware components offer multiple “knobs” that affect power demand as well as performance. These configurations include processor frequencies and sleep states with varying wake-up latencies. Identification of the most efficient states while maintaining desired real-time behavior is therefore a complex problem [16].

To analyze the relation between energy demand and performance, we trace packets with X-LAP and measure the end-to-end latency. Simultaneously, we measure the energy demand of the sender and the receiver hosts during the experiment. We repeat these measurements for the various configurations of the hardware and thereby identify the most energy-efficient setup.

3.6 Slowdown Analysis

Besides finding the most energy-efficient hardware configuration, Delta offers additional fine-grained timing information, giving insight into the detailed behavior of individual protocol parts.

To analyse how protocol parts behave under different hardware power states, we measure a *fast* and a *slow* experiment run. Then, we reconstruct the control flow and, for each code segment X , we compute the average delays $D_{X,fast}$ and $D_{X,slow}$, from the fine-grained timing information of the respective experiment runs. We compute the *slowdown* S_X of each segment, and we normalize the slowdown values using the end-to-end slowdown:

$$S_X = \frac{D_{X,slow}}{D_{X,fast}} \quad S_X^* = \frac{S_X}{S_{E2E}}$$

The normalized slowdown reveals whether the duration a code segment depends on the processor speed. If S_X^* is above 1, the segment latency suffers from the slow hardware configuration. A value below 1, however, means that the performance loss at a code segment is less than the end-to-end performance loss.

4 EVALUATION

We evaluate our analysis approaches using X-LAP and PRRT. We obtain detailed package traces and submit them to the Delta analysis techniques described in Section 3. The purpose of this evaluation is to find out whether the analysis techniques presented in Section 3 produce meaningful results. We therefore apply them to a real-time networking protocol, PRRT.

4.1 Experimental Setup

We run all experiments on two nodes in a dedicated networking testbed to make sure that only minimal interference on the hosts and the network disturb our experiments.

4.2 Control Flow Reconstruction

To evaluate the control flow reconstruction, we have traced 4095 packets with X-LAP. Figure 1 depicts a representative sub-graph of the computed *happens-directly-before* relation. For events that happen in the same thread, the analysis reconstructs the control flow reliably. However, one computed flow was a *false-positive* because, by chance, an event in one thread always happened before another event in another thread. We also encountered one *false-negative*, where two events happened nearly simultaneously and the clock resolution was not sufficient to order the two timestamps.

4.3 Latency Criticality Evaluation

We evaluate the latency criticality of all code segments identified by the control flow reconstruction. Figure 2 visualizes results for the most-critical code segments. The values correspond loosely to the average segment latency, because jitter is often proportional to the duration. Therefore, long segments correlate stronger with the

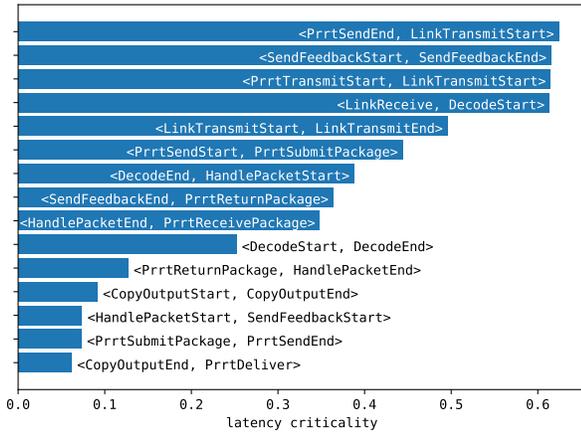


Figure 2: Latency criticality of protocol parts

end-to-end latency. However, some segments have low criticality because of concurrency in the protocol stack.

4.4 Timing Predictability Evaluation

To evaluate the timing predictability analysis, we have traced two samples with 100 packets each. Both samples use identical code revisions, and the hardware frequency was fixed to 3 GHz. We run the Anderson-Darling-Test on both data sets, which reveals that 9 code segments have non-reproducible timing, 5 of which were expected because of concurrency or interaction with the network hardware. Out of the 4 remaining segments, only one has a relatively high latency criticality (<PrprtSendStart, PrprtSubmitPackage>). Thus, the analysis gives a hint where the protocol implementation can be optimized to improve timing predictability.

4.5 Modification Tracking and Evaluation

We repeat the timing predictability analysis, but we add an artificial *sleep* call as a dummy code modification that impacts the timing behavior of a single code segment. The sleep call was already in a segment with non-reproducible timing, at the sender side. We compare the results with the timing reproducibility analysis to eliminate segments where we know that the timing changes between experiment runs. The analysis reveals that two code segments at the receiver side change their timing behavior, according to the Anderson-Darling-Test. Thus, our analysis reveals interference between protocol parts.

4.6 Energy Efficiency

We evaluate the system energy efficiency at three performance states: We configure our system to run at a fixed speed of 1 GHz, 2 GHz, or 3 GHz, for each evaluation run. Thereby, we use X-LAP to measure the end-to-end package latency. Simultaneously, we measure the energy demand with RAPL [9].

We compute the average ET² metric [15] per packet, because it is a fair comparison for DVFS settings. Figure 3 summarizes the evaluation results. Thereby, the most energy-efficient configuration is neither the fastest nor the slowest. This result is well-aligned

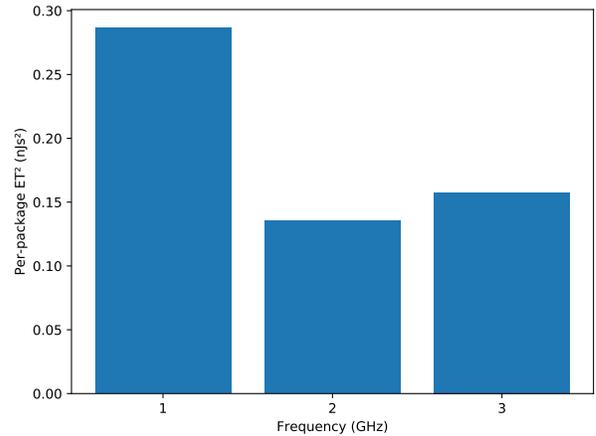


Figure 3: Energy efficiency at different processor speed settings (lower is better)

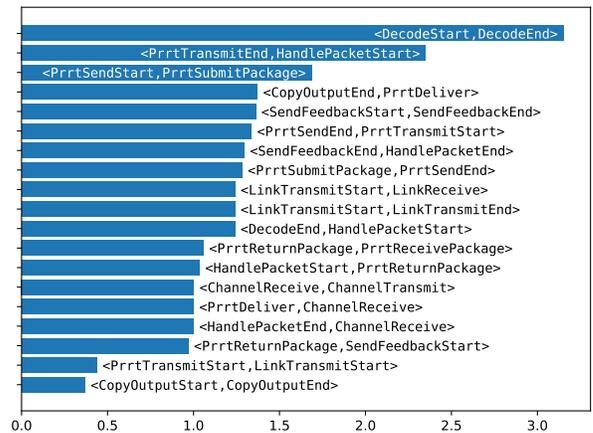


Figure 4: Normalized slowdown for code segments

with other research [11]. This outcome also indicates that energy efficiency needs actual energy measurements to find the optimal hardware configuration.

4.7 Normalized Slowdown

We further analyse the precise packet timings at 2 GHz and 3 GHz and compute the normalized slowdown for each code segment, using the control flow reconstruction to identify relevant code segments. The results are summarized in Figure 4, showing that on the one hand, we find a relatively high normalized slowdown at the <DecodingStart, DecodingEnd> segment, because the *decoding* operation is CPU bound. On the other hand, memory-bound segments, such as <CopyOutputStart, CopyOutputEnd>, exhibit a low normalized slowdown. These results indicate that different error codes (or optimized versions of it) might allow to use a more energy-efficient hardware configuration.

5 CONCLUSION

This paper has introduced Δ , a collection of analysis techniques for cyber-physical networks. In summary, these novel techniques offer detailed insight on the fine-grained timing behavior and energy demand. We automatically reconstruct control flows from time measurements. We utilise this information in a differential analysis approach to detect code segments with unpredictable timing. Additionally, our analysis detects the impact on code revisions on the overall system behavior. The evaluation shows that minor changes in one processing step can indeed change the timing of various other, seemingly unrelated, parts of the protocol. We further combine precise timing information with actual energy measurements. Our analysis identifies the most energy-efficient hardware configuration, and evaluates quantitatively which code segments tolerate low processor speeds.

In summary, our proposed approach helps designers of cyber-physical networks to verify timing properties, to reduce the end-to-end latency and jitter, and to increase the energy efficiency of the overall systems.

In future work, we are going to apply these analyses to improve transport protocol stacks for cyber-physical networks, to achieve predictably low end-to-end latency and a high energy efficiency.

ACKNOWLEDGMENTS

The work is supported by the German Research Foundation (DFG) as part of SPP 1914 “Cyber-Physical Networking” under grants HE 2584/4-1 and SCHR 603/15-1.

REFERENCES

- [1] Mohammad Alizadeh, Abdul Kabbani, Tom Edsall, Balaji Prabhakar, Amin Vahdat, and Masato Yasuda. 2012. Less is More: Trading a little Bandwidth for Ultra-Low Latency in the Data Center. In *Proceedings of the 9th Conference on Networked Systems Design and Implementation (NSDI)*. USENIX, 1–14.
- [2] Luiz Barroso, Mike Marty, David Patterson, and Parthasarathy Ranganathan. 2017. Attack of the Killer Microseconds. *Communications of the ACM* 60, 4 (March 2017), 48–54.
- [3] Pete Beckman, Kamil Iskra, Kazutomo Yoshii, and Susan Coghlan. 2006. The Influence of Operating Systems on the Performance of Collective Operations at Extreme Scale. In *Proceedings of the 8th IEEE Annual International Conference on Cluster Computing (CLUSTER 2006)*. IEEE, 1–13.
- [4] Jeffrey Dean and Luiz André Barroso. 2013. The Tail at Scale. *Commun. ACM* 56, 2 (Feb. 2013), 74–80.
- [5] Thomas Ferrandiz, Fabrice Francès, and Christian Fraboul. 2009. A method of computation for worst-case delay analysis on SpaceWire networks. In *Proceedings of the 11th IEEE International Symposium on Industrial Embedded Systems (SIES 2009)*. IEEE, 19–27.
- [6] Gerhard P. Fettweis. 2014. The tactile internet: Applications and challenges. *IEEE Vehicular Technology Magazine* 9, 1 (March 2014), 64–70.
- [7] Google Chromium Project. 2017. QUIC, a multiplexed stream transport over UDP. (2017). <https://www.chromium.org/quic>
- [8] Manuel Gorius. 2012. *Adaptive Delay-constrained Internet Media Transport*. Ph.D. Dissertation. Saarland University.
- [9] Intel Corporation. 2018. *Intel Architecture Software Developer’s Manual*. Intel Corporation, Santa Clara, California, USA.
- [10] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (July 1978), 558–565.
- [11] Etienne Le Sueur and Gernot Heiser. 2011. Slow Down or Sleep, That is the Question. In *Proceedings of the USENIX Annual Technical Conference (ATC 2011)*. USENIX, 1–6.
- [12] Edward A Lee. 2008. Cyber physical systems: Design challenges. In *Proceeding of the 11th IEEE International Symposium on Object Oriented Real-Time Distributed Computing (ISORC)*. IEEE, 363–369.
- [13] I-Chung Lee, Cheng-Shang Chang, and Ching-Ming Lien. 2005. On the throughput of multicasting with incremental forward error correction. *IEEE Transactions on Information Theory* 51, 3 (March 2005), 900–918.
- [14] Meng Liu, Matthias Becker, Moris Behnam, and Thomas Nolte. 2016. Using Segmentation to Improve Schedulability of Real-Time Traffic over RRA-based NoCs. *ACM SIGBED Review* 13, 4 (Sept. 2016), 20–24.
- [15] Alain J. Martin, Mika Nyström, and Paul I. Pénzes. 2002. ET²: A metric for time and energy efficiency of computation. In *Power Aware Computing*, Robert Graybill and Rami Melhem (Eds.). Kluwer Academic Publishers, 293–315.
- [16] Akihiko Miyoshi, Charles Lefurgy, Eric Van Hensbergen, Ram Rajamony, and Raj Rajkumar. 2002. Critical Power Slope: Understanding the Runtime Effects of Frequency Scaling. In *Proceedings of the 16th International Conference on Supercomputing (ICS)*. ACM, 35–44.
- [17] Stefan Reif, Andreas Schmidt, Timo Hönig, Thorsten Herfet, and Wolfgang Schröder-Preikschat. 2017. X-Lap: A Systems Approach for Cross-Layer Profiling and Latency Analysis for Cyber-Physical Networks. In *Proceedings of the 15th International Workshop on Real-Time Networks (RTN)*. 6.
- [18] Joao Rodrigues, Hugo Miranda, João Ventura, and Luis Rodrigues. 2001. The design of RT-Appia. In *Proceedings 6th International Workshop on Object-Oriented Real-Time Dependable Systems (WORDS)*. IEEE, 261–268.
- [19] Andreas Schimmel and Alois Zoitl. 2010. Real-time communication for IEC 61499 in switched ethernet networks. In *Proceedings of the 2nd IEEE International Congress on Ultra Modern Telecommunications and Control Systems and Workshops (ICUMT)*. IEEE, 406–411.
- [20] F. W. Scholz and M. A. Stephens. 1987. K-Sample Anderson-Darling Tests. *J. Amer. Statist. Assoc.* 82, 399 (Sept. 1987), 918–924.
- [21] Dan Tsafir, Yoav Etsion, Dror Feitelson, and Scott Kirkpatrick. 2005. System Noise, OS Clock Ticks, and Fine-grained Parallel Applications. In *Proceedings of the 19th Annual International Conference on Supercomputing (ICS)*. ACM, 303–312.
- [22] Reinhard Wilhelm, Jakob Engblom, Andreas Ermedahl, Niklas Holsti, Stephan Thesing, David Whalley, Guillem Bernat, Christian Ferdinand, Reinhold Heckmann, Tulika Mitra, Frank Mueller, Isabelle Puaut, Peter Puschner, Jan Staschulat, and Per Stenström. 2008. The Worst-case Execution-time Problem—Overview of Methods and Survey of Tools. *Transactions on Embedded Computing Systems (TECS)* 7, 3 (May 2008), 36:1–36:53.