Report from Dagstuhl Seminar 15112

Network Calculus

Edited by

Florin Ciucu[1], Markus Fidler[2], Jörg Liebeherr[3], and Jens Schmitt[4]

1 **University of Warwick, GB,** florin@dcs.warwick.ac.uk
2 **Leibniz Universität Hannover, DE,** markus.fidler@ikt.uni-hannover.de
3 **University of Toronto, CA,** jorg@comm.utoronto.ca
4 **University of Kaiserslautern, DE,** jschmitt@cs.uni-kl.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 15112 "Network Calculus". At the seminar, about 30 invited researchers from academia and industry discussed the promises, approaches, and open challenges of the Network Calculus. This report gives a general overview of the presentations and outcomes of discussions of the seminar.

Seminar March 8–11, 2015 – http://www.dagstuhl.de/15112
1998 ACM Subject Classification C.2.1 Network Architecture and Design – Packet switching networks
Keywords and phrases Deterministic, Stochastic Network Calculus, Queueing Theory, Effective Bandwidth, Performance Evaluation
Digital Object Identifier 10.4230/DagRep.5.3.63
Edited in cooperation with Hao Wang, University of Kaiserslautern, DE

1 Executive Summary

Florin Ciucu
Markus Fidler
Jörg Liebeherr
Jens Schmitt

The network calculus has established as a versatile methodology for the queueing analysis of resource sharing based systems. Its prospect is that it can deal with problems that are fundamentally hard for alternative methodologies, based on the fact that it works with bounds rather than striving for exact solutions. The high modelling power of the network calculus has been transposed into several important applications for network engineering problems, traditionally in the Internet's Quality of Service proposals IntServ and DiffServ, and more recently in diverse environments such as wireless networks, sensor networks, switched Ethernets, Systems-on-Chip, as well as smart grids.

The goal of this Dagstuhl seminar was to gather the deterministic and stochastic network calculus community, to discuss recent research activities, to identify future research questions, and to strengthen cooperation. Topics of this Dagstuhl seminar were:

Wireless systems: for the analysis of wireless networks, a question of interest is how the stochastic properties of wireless channels impact delay and backlog performance. The usual statistical models for radio signals in a propagation environment do not lend themselves easily to a queueing model. Promising methods that were elaborated in the seminar are effective capacities and a recent network calculus of fading channels.

Lower bounds and tightness of bounds: based on the ability to solve some fundamentally hard queueing problems, the stochastic network calculus is regarded as a valuable alternative to the classical queueing theory. The derivation of performance bounds in the stochastic network calculus, e.g., for backlog, and delay, frequently exploits well known tail estimates, such as Chernoff bound and others. The tightness of these bounds and alternative more accurate models and techniques, such as Martingale bounds, were a topic of the seminar.

Network topology: a remarkable quality of the network calculus is that it includes a variety of systems that can be composed to arbitrary network topologies. Various analytical as well as numerical approaches have been explored to analyze different types of topologies, such as line topologies or feed-forward networks. The goal of this seminar was to identify relevant classes of topologies, their defining properties, and corresponding methods.

Parallel systems: the area of performance evaluation of parallel systems has recently become increasingly important due to the prevalence of modern parallel computational models. It is thus a great opportunity for the network calculus community to develop new models and methods which can enable a fundamental and broad understanding of the performance of parallel systems. At the seminar, recent approaches to parallel systems have been discussed.

Related methods: the network calculus has a number of rather unexplored and unexploited connections to related methods in the areas of competitive analysis, adversarial queueing theory, and robust queueing theory that may offer a significant potential for future research. At the seminar, researchers from related fields provided valuable new input to the network calculus community.

During the seminar, we discussed and (partly) answered the following questions:

What are the requirements on a wireless network calculus? Given the increasing importance of wireless communications, the seminar featured two sessions comprising seven presentations on wireless systems, where different approaches and their applications were discussed. Subsequently, a wireless roadmap discussion was centered around the following questions:

- How to model wireless channels and systems?
- What are the most relevant future systems and technologies?
- Which assumptions are needed, which can be safely made?
- What kind of results are needed, which theories can provide these?

With regard to the questions above, we highlight some of the main aspects that were elaborated on during the seminar. The methods that were presented include

- effective capacities,
- impairment models (duality with left-over service curves of scheduling),
- (min, x)-calculus for fading channels,
- capacity-delay-error boundaries,
- central limit theorem,
- Martingale bounds.

Providing different pros, a common basis of many of these methods was found to be due to the prevailing use of moment generating functions (Laplace transforms or Mellin transforms). Relevant systems that were discussed are cognitive radio, 3GPP, MIMO, spatial multiplexing, automatic repeat request, and medium access control. Some fundamental aspects of modelling wireless systems are the assumptions that are required today. Typical choices include

- service increments:
  - independent,
  - Markovian, Gilbert-Elliott channel,
- in-order delivery,
- error-free, instantaneous feedback channel,
- instantaneous retransmission of erroneous data,
- channel state information.

During the discussion, the need for transfer domains beyond Gilbert-Elliott models was raised. Also, the introduction of a notion of time into information-theoretic concepts, such as channel capacity, was discussed and finite-block length capacity results were brought up. Topics of further interest included spatial aspects of wireless networks, interference, and multi-hop networks in general. Regarding the solutions that can be obtained, a tradeoff between exactness and analytical closed forms became apparent. In particular, in system optimization analytical solutions were mentioned to be most useful to obtain derivatives of relevant performance measures. The discussion also touched upon some more general aspects such as qualitative vs. quantitative results, where many practical applications may not require exact results but can benefit from measurable rules of thumb.

What are most promising future research topics in the network calculus? This question was elaborated on in group work sessions, where the task was to identify an upcoming, relevant research topic where performance evaluation can be expected to make a key contribution. The discussion was guided by the following questions:

- What are the requirements for theory, which assumptions can be made?
- Which results would be needed from theory?
- How would a model/approach look like?
- What would be the best case outcome?
- Which body of theory could provide such results?
- What would be a good topic/method/approach for a PhD dissertation in this area?

Relevant topics in the network calculus were found to include cross-layer design, industrial communication, systems on chip, networks on chip, data center communication, and big data. A strategic orientation may also focus on new and unorthodox problems such as

- just-in-time manufacturing,
- renewable energy, smart grid,
- caching,
- financial engineering,
- road traffic,

where the intuitive concept of envelopes as used by the network calculus may be beneficial for many applications in industry. Methodological aspects that may pose relevant and interesting challenges were discussed in the areas of:

- re-entrant lines, particularly stability of such systems,
- max-min problems,
- derivative constraints, e.g., in modelling of batteries,
- network topologies, particularly non-feed forward networks.

Making network calculus happen: computational aspects, application modelling, tool support. Clearly, for network calculus to become a standard technique in performance modelling and analysis of networked and distributed systems it is crucial to arrive at computable solutions, demonstrate its strengths in diverse applications and provide software tools to support performance engineers in their daily tasks. As these different issues are interrelated on many levels two sessions with nine presentations were devoted to them. Among the different issues raised during these presentations and the corresponding discussions were the following:

- What are suitable novel application domains for network calculus? What are their requirements?
- How can network calculus computations be made more scalable? Where are fundamental limits for the network analysis? How do current software tools perform?
- What is the "killer" application for network calculus, and, in particular, for stochastic network calculus?
- How can network calculus' scope be extended to open up for new application domains?

Some (partial) answers to these important questions could be hinted at by the presentations and the subsequent discussions:

- Currently, some of the most promising application domains of (deterministic) network calculus were identified in industrial control, automotive and aerospace industries; also, interesting steps using (stochastic) network calculus in the modelling of smart energy grids were presented.
- The hardness of feedforward network analysis is by now understood, good heuristic approaches are on the way; however, cyclic dependencies and feedback systems are still open problems to some degree.
- The modelling of parallel systems using network calculus seems a promising building block to address novel attractive applications.
- Software tool support for network calculus, in particular for the stochastic version, is under construction and requires a community effort.

Looking over the fence: related methods. The research goals of network calculus and its methodologies, such as system performance evaluation, Markov chain analysis, or large deviations, intersect with those of other research communities. The objective of the session "Related Methods" was to create a forum where researchers from diverse research communities present their research approaches and discuss them with network calculus researchers. Thus, the session exposed the network calculus community to recent trends in system performance evaluation. Moreover, since speakers in this session had previously no or only limited exposure to network calculus, the session created an opportunity to disseminate the network calculus research agenda to other communities. The session was subtitled as "Looking over the fence", indicating an interest in learning new methodologies and the desire for cross- and interdisciplinary interactions. The session featured speakers from four countries (Canada, France, Israel, USA), from three disciplines (mathematics, theoretical computer science, operations research), presenting recent research on approaches on topics such as robust queueing theory, adversarial queueing theory, and competitive analysis.

This report provides an overview of the talks that were given during the seminar. Also, the seminar comprised a one minute madness session for introduction and for statements on the network calculus, a breakout session for group work on promising future research topics in the network calculus, as well as a podium discussions on wireless network calculus. The discussions, viewpoints, and results that were obtained are also summarized in the sequel.

We would like to thank all presenters, scribes, and participants for their contributions and lively discussions. Particular thanks go to the team of Schloss Dagstuhl for their excellent organization and support.

2 Table of Contents

3 Overview of Talks

3.1 Network Calculus for Parallel Processing

George Kesidis (The Pennsylvania State University, US)

We begin with an overview of classical Markovian results in fork-join queues and cloud-computing jargon. We then present preliminary results on the use of network calculus for parallel processing (fork join) systems such as MapReduce. We derive a probabilistic bound on delay through a single parallel processing stage. We also provide a numerical result using a publicly available dataset of a Facebook data-center that includes the total job arrival rate and workload statistics of the tasks of different types of MapReduce jobs at both the mapper and reducer stages. Finally, we discuss how to extend to tandem queues.

3.2 Wireless Network Calculus

Yuming Jiang (NTNU – Trondheim, NO)

In this talk, an overview of the difficulty, the key underlying issues and an overall picture of Wireless Network Calculus, i.e. extension/application of SNC to wireless networks, is first presented. This is followed by a brief introduction of our achieved research results in Wireless Network Calculus. The last part is devoted to the introduction of a fundamental problem in Wireless Network Calculus, which is end-to-end (e2e) QoS analysis of wireless networks where there is interference among neighbor hops. Some preliminary ideas to deal with this analysis are presented.

3.3 Energy Efficient Effective Capacity for 5G Networks

Eduard Jorswieck (TU Dresden, DE)

URL http://5glab.de/

In wireless 5G networks a paradigm change of services and applications to machine-to-machine low delay communications (tactile internet) requires to guarantee round trip times below 1–10 ms. On the other hand, the energy efficiency in 5G should be improved by a factor of 1000, too. Therefore, we propose a new performance metric which combines both conflicting objectives into the efficient effective capacity defined as the ratio of effective capacity to total consumed energy. The maximization of this metric leads to a fractional programming problem which can be solved efficiently by the Dinkelbach algorithm. The extension of the efficient effective capacity to the elements of multiuser networks, i.e., to the multiple access and broadcast channel is not available yet, because expressions for the effective capacity region are missing. In order to develop 5G networks with latency requirements/guarantees, we need to solve the following problems:

1. Derive the effective capacity region for multiple access channels.
2. Compute the effective capacity region for broadcast channels.
3. Derive the effective capacity for multihop (relaying) networks for different relaying protocols (amplify-and-forward, decode-and-forward, compress-and-forward, compute-and-forward, noisy network coding).

3.4 Effective Capacity – Through Physical and Data-Link Layers

Sami Akin (Leibniz Universität Hannover, DE)

Alongside the growth in social networks, mobile computing and pervasive communications, and the innovations in lower layer technologies, we see the need to re-visit network design strategies and develop better protocols. Can we design better higher layer strategies that inform, or are informed by, the underlying physical layer? With sufficient co-existence mechanisms, what novel cognitive radio network architectures are required? Hence, in this presentation, we discuss Effective Capacity from a physical layer perspective and investigate the effects of physical layer features on buffer performance in data-link layers by considering the cognitive radio framework as a working ground.

3.5 Performance of In-Network Processing for Visual Analysis in Wireless Sensor Networks

Hussein Al-Zubaidy (KTH Royal Institute of Technology, SE)

Joint work of Al-Zubaidy, Hussein; Gyorgy, Dan; Viktoria, Fodor
Main reference H. Al-Zubaidy, D. Gyorgy, F. Viktoria, "Performance of in-network processing for visual analysis in wireless sensor networks," in Proc. of the 14th IFIP Networking Conf. (Networking 2015), pp. 1–9, IEEE, 2015.
URL http://dx.doi.org/10.1109/IFIPNetworking.2015.7145292

Nodes in a sensor network are traditionally used for sensing and data forwarding. However, with the increase of their computational capability, they can be used for in-network data processing, leading to a potential increase of the quality of the networked applications as well as the network lifetime. Visual analysis in sensor networks is a prominent example where the processing power of the network nodes needs to be leveraged to meet the frame rate and the processing delay requirements of common visual analysis applications. The modelling of the end-to-end performance for such networks is, however, challenging, because in-network processing violates the flow conservation law, which is the basis for most queuing analysis. In this work we propose to solve this methodological challenge through appropriately scaling the arrival and the service processes, and we develop probabilistic performance bounds using stochastic network calculus. We use the developed model to determine the main performance bottlenecks of networked visual processing. Our numerical results show that an end-to-end delay of 2–3 frame length is obtained with violation probability in the order of 10^{-6}. Simulation shows that the obtained bounds overestimates the end-to-end delay by no more than 10%.

3.6 Capacity-Delay-Error Boundaries: A Composable Model of Sources and Systems

Nico Becker (Leibniz Universität Hannover, DE)

Joint work of Fidler, Markus; Lübben, Ralf; Becker, Nico
Main reference M. Fidler, R. Lubben, N. Becker, "Capacity-Delay-Error Boundaries: A Composable Model of Sources and Systems," IEEE Trans. onWireless Communications, 14(3):1280–1294, 2014.
URL http://dx.doi.org/10.1109/TWC.2014.2365782

It is presented a notion of capacity-delay-error (CDE) boundaries as a performance model of networked sources and systems. It is shown that the model has the property of additivity, which enables composing CDE boundaries obtained for sources and systems as if in isolation. Results for essential sources, channels and for the composition of sources and channels coders are presented.

3.7 Service-Martingales: Theory and Applications to the Analysis of Random Access Protocols

Felix Poloczek (TU Berlin, DE)

Joint work of Poloczek, Felix; Ciucu, Florin
Main reference F. Poloczek, F. Ciucu, "Service-martingales: theory and applications to the delay analysis of random access protocols," in Proc. of the 2015 IEEE Conf. on Computer Communications (INFOCOM'15), pp. 945–953, IEEE, 2015.
URL http://dx.doi.org/10.1109/INFOCOM.2015.7218466

We propose a martingale extension of effective capacity, a concept which has been instrumental in the teletraffic theory to model the link-layer wireless channel and to analyze QoS metrics. Together with a recently developed concept of an arrival-martingale the proposed *service-martingale* concept enables the queuing analysis of a bursty source sharing a MAC channel. In particular, we derive the first rigorous stochastic delay bounds for a Markovian source sharing either an ALOHA or CSMA/CA channel. By leveraging the powerful martingale methodology, the obtained bounds are remarkably tight.

3.8 Queuing Analysis of Wireless Systems: A Waste of Time?

James Gross (KTH Royal Institute of Technology, SE)

For some time now, there is a significant research activity with respect to queuing analysis of wireless systems based on effective capacity. These contributions follow a certain pattern: Identify what is hot in information theory and provide the corresponding queuing analysis. However, such contributions are limited by the additional insight they provide (in comparison to the original publication), while on the other hand the models are usually too theoretic to have practical value. In this talk I mainly illustrate these circumstances based on my own work, and intend to provoke discussions around the future value of queuing-related analysis of wireless systems. A few possible ways forward are finally presented, too.

3.9 SLA Calculus – Modelling Software Systems with Network Calculus

Peter Buchholz (TU Dortmund, DE)

Quantitative properties of modern software systems are often defined as part of a service level agreement (SLA) that fixes the maximal load and the maximal delay. Evaluation of the software system in order to validate the SLA is a challenging task since the system is to a large extend unknown and unpredictable. Thus, performance analysis has to be based on the SLAs without additional information about the basic system. The talk presents a new approach to analyze software architectures based on the ideas available in network-order real time calculus. In this way, bounds for departure processes are computed from available bounds for the arrival and delay processes. With the technique systems of composed services can be easily analyzed resulting in SLAs for the composed service. It is shown, how the solutions can be used to help a user and a provider to analyze and determine SLAs. Furthermore, open questions and limitations of the proposed approach are outlined.

3.10 Modelling Avionics Communicating Systems: Successes, Failures, Challenges

Marc Boyer (ONERA – Toulouse, FR)

This talk gave some perspectives on "the application modelling side, what is required from NC, what is still missing, what are success and failure stories". The talk presented how the modelling of AFDX has been done in an accurate way, whereas the one of SpaceWire has not. Thereafter, seven challenges on modelling are listed.

3.11 Industrial Application of Network Calculus

Kai-Steffen Jens Hielscher (Universität Erlangen-Nürnberg, DE)

Joint work of Herpel, Thomas; Hielscher, Kai-Steffen; Klehmet, Ulrich; German, Reinhard
Main reference T. Herpel, K.-S. Hielscher, U. Klehmet, R. German, "Stochastic and deterministic performance evaluation of automotive CAN communication," Computer Networks, 53(8):1171–1185, 2009.
URL http://dx.doi.org/10.1016/j.comnet.2009.02.008

In this talk we present the application of deterministic Network Calculus for two real-world examples: Communication of embedded controllers in automotive networks in cooperation with Audi AG and industrial Ethernet communication for industry automation in cooperation with Siemens AG. In the automotive example, the industry partner provided the topology and information about periodic CAN and FlexRay messages. The goal was to decide on which of the different busses inside a car interoperating electronic control units (ECUs) should be placed to avoid the violation of the hard real-time bounds. To achieve this, the CAN media access method had to be modelled in Network Calculus. Since the busses are interconnected by a central gateway, the service of this gateway also has to be modeled. Besides the scheduling strategy, this involved considering the aggregation of numerous interfering flows.

Industrial automation today mainly uses variants of industrial Ethernet like Profinet RT. Since these technologies do not provide guarantees like traditional field busses, our industry partner Siemens uses Network Calculus to calculate bounds for real-time traffic. For this purpose, they are integrating a Network Calculus Engine into their existing network planning tool. The tool already contains topology information and necessary information to generate arrival curves for scheduled flows. Other flows generated by user programs can be integrated by semi-automatic static code analysis. Since the end users often integrate hardware like web cams and HMI terminals into the network that generates non-real-time traffic, traffic profiles for these applications have been defined. To ensure that the limits provided in the profiles are not exceeded, traffic shaping has to be introduced into the network for the non-real-time flows.

References

1. T. Herpel, K.-S. Hielscher, U. Klehmet, and R. German. Stochastic and deterministic performance evaluation of automotive CAN communication, in *Computer Networks*, vol. 53, no. 8, pp. 1171–1185, 2009, Performance modelling of Computer Networks: Special Issue in Memory of Dr. Gunter Bolch.
2. S. Kerschbaum, K.-S. Hielscher, U. Klehmet, and R. German, A framework for establishing performance guarantees in industrial automation networks, in *Measurement, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance*, ser. Lecture Notes in Computer Science, K. Fischbach and U. Krieger, Eds., Springer International Publishing, 2014, vol. 8376, pp. 177–191.

3.12 On the Scalability of Real Time Calculus

Kai Lampka (Uppsala University, SE)

Joint work of Lampka, Kai; Perathoner, Simon; Suppiger, Urban; Thiele, Lothar
Main reference U. Suppiger, S. Perathoner, K. Lampka, L. Thiele, "A simple approximation method for reducing the complexity of modular performance analysis," Technical Report 329, ETH Zurich, August 2010.
Main reference http://www.tik.ee.ethz.ch/db/public/tik/?db=publications&form=report_single_publication&publication_id=3494

With Real Time Calculus and the related tool-support [8], it can be observed that the computation of the commonly used piece-wise linear pseudo-periodic functions, may require significant demands of computation and memory resources. The resulting overheads might render system analysis inefficient, if not infeasible, or often enforce simplifications, respectively overapproximations in the modelling. Simplifying overapproximations of signal frequencies or processing patterns, yield analysable models and guarantees conservativeness of results. However, it results in a non-tight bounding of performance metrics and ultimately yields potentially over-provisioned system designs. This shortcoming is the starting point for precise prefixing of bounding functions, as it only exploits overapproximations on the unneeded parts of functions. In order to achieve this the presentation presents the following innovations.

- The presentation introduces the concept of curve prefixing. This allows one to present curves precisely only on the interval $[0, c]$. For the range $(c, +\infty)$ the concept uses simplifying overapproximations which makes periodic tail descriptions of curves obsolete.

- The presentation formally establishes the framework for computing backlog and delay bound which are as tight as if one would have used the original curve representation. It thereby lifts the presenter's previous work in this direction from the level of an approximation method to the level of a precise analysis technique.
- The presentation contains an industrial, real-time constraint communication system. The system contains over 200 devices and integrates different real-time applications in a single (non-partitioned) architecture.

The concept of curve prefixing and tail overapproximations makes a clear distinction to today's implementations of Network or Real-time Calculus, e.g., as provided by the Matlab-based MPA-toolbox [5]. There, curve prolongation is the default behaviour, at each component the least common multiple of the periods of two input curves gives the period of the resulting output curve. The proposed approach therefore clearly increases the scalability of RTC-based system analysis as demonstrated by the industrial case study. But most importantly, it works on top of the existing tools and thereby avoids re-implementation of RTC.

References

1. J.-Y. L. Boudec and P. Thiran. Network Calculus: a theory of deterministic queuing systems for the Internet, volume 2050 of *LNCS*. Springer, 2001.
2. A. Bouillard and E. Thierry. An algorithmic toolbox for network calculus, in *Journal of Discrete Event Dynamic Systems (JDEDS)*, 18(1):3–49, 2008.
3. R. L. Cruz. A calculus for network delay. part i: Network elements in isolation and part ii: Network analysis, in *IEEE Transactions on Information Theory*, 37(1):114–141, January 1991.
4. R. Henia, A. Hamann, M. Jersak, R. Racu, K. Richter, and R. Ernst. System level performance analysis – the SymTA/S approach, chapter 2, pages 29–72, The Institution of Electrical Engineers, London, United Kingdom, 2006.
5. Modular performance analysis framework. http://www.mpa.ethz.ch.
6. U. Suppiger, S. Perathoner, K. Lampka, and L. Thiele. Modular performance analysis of large-scale distributed embedded systems: an industrial case study, Technical Report 330, ETH Zurich, November 2010.
7. U. Suppiger, S. Perathoner, K. Lampka, and L. Thiele. A simple approximation method for reducing the complexity of modular performance analysis, Technical Report 329, ETH Zurich, August 2010.
8. E. Wandeler, L. Thiele, M. Verhoef, and P. Lieverse. System architecture evaluation using modular performance analysis – a case study, in *International Journal on Software Tools for Technology Transfer*, 8(6):649–667, October 2006.

3.13 Network Calculus Tool Support – Expectations and Reality

Steffen Bondorf (University of Kaiserslautern, DE)

The first part of this talk will be covering the Disco Deterministic Network Calculator (DiscoDNC), an open-source network calculus tool [1].

Steffen Bondorf has been working with the network calculus tool support offered by the DISCO group for some time [2] before he eventually took over the role as its maintainer.

Since then, he has put effort into improving the tool in different aspects [3] – one of which is lowering the barrier to start working with deterministic network calculus. For that, the code has been restructured, the API reworked, functional tests have been created and their computations have been documented in detail.

This work led to several inquiries from researchers seeking to make use of network calculus results in order to evaluate their work. Unfortunately, there is a gap between the expectations that those researcher had regarding tool support and the reality at hand. Most notably, the preceding modelling step required to apply network calculus emerged as the single most problematic hurdle on the way towards deriving delay and backlog bounds. The DiscoDNC, however, strictly depends on the network calculus model, i.e., service curves and arrival curves need to be given in order to analyze a network.

In his talk, Steffen shares his experiences from being approached by academics making their first steps in the area of network calculus. He will depict common misconceptions along the lines of an example, showing that the effort to analyze a "simple" network with roughly 200 nodes can result in actually analyzing a so-called server graph connecting 1140 queues (servers) connected by nearly 7000 links. This observation motivates Steffen's work on improving the efficiency of network calculus analyses.

The second part of this talk will depict several enhancements to the computational efficiency of network calculus analyses. These improvements can be divided into two groups:

- Technical solutions allowing the DiscoDNC to derive bounds faster and
- Conceptual improvements in network calculus itself.

The former part covers the reuse intermediate results and the potential to parallelize the execution of a network analysis – both possible thanks to the modularity of (algebraic) deterministic network calculus.

The latter part will conclude the talk by providing some insight into an upcoming result [4] enabling to significantly reduce the analysis effort in sink trees with token-bucket arrival curves and rate-latency service curves.

References

1. The Disco Deterministic Network Calculator. http://disco.cs.uni-kl.de/index.php/projects/disco-dnc
2. S. Bondorf and J. Schmitt. Statistical Response Time Bounds in Randomly Deployed Wireless Sensor Networks, in *Proceedings of the 35th IEEE Conference on Local Computer Networks (LCN 2010).*
3. S. Bondorf and J. Schmitt. The DiscoDNC v2 – A Comprehensive Tool for Deterministic Network Calculus, in *Proceedings of the 8th International Conference on Performance Evaluation Methodologies and Tools (ValueTools 2014).*
4. S. Bondorf and J. Schmitt. Boosting Sensor Network Calculus by Thoroughly Bounding Cross-Traffic, in *Proceedings of the 34th IEEE International Conference on Computer Communications (INFOCOM 2015).*

3.14 Exact Delays in Networks

Anne Bouillard (ENS/INRIA, FR)

In this talk, we present a method based on linear programming to compute exact worst-case delay bounds under network calculus assumptions. We assume that the network is feed-forward; that the arrival/service curves are piecewise linear concave/convex and that the service policy is FIFO. The proposed method encodes every NC constraint into linear constraints, possibly with boolean variables. Then the solution of the LP is the exact worst-case delay. This algorithm is compared against existing method; derived into two simpler LPs that respectively compute good approximations of the upper bound and lower bound of the exact worst-case delay.

3.15 Optimal Joint Path Computation and Rate Allocation for Real-time Traffic

Giovanni Stea (University of Pisa, IT)

Computing network paths under worst-case delay constraints has been the subject of abundant literature in the past two decades. Assuming Weighted Fair Queueing scheduling at the nodes, this translates to computing paths and reserving rates at each link. The problem is NP-hard in general, even for a single path; hence polynomial-time heuristics have been proposed in the past, that either assume equal rates at each node, or compute the path heuristically and then allocate the rates optimally on the given path. In this paper we show that the above heuristics, albeit finding optimal solutions quite often, can lead to failing of paths at very low loads, and that this could be avoided by solving the problem, i.e., path computation and rate allocation, *jointly* at *optimality*. This is possible by modelling the problem as a mixed-integer second-order cone program and solving it optimally in split-second times for relatively large networks on commodity hardware; this approach can also be easily turned into a heuristic one, trading a negligible increase in blocking probability for one order of magnitude of computation time. Extensive simulations show that these methods are feasible in today's ISPs networks and they significantly outperform the existing schemes in terms of blocking probability.

3.16 How Can Network Calculus Help Smart Grids?

Yashar Ghiassi (University of Waterloo, CA)

This work is motivated by the challenges that arise when integrating large scale renewable energy integration. The significant fluctuations injected to the grid by renewable energy sources must be captured by storage systems. The role of storage in smart grids resembles the role of buffers and shapers in computer networks. We use this analogy to employ the buffer-overflow bounds from network calculus to size storage systems for given maximum loss of power and rate of power probabilities. This framework applies to a large range of applications in smart grids given that storage is an integral element in smart grids.

3.17 Computable Bounds in Fork-Join Queueing Systems

Amr Rizk (University of Massachusetts – Amherst, US)

Joint work of Rizk, Amr; Poloczek, Felix ; Ciucu, Florin
Main reference A. Rizk, F. Poloczek, F. Ciucu, "Computable Bounds in Fork-Join Queueing Systems," SIGMETRICS Perform. Eval. Rev., 43(1):335–346, ACM, 2015.
URL http://dx.doi.org/10.1145/2745844.2745859

A Fork-Join (FJ) queueing system is characterized by an upstream fork station that splits incoming jobs into N tasks to be further processed by N parallel servers, each with its own queue; the response time of one job is determined, at a downstream join station, by the maximum of the corresponding tasks's response times. FJ queueing systems help modelling multi-service systems subject to synchronization constraints. One prominent example are MapReduce clusters. In this work we provide first computable stochastic bounds on the waiting and response time distributions in FJ systems for renewal and non-renewal arrivals. Further, we consider blocking and non-blocking server behavior and prove that delays scale as $\mathcal{O}(\log N)$ in the non-blocking case, a law which is known for first moments under renewal input only. We show simulation results indicating that our bounds are tight, especially at high utilizations.

3.18 Window Flow Control in Network Calculus

Michael Beck (University of Kaiserslautern, DE)

This talk is concerned with the long-standing problem of feedback in Network Calculus (NC), in particular Stochastic Network Calculus (SNC). While there are plenty and elegant results on the deterministic side of NC, corresponding theorems are missing in SNC. This in turn limits the areas where SNC could be applied. In this talk – presenting preliminary work – the feedback-inequality in its original form and its connection to a Window Flow Controller is given. An overview follows, presenting the generalizations on the feedback-inequality. This is concluded with the solution to the feedback-inequality for the continuous-time and bivariate case. While this is a necessary step, it is not sufficient for a full analysis of a WFC. It provides, however, some insights, especially on a paradox behavior concerning dynamic window sizes, which perform worse compared to their static window counterparts. At last it is shown that under (very) strict assumptions an analysis of the stochastic WFC is possible.

3.19 Scaling Laws in the Network Calculus Bounds vs. Exact results

Almut Burchard (University of Toronto, CA)

In this talk, I described how exact results can be recovered from performance bounds in the stochastic network calculus in certain important limits. For example, it is well-understood that the output bound agrees, in the long-time limit, with the arrival rate; similarly, the exact fail decay of the backlog can be recovered from the delay bound. A more delicate question is the growth of end-to-end delays with the path length. For heavy-tailed and self-similar processes, such delays grow with a power-law, but the exact power is not known. I illustrated the importance of simple scaling laws for the evaluation of simulations.

3.20 Routing and Scheduling for Bursty Adversarial Traffic – Adversarial Queuing Theory

Adi Rosén (CNRS / University Paris-Diderot, FR)

In this talk we mainly consider the setting of Adversarial Queuing Theory.

The main part of the talk gives a simple, deterministic, local-control routing and scheduling protocol that applies to any network topology. This protocol guarantees that, for any input traffic for which stability is possible, stability is indeed achieved, and moreover the buffers at the nodes are polynomially-bounded as well as each packet has polynomially-bounded delivery time. This part of the talk is based on the paper [1].

This main part of the talk is complemented by a short (partial) survey of results in Adversarial Queueing Theory, as well as results on the achievable throughput in networks with fixed, bounded buffers. The latter results, compared to the results in Adversarial Queuing Theory, suggest that the question of stability and the question of throughput under bounded buffers are different questions with answers that do not relate to each other.

References

1 W. Aiello, E. Kushilevitz, R. Ostrovsky, and A. Rosén. Adaptive packet routing for bursty adversarial traffic, in *JCSS*, vol. 60, no. 3, pp. 482–509, 2000.

3.21 Managing Queues with Bounded Buffers: Micro-decisions from a Competitive Lens

Gabriel Scalosub (Ben Gurion University – Beer Sheva, IL)

Network Calculus has traditionally assumed flow conservation, and that no traffic is lost while traversing the network. In consequence, it has primarily focused on understanding the performance of systmes in terms of delay, and provisioning of capacity. Nevertheless, packet loss is a feature of common networks, most predominantly the Internet, where packets are dropped due to buffer overflows, congestion control mechanisms, and and ever growing demand for more bandwidth which is not always available. In this talk we present models and algorithms for dealing with such packet loss, focusing primarily on buffer-management mechanisms. These results are cast within a competitive framework, which subscribes to other related models, such as AQT. We present both classic results in this framework, as well as some more recent results, and emphasize the characteristics of performing buffer-management, which are sometimes counter-intuitive, and sometimes lead to surprising results. Among other things, this talk may also serve as a teaser for Network Calculus to try and incorporate packet loss (and working with bounded buffers) into its framework.

3.22 Robust Queueing Theory

Nataly Youssef (MIT – Cambridge, US)

Joint work of Bandi, Chaithanya; Bertsimas, Dimitris; Youssef, Nataly
Main reference C. Bandi, D. Bertsimas, N. Youssef, "Robust Queueing Theory," Operations Research, 63(3):676–700, 2015.
URL http://dx.doi.org/10.1287/opre.2015.1367

We propose an alternative approach for studying queues based on robust optimization. We model the uncertainty in the arrivals and services via polyhedral uncertainty sets which are inspired from the limit laws of probability. Using the generalized central limit theorem, this framework allows to model heavy-tailed behavior characterized by bursts of rapidly occurring arrivals and long service times. We take a worst-case approach and obtain closed form upper bounds on the transient and steady-state system time in multi-server queues and feedforward networks. These expressions provide qualitative insights which mirror the conclusions obtained in the probabilistic setting for light-tailed arrivals and services and generalize them to the case of heavy-tailed behavior. We also develop a calculus for analyzing a steady-state network of queues based on the following key principle: (a) the departure from a queue, (b) the superposition, and (c) the thinning of arrival processes have the same uncertainty set representation as the original arrival processes. The proposed approach (a) yields results with error percentages in single digits relative to simulation, and (b) is to a large extent insensitive to the number of servers per queue, network size, degree of feedback, traffic intensity, and somewhat sensitive to the degree of diversity of external arrival distributions in the network.

4 Working Groups

4.1 Working group A

Steffen Bondorf

The working group started with discussing recent developments combining network calculus (NC) with optimization. In deterministic network calculus (DNC), work in this area started as early as 2008 when the basic problem of existing algebraic tandem analyses was identified.

It was suggested to derive an optimization problem from the network calculus "constraints" to solve it. Since then, the optimization-based analysis has been advanced to ultimate tightness, i.e., the best results possible with the given NC constraints, and has been extended to encompass the entire network instead of a tandem of servers only. In the stochastic network calculus (SNC) branch, work recently suggested to make use of its modelling capabilities in robust optimization. Thus creating a robust queueing theory.

These developments were caused by deficiencies in network calculus that are not easy to overcome. The discussion identified the following major issues:

- Lack of decision variables: Being restricted to the analysis of systems, NC relies on the provision of an exact model. It lacks the capability to directly support system engineering by finding assignments for open parameters such that a given requirement is fulfilled. Complementary methodologies as add-ons can help NC to increase its applicability in this area.
- Bounds instead of exact results: Network calculus itself is concerned with bounding a performance indicator instead of providing an exact result. Quality of bounds is a problem of both branches, DNC and SNC.
- Computational effort: Moreover, the computational effort involved in deriving bounds can be very high. For example, in the DNC analysis of tandems of FIFO multiplexing servers can already be very involved, yet, it is not ultimately tight. Although NC only derives bounds, oftentimes an additional tradeoff is required to derive results at all – especially in the analysis of reasonably sized networks.

The working group then turned to the aspect thought to be the common cause of the above problems: The model used for network calculus. From the beginning, i.e., Cruz' first papers, its simplicity was considered as defining NC's beauty. NC does not take many assumptions into account whereas classic performance tend to have too many to keep track of all of them properly. However, in order to overcome the problems identified in the discussion, the model may be considered simplistic. It is, e.g., even simpler than visual models for simulation as used by OMNeT++ or others. Summarizing, this kind of beauty defines the limitations of network calculus as well. The constrained expressiveness hampers the ambition to derive better bounds while simultaneously not allowing for fast and easy derivations either.

Given that the NC model's expressiveness currently seems to restrict leaps forward, the question about creating a potentially better model appeared. The group members asked themselves if it was possible to take such a disruptive step that incorporates the knowledge and experience the community generated over the past years. I.e., can we redesign the calculus such that its main deficiencies will be gone for good?

4.2 Working group B

Yashar Ghiassi

We started the meeting by discussing the application domains of network calculus. We classified the applications to two groups: emerging applications and traditional ones. Examples of emerging applications are vehicular transportation, energy systems (battery and EVs), financial engineering, and inventory control and manufacturing systems. Examples of traditional applications are communication networks and computation networks (e.g., cloud, embedded).

In the second half of our meeting we tried to discuss possible interesting problems (in each of the applications listed above) for which network calculus is helpful. The first interesting set of problems are facility location and dynamic topology; e.g., how do we optimally size and locate storage systems in smart grids? As another important and open problem we discussed feedback networks problems and the possibility that network calculus extends to that area of research. Routing algorithms was the third possible research direction that we discussed. Finally, we discussed a series of control related problems: state-dependent scheduling, traffic lights/signals, avoiding underflow (finance), and ramp limitations (electricity).

4.3 Working group C

Amr Rizk

The main focus of the discussion within this working group was on identifying requirements for advancing the Network Calculus (NC), as well as, main technical problems that are not solved (yet!) in the NC framework. We identified two pillars that would help the advancement of Network Calculus in the sense of increasing the user, as well as, the researcher community, i.e., (i) bringing NC to standardization and (ii) teaching NC at a graduate level. One success story of a related performance evaluation research topic that made a key contribution through the transition to (IETF) standardization is fair scheduling. Hence, it is of utmost importance to visualize the impact of the NC framework with implementations and case studies of actual deployment. A particular strength of NC lies in providing fundamental characterizations (limits) on basic networking elements that can be compiled into complex communication scenarios. We believe that a collection of such results with appropriate deployment examples would be very instructive for potential adoption. However, there are still basic elements and protocols that do not lend themselves to the (stochastic) Network Calculus, such as feedback and lossy systems. The conclusion of the discussion was that there are still many open challenges/problems to be solved within the Network Calculus framework.

5 Seminar Programme

Monday	
09:00-09:30	Welcome and general introduction
09:30-10:30	One minute madness: introduction of participants
11:00-12:00	Seed talk: George Kesidis
14:00-15:30	Wireless network calculus Yuming Jiang Eduard Jorswieck Sami Akin Hussein Al-Zubaidy
16:00-17:00	Wireless network calculus Nico Becker Felix Poloczek James Gross
17:00-17:45	Wireless network calculus: roadmap discussion
evening	Network calculus pub quiz

Tuesday	
09:00-10:00	Seed talk: Peter Buchholz
10:00-12:00	Group work: future network calculus topics
14:00-15:30	Making network calculus happen: computational aspects application modelling and tool support (CAT) Marc Boyer Kai-Steffen Hielscher Kai Lampka Steffen Bondorf
16:00-17:45	Making network calculus happen: computational aspects application modelling and tool support (CAT) Anne Bouillard Giovanni Stea Yashar Ghiassi-Farrokhfal Amr Rizk Michael Beck

Wednesday	
09:00-10:30	Looking over the fence: related methods Almut Burchard Adi Rosen Gabriel Scalosub Nataly Youssef
11:00-12:00	Feedback from group work
12:00-12:15	Seminar resume and farewell

Participants

- Sami Akin
Leibniz Univ. Hannover, DE
- Hussein Al-Zubaidy
KTH Royal Institute of Technology, SE
- Michael Beck
University of Kaiserslautern, DE
- Nico Becker
Leibniz Univ. Hannover, DE
- Daniel Berger
University of Kaiserslautern, DE
- Steffen Bondorf
University of Kaiserslautern, DE
- Anne Bouillard
ENS – Paris, FR
- Marc Boyer
ONERA – Toulouse Research Center, FR
- Peter Buchholz
TU Dortmund, DE
- Almut Burchard
University of Toronto, CA
- Florin Ciucu
University of Warwick, GB
- Markus Fidler
Leibniz Univ. Hannover, DE
- Reinhard German
Univ. Erlangen-Nürnberg, DE
- Fabien Geyer
TU München, DE
- Yashar Ghiassi-Farrokhfal
University of Waterloo, CA
- James Gross
KTH Royal Institute of Technology, SE
- Kai-Steffen Jens Hielscher
Univ. Erlangen-Nürnberg, DE
- Yuming Jiang
NTNU – Trondheim, NO
- Eduard Jorswieck
TU Dresden, DE
- George Kesidis
Pennsylvania State University – University Park, US
- Kai Lampka
Uppsala University, SE
- Jörg Liebeherr
University of Toronto, CA
- Krishna S. Pandit
TU Darmstadt, DE
- Felix Poloczek
TU Berlin, DE
- Amr Rizk
University of Massachusetts – Amherst, US
- Adi Rosén
CNRS / Univ. Paris-Diderot, FR
- Gabriel Scalosub
Ben Gurion University – Beer Sheva, IL
- Jens Schmitt
University of Kaiserslautern, DE
- Giovanni Stea
University of Pisa, IT
- Hao Wang
University of Kaiserslautern, DE
- Nataly Youssef
MIT – Cambridge, US