

Presentation Title: Network Calculus-Based Co-Design for Real-Time Performance and Resource Optimization in TSN

Abstract:

Time-Sensitive Networking (TSN) is designed to enable real-time communication in time-critical domains by providing a foundational framework of sub-standards. However, the practical realization of real-time communication heavily depends on the co-design of real-time performance verification and network resource configuration. Traditional approaches to co-designing resource configuration optimization and real-time performance verification rely on an iterative feedback paradigm, which renders the overall optimization process highly time-consuming. This abstract highlights our recent work on a tightly coupled co-design model based on Network Calculus (NC) theory for optimizing network resources in various TSN scheduling mechanisms. By introducing an analytical method grounded in optimization theory, we derive a closed-form solution that reduces the complexity of co-designing real-time performance and configuration optimization while ensuring reliable real-time guarantees and improving resource utilization.

Presenter: Luxi Zhao, BUAA, China

Presentation Title: Telescoping-Tightness Single-node Performance Bounds

Abstract:

In this talk, we contrast the state of the art in stochastic bounds for single node queueing models with a new class of bounds with a telescoping property. This property allows expressing the CCDF of the steady-state waiting time $P[W > \sigma]$ exactly for simple (GI/G/1 or AR/G/1) queues as well as for Markovian fluid queues. Remarkably, this expression that takes the form of an infinite sum with positive terms allows deriving closed-form bounds by evaluating a finite number of terms. Keys to this expression are (i) an exponential change of measure that reverses the sign of the drift of the underlying random walk, and (ii) a representation of the overshoot $R_\sigma = \sum_{i=1}^T X_i - \sigma$ of a point process. In this talk, we will discuss the main idea behind this class of bounds in comparison to the weakness of current Martingale waiting time bounds. This talk is based on a previously published paper [1].

[1] F. Ciucu, S. Mehri and A. Rizk, "On Ultra-Sharp Queueing Bounds," IEEE Conference on Computer Communications (INFOCOM), 2024, pp. 1980-1988, doi: 10.1109/INFOCOM52122.2024.10621177.

Presenter: Amr Rizk, Leibniz University Hannover, Germany

Presentation Title: A Network Calculus Model of Congestion Control in Data Center Networks

Abstract:

Congestion control is essential for managing data center traffic and preventing network overload. In many data center environments, certain parameters, such as propagation delay between hosts or available bandwidth along a given path, remain constant during operation. We introduce a novel congestion control algorithm that leverages these invariances by modelling the network as a multi-flow window flow control system with a common bottleneck switch and homogeneous feedback delay.

Using network calculus, we derive sufficient conditions for our algorithm to unthrottle the network, ensuring high throughput in dynamic traffic scenarios. We further analyze the conditions under which the network achieves max-min fairness, particularly when the bottleneck switch employs a first-in-first-out (FIFO) scheduling discipline. Additionally, by examining scenarios where flows exhibit periodic behaviour, we develop a mechanism to mitigate traffic burstiness while imposing minimal rate limiting.

Our approach provides analytical guarantees regarding throttling and fairness and offers new insights into designing congestion control strategies for data center networks. These findings can serve as a foundation for future work in optimizing networks through parameter-aware control mechanisms.

Presenter: Natchanon Luangsomboon, University of Toronto, Canada

Presentation Title: Nancy.Expressions: a library for automated optimizations

Abstract:

We present Nancy.Expressions, an open-source library that extends the existing Nancy tool for DNC. Building on Nancy's single-operation implementations, Nancy.Expressions introduces a novel computer algebra system enabling users to construct, manipulate, and simplify expressions involving DNC operations. This system supports automated optimizations through equivalence rules, such as simplifying subadditive functions using $f \otimes f = f$, as well as re-usable implementation of known techniques such as Finitary RTC. This tool not only enhances the analysis of network systems but also opens new avenues for heuristic algorithms to optimize DNC expressions with minimal overhead.

Presenter: Raffaele Zippo, University of Pisa, Italy

Presentation Title: Statistical delay QoS optimization for XR traffic over WiFi networks

Abstract:

The ability to seamlessly merge physical and digital worlds has made eXtended Reality (XR) a highly researched topic in academia and industry. To ensure the immersion and presence, XR applications impose stringent requirements on statistical delay, including the delay bound and delay bound violation probability. Meanwhile, WiFi has experienced rapid proliferation and is now widely used for multimedia content distribution. However, reliable quality of service (QoS) support over WiFi is a challenging endeavor. This leads to a compelling question: Can WiFi effectively support XR applications? In this talk, we firstly introduce a feasible analysis and optimization framework of statistical delay QoS for WiFi. In this framework, the arrival process of XR traffic is formulated as auto regressive (AR) model, and the service process of WiFi is modelled as the Markov modulated Bernoulli process. The service model is concise and captures essential features of carrier-sense multiple access with collision avoidance (CSMA/CA), including random backoff. Moreover, it facilitates optimization using reinforcement learning by integrating easily measurable parameters such as the transmission probability and the successful transmission probability. Then, we construct a martingale with respect to the arrival and service sequences, and the martingales theory is utilized to derive the upper bound of delay bound violation probability. Finally, the statistical-delay-optimal CSMA/CA algorithm is formulated as a Markov decision process (MDP). The VDN (Value-Decomposition Networks) based Q-learning approach is utilized to solve the problem. As a result, the initial contention window is dynamically optimized to minimize the average of delay bound violation probabilities in a distributed way. Simulation results show the superiority of statistical-delay-optimal CSMA/CA which can meet the stringent demands of XR applications.

Presenter: Linlin ZHAO, Jilin University, China

Presentation Title: Delay Bounds in IEEE 802.11 Networks Under Non-Saturated Conditions: Do Practical Models Improve Accuracy?

Abstract:

Understanding delay performance in networked systems with contention-based channel access is crucial for the development and deployment of delay-sensitive applications in wireless scenarios. Stochastic Network Calculus (SNC) provides a mathematical framework for analyzing performance guarantees under uncertainty. Instead of directly handling probability distributions, the Moment Generating Function (MGF)-based SNC leverages exponential tail bounds to derive probabilistic performance metrics, such as backlog and delay. In this work, we evaluate the tightness of delay bounds in an IEEE 802.11 single-cell network under non-saturated conditions, using different MGFs to model the service provided by the Distributed Coordination Function (DCF). These models range from theoretical formulations to practical approximations obtained via network simulations. Special attention is given to the independence assumption and the overestimation of probabilities induced by the application of Boole's inequality and the Chernoff bound when deriving an upper bound on the probability that the delay exceeds a given threshold.

Presenter: Orangel Azuaje Contreras, University of Porto, Portugal

Presentation Title: Stochastic Performance Bounds Under Weak Service - Minimal Arrival Guarantees as Enabler

Abstract:

Understanding and analyzing performance metrics in queueing systems is crucial for evaluating the timeliness and reliability of networked and distributed environments. Stochastic network calculus (SNC) provides rigorous performance bounds, but many existing results rely on the assumption that servers are work-conserving. While common, this assumption can be overly restrictive, particularly when deriving per-flow bounds - an essential aspect and strength of SNC.

In this presentation, we extend SNC by relaxing the work-conserving requirement, broadening its applicability to more general service models. This is achieved through minimum arrival guarantees, formulated as Laplace bounds on arrival processes. By leveraging these guarantees, we derive new per-flow delay bounds that enhance the potential of SNC analysis. To that end, we establish Laplace bounds for Markov-modulated arrivals, demonstrating the versatility of our new approach.

Presenter: Vlad-Cristian Constantin, TU Kaiserslautern, Germany

Title: Improving arrival curves for flow splitting

Abstract

We present a novel method for refining the arrival curve of flows in a system where packets are divided into subflows based on packet types, e.g. due to destination address or load balancing. Previous attempts at modelling load balancing have found either overly pessimistic results (i.e., the same arrival curve as the input) or wrong results (of which we recap the counterexamples). In this work we consider a model where packets are generated in cycles, each containing a variable number of requests of different types, with partial knowledge of their spatial distribution within each cycle. Leveraging this, we find improved arrival curves for each type of packets, which can be used to improve the worst-case analysis of these systems.

Presenter: Edoardo Loni, University of Pisa, Italy

Presentation Title: Per-Flow Performance Bounds for a FIFO Server Beyond Token-Bucket-Constrained Arrivals

Abstract:

Network Calculus (NC) is a powerful tool for analyzing various scheduling disciplines; however, FIFO remains a particularly challenging case. Existing NC analyses of FIFO systems are mostly limited to traffic constrained by simple token bucket arrival curves. In this presentation, we extend the NC analysis of a FIFO server to more general arrival curves, particularly piecewise linear (PWL) curves. We show that delay-optimal residual service curves can be derived for concave arrival curves by leveraging the FIFO property that the delay bound of an aggregate of flows equals the per-flow delay bounds. To obtain backlog-optimal residual service curves, we separate PWL arrival curves into their underlying token bucket components and utilize these to systematically derive the corresponding service curve.

Presenter: Lukas Wildberger, Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau (RPTU), Germany

Presentation Title: Using Minimal Arrival Curves to Derive Per-Flow Performance Bounds in (Tandem) Networks with Complex Feedback Structures

Abstract:

In [1], a way to extend network calculus by trading the strictness assumption of a server for a minimal arrival curve of the flow of interest has been presented. Here, the per-flow analysis of a small network with 2 nodes and a feedback structure was explored. We saw that the analysis is easier to employ and is more flexible w.r.t. scheduling policies using minimal arrival curve assumptions. We expand on this result to allow the derivation of per-flow performance bounds for tandem networks with a complex feedback structure. First, we isolate canonical feedback structures and derive open-loop representations for these structures with equivalent service curves. Then, we illustrate how these canonical structures can be applied to larger feedback structures, including nested and overlapping feedback loops. We show in which order the open-loop transformations have to be performed to maintain the existing feedback structures of the closed-loop system. In a numerical example, we illustrate how to transform a closed-loop tandem network into an equivalent open-loop one and derive per-flow performance bounds.

[1] "Extending network calculus to deal with min-plus service curves in multiple flow scenarios", Anja Hamscher, Vlad-Cristian Constantin, Jens B. Schmitt. IEEE 30th Real-Time and Embedded Technology and Applications Symposium (RTAS), 2024.

Presenter: Anja Hamscher, RPTU Kaiserslautern-Landau, Germany

(Back up presentation by Yuming if there is space.)

Presentation Title: Impact of Packetization on Network Calculus Analysis

Abstract:

This talk is focused on the impact of packetization on network calculus analysis. A simple single node case is considered. The discussion starts from the known result that the effect of packetizer can be ignored in delay bound analysis. It is also shown that this result is applicable even for the stochastic setting. The discussion is then extended to service curve and output characterization analysis. In contrast to that in delay bound analysis, the discussion will show that the effect of packetization on service curve characterization and output characterization cannot / should not be ignored.

Presenter: Yuming Jiang, NTNU, Norway