10051 Abstracts Collection
Quantitative and Qualitative Analysis of Network Protocols
— Dagstuhl Seminar —

Bengt Jonsson[1], Jörg Kreiker[2] and Marta Kwiatkowska[3]

[1] University of Uppsala, SE
Bengt.Jonsson@it.uu.se
[2] TU München, DE
joba@model.in.tum.de
[3] University of Oxford, GB
Marta.Kwiatkowska@comlab.ox.ac.uk

Abstract. From Jan 31, 2010 to Feb 5, 2010, the Dagstuhl Seminar 10051 "Quantitative and Qualitative Analysis of Network Protocols " was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Network protocols, verification, static analysis, quantitative modeling and model checking, graph transformation, process calculi

10051 Executive Summary – Quantitative and Qualitative Analysis of Network Protocols

This is the executive summary for the seminar Quantitative and Qualitative Analysis of Network Protocols held from Jan 31, 2010 to Feb 5, 2010.

Keywords: Executive summary

Joint work of: Jonsson, Bengt; Kreiker, Joerg; Kwiatkowska, Marta

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2010/2516

Dagstuhl Seminar Proceedings 10051
Quantitative and Qualitative Analysis of Network Protocols
http://drops.dagstuhl.de/opus/volltexte/2010/2518

Computing the Leakage of Information-Hiding Systems

Miguel E. Andres (Radboud University Nijmegen, NL)

We address the problem of computing the information leakage of a system in an efficient way. We propose two methods: one based on reducing the problem to reachability, and the other based on techniques from quantitative counterexample generation. The second approach can be used either for exact or approximate computation, and provides feedback for debugging. These methods can be applied also in the case in which the input distribution is unknown. We then consider the interactive case and we point out that the definition of associated channel proposed in literature is not sound. We show however that the leakage can still be defined consistently, and that our methods extend smoothly.

Keywords: Specification and verification techniques for finite state systems, Software verification, Model-checking, Information theory, Anonymity, probabilistic automata

Joint work of: Andres, Miguel E.; Palamidessi, Catuscia; van Rossum, Peter; Smith, Geoffrey

Probabilitic automata over infinite words

Christel Baier (TU Dresden, DE)

Probabilistic ω-automata are variants of nondeterministic automata for infinite words where all choices are resolved by probabilistic distributions. Acceptance of an infinite input word can be defined in different ways: by requiring that (i) the probability for the accepting runs is positive (probable semantics), or (ii) almost all runs are accepting (almost-sure semantics), or (iii) the probability measure of the accepting runs is greater than a certain threshold (threshold semantics).

The underlying notion of an accepting run can be defined as for standard ω-automata by means of a Büchi condition or other acceptance conditions, e.g., Rabin or Streett conditions.

In this paper, we put the main focus on the probable semantics and provide a summary of the fundamental properties of probabilistic ω-automata concerning expressiveness, efficiency, and decision problems.

Keywords: Büchi automata, probabilistic automata, expressiveness, efficiency

Mean-Field Analysis for Large Dynamic Networks

Rena Bakhshi (Vrije Universiteit - Amsterdam, NL)

We consider large dynamic networks, in which nodes move around and interact with each other in a gossip or peer-to-peer fashion.

In this talk, we present an abstraction method for the performance evaluation of such dynamic systems, called the mean-field approximation. This method allows us to easily evaluate systems with very large numbers of nodes, that is, systems of such a size that traditional performance evaluation methods would fall short. The talk covers some results we obtain by applying this method, and by automating it.

Keywords: Large-scale probabilistic systems, gossip protocols, mean-field approximation

Verification of Graph Transformation Systems: An Unfolding-based Approach

Paolo Baldan (University of Padova, IT)

With the advent of mobile and ubiquitous computing, modern software and computer systems are frequently characterised by a high level of dynamicity. Features such as flexible topologies, the dynamic creation and deletion of objects, and an infinite-state space make them very hard to analyse and verify. In this context, graph transformation systems (GTSs) emerge as an expressive specification formalism for concurrent, distributed and mobile systems, generalising another classical model of concurrency, namely Petri nets. While the semantic theory of GTSs is relatively well-understood, the problem of developing automatic verification techniques for GTSs is still a stimulating and promising research direction.

In this presentation we outline a framework for the verification of infinite and finite-state GTSs based on their unfolding semantics.

For general, possibly infinite-state, GTSs one can construct finite under- and over- approximations of the (infinite) unfolding, with arbitrary accuracy. Such approximations can be used to check behavioural properties of a GTS, expressed in a suitable temporal graph logic. For finite-state GTSs, we propose a variant of McMillan's complete prefix approach (originally developed for Petri nets), discussing some issues related to the construction of the prefix and its use.

Keywords: Graph transformation, Petri nets, unfolding

Joint work of: Baldan, Paolo; König, Barbara

Software Verification for TinyOS

Doina Bucur (University of Oxford, GB)

We provide a first verification tool for applications written for the mainstream sensor operating system, TinyOS. We statically verify a TinyOS application against both standard programming errors (e.g., memory violations) and higher-level safety specifications written in assertion form.

Keywords: Sensors, software verification, TinyOS

Joint work of: Bucur, Doina; Kwiatkowska, Marta

Depletable Channels: Dynamics and Behaviour

Pietro Cenciarelli (University of Rome "La Sapienza", IT)

A simple model of multi-hop communication in ad-hoc networks is considered. Similar models are often adopted for studying energy efficiency and load balancing of different routing protocols. We address an orthogonal question never considered by the networking community: whether, regardless of specific protocols, two networks may be considered as equivalent from the viewpoint of the communication service they provide.

In particular, we consider equivalent two networks with identical maximum and minimum inhibiting flow, and prove that this notion of equivalence coincides with a standard trace-based notion of equivalence borrowed from the theory of concurrency. We study the computational complexity of the proposed equivalence and discuss possible alternatives. Relationship with the Braess paradox is duscussed.

Joint work of: Cenciarelli, Pietro; Gorla, Daniele; Salvo, Ivano

Full Paper:
http://www.springerlink.com/content/04614t3572307472/

Revising probabilistic model checking for distributed systems

Pedro R. D'Argenio (National University - Cordoba, AR)

Probabilistic model checking computes the probability values of a given property quantifying over all possible schedulers (also known as adversaries or policies). It turns out that maximum and minimum probabilities calculated in such a way are overestimations on models of distributed systems in which components are loosely coupled and share little information with each other (and hence arbitrary schedulers may result too powerful). In this talk I will revise the idea of schedulers in distributed system and propose a definition. I will also discuss decidability issues and algorithmic results.

Keywords: Probabilistic model-checking, distributed systems, minimal and maximal probabilities

Analysis and Implementation of CaPiTo

Han Gao (Technical University of Denmark, DK)

We show an approach, CaPiTo, to model service-oriented applications using process algebras such that, on the one hand, we can achieve a certain level of abstraction without being overwhelmed by the underlying implementation details and, on the other hand, we respect the concrete industrial standards used for implementing the service-oriented applications. By doing so, we will be able to not only reason about applications at different levels of abstractions, but also to build a bridge between the views of researchers on formal methods and developers in industry. We apply our approach to the financial case study taken from EU project Sensoria. Finally, we develop a static analysis to analyse the security properties as they emerge at the level of concrete industrial protocols.

Keywords: Process Algebra, Protocol, Static Analysis

Joint work of: Gao, Han; Nielson, Flemming; Nielson, Hanne Riis

Reliability Modelling of Disk Subsystems with Probabilistic Model Checking

Kanchi Gopinath (Indian Inst. of Science - Bangalore, IN)

We discuss how probabilistic model checking, a variant of model checking, can be useful in modelling multiple failures in disk subsystems, as it allows specification of probabilities or rates of transitions also. Probabilistic model checking can be used to not only check the correctness of formulae that are augmented with a probability operator, but also to compute, as a side effect, probabilities and Markov rewards in response to queries expressed in an expressive logic. With exponential distributions for failures, it is possible to accurately compute the reliability of many non-trivial disk systems. Exploiting the technique of symmetry reduction, it is possible to handle even larger fully symmetric systems. For non-exponential distributions, such as Weibull models of disk reliability, we can analyze reasonable sized systems (such as RAID4/5/6) if good approximations based on exponomials are possible. We report our experiences and compare our results with those from simulation.

Keywords: Probabilistic model checking, disk storage systems, continuous time Markov chains, RAID, Weibull models

Tutorial on Qualitative Properties of Protocols and Process Algebras

Daniele Gorla (University of Rome "La Sapienza", IT)

Process algebras have longly been used for formalizing distributed systems, network protocols and their properties. They are usually equipped with a number of verification techniques, ranging from type systems to behavioural equivalences. General purpose process algebras (like CCS and pi-calculus) are often used, but new and more specilized process algebra have appeared in the last decade to better model process distribution and mobility, dynamically evolving network links, cryptography, etc. In this tutorial we shall give some hints and samples on this research line, mostly focusing on security protocols and the verification of some classical properties (viz., authentication and secrecy).

Did we do the right things to serve the networking systems community?

Boudewijn Haverkort (University of Twente, NL)

Both the verification and the performance evaluation community have developed very good techniques over the last 40 years. Also the networking community has developed enormoulsy, and has given us great systems, that billions of people are using nowadays. However, is there any reason to beliive that the performance and verification community has contributed to that? To be honest, I think this contribution has been negligible. Time for some introspection: did we do the right things?

Keywords: Performance, dependability, verification, networking

Beyond PRISM

Holger Hermanns (Universität des Saarlandes, DE)

The talk discusses some variations of the PRISM language, and discusses how the resulting richer models can be analysed by mechanised tool support. We cover infinite state models (introducing PASS and INFAMY), parametric models (introducing PARAM) and timed probabilistic model (introducing mcpta).

Keywords: Probabilistic model checking, infinite state models, abstraction refinement, parametric models, Markov chains, Markov decision processes

Joint work of: Hahn, E. Moritz; Hartmanns, Arnd; Wachter, Björn; Zhang, Lijun

Graph Grammar Modeling and Verification of Ad Hoc Routing Protocols

Bengt Jonsson (University of Uppsala, SE)

We present a technique for modeling and automatic verification of network protocols, based on graph transformation. It is suitable for protocols with a potentially unbounded number of nodes, in which the structure and topology of the network is a central aspect, such as routing protocols for ad hoc networks. Safety properties are specified as a set of undesirable global configurations.

We verify that there is no undesirable configuration which is reachable from an initial configuration, by means of symbolic backward reachability analysis. In general, the reachability problem is undecidable. We implement the technique in a graph grammar analysis tool, and automatically verify several interesting nontrivial examples. Notably, we prove loop freedom for the DYMO ad hoc routing protocol. DYMO is currently on the IETF standards track, to potentially become an Internet standard.

Keywords: Graph grammars, ad hoc networks, routing protocols, formal verification, model checking

Joint work of: Mayank Saksena, Oskar Wibling, Jonsson, Bengt

Full Paper:
http://dx.DOI.org/10.1007/978-3-540-78800-3_3

Tutorial on Graph Transformation

Barbara Koenig (Universität Duisburg-Essen, DE)

One potential method to verify network protocols is to model the protocol by graph transformation and to apply techniques for the verification of graph transformation systems. Such systems are well-suited to model systems in a natural way with a high degree of dynamicity. This talk gives an introduction to graph transformation and continues with an overview over existing verification techniques. We distinguish between methods for the verification of finite-state and infinite-state graph transformation systems and briefly consider the relevant logics.

API Authentication and Authorization Protocols

David B. Lee (Ohio State University, US)

API authentication and authorization is a process of multi-party authentication and authorization through Application Programming Interfaces (API).

A general and formal API authentication and authorization protocol model is developed from commercially available protocols, including OAuth, AOL OpenAuth, Yahoo! BBAuth and Flickr Authentication API. Security properties are discussed with OAuth as a case study.

Keywords: Authentication, Authorization, Application Programming Interface, Security Properties

Joint work of: Lee, David; Hsu, Yating

Wanted: Formal methods for Adaptive Routing in MANETs

Annabelle McIver (Macquarie University - Sydney, AU)

Crucial to the performance of MANETs is route-discovery; unfortunately there is a trade off between the amount of information that should be disseminated by protocols for route-discovery and the quality of the resulting link decisions. Typically the outcome for the user is poor quality performance.

Emerging protocols attempt to overcome this problem by using different metrics to evaluate route quality, varying them depending on the operating context. The range of metrics with which to instrument these novel strategies is immense and a key research challenge is to identify "optimal" metrics, namely those which will deliver the best overall performance rating.

This raises an exciting opportunity for formal methods: to contribute to that identification. In this short talk I shall review some of the problems faced by protocol designers in this important research area.

Keywords: MANETs, performance metrics

A Chemical Approach to Network Protocol Design and Analysis

Thomas Meyer (Universität Basel, CH)

A Chemical Approach to Network Protocol Design and Analysis Thomas Meyer and Christian Tschudin University of Basel, Switzerland

In this talk, we show how to design network protocols based on molecule-like entities such that the corresponding execution flows can by analyzed as if they were chemical processes. We introduce the metaphor of chemical networking protocols and highlight its benefits by a formal convergence and stability analysis of a gossip-style protocol.

In a second part we demonstrate self-healing programs that treat the loss of code fragments in a similar way than networking protocols recover from packet loss. As a first example, we present a self-healing load balancing protocol that is resilient to code removal attacks. Our long term goal is to enable programs to run in unreliable execution environments.

Keywords: Chemical Networking, Fraglets, Gossiping, Stability, Load Balancing

Joint work of: Meyer, Thomas; Tschudin, Christian
Full Paper:
http://conferences.sigcomm.org/hotnets/2009/papers/hotnets2009-final66.pdf

See also: Thomas Meyer and Christian Tschudin: "Chemical Networking Protocols", Proceedings of the 8th ACM Workshop on Hot Topics in Networks (HotNets-VIII), New York, NY, October 22-23, 2009.

Characterising Temporal Distance and Reachability in Mobile and Online Social Networks

Mirco Musolesi (University of St Andrews, GB)

The analysis of social and technological networks has attracted a lot of attention as social networking applications and mobile sensing devices have given us a wealth of real data. Classic studies looked at analysing static or aggregated networks, i.e., networks that do not change over time or built as the results of aggregation of information over a certain period of time. Given the soaring collections of measurements related to very large, real network traces, researchers are quickly starting to realise that connections are inherently varying over time and exhibit more dimensionality than static analysis can capture.

We propose new temporal distance metrics to quantify and compare the speed (delay) of information diffusion processes taking into account the evolution of a network from a local and global view. We also characterise network reachability with the concepts of in- and out-components. We show how these metrics are able to capture the temporal characteristics of time-varying graphs, such as delay, duration and time order of contacts (interactions), compared to the metrics used in the past on static graphs.

Full Paper:
http://www.cs.st-andrews.ac.uk/ mirco/papers/CCR10.pdf

Keywords: Temporal distance, mobile networks, social networks

Formal Models of Wireless Networks

Sebastian Nanz (ETH Zürich, CH)

In this talk I'll give an overview of recent approaches to modeling and analyzing wireless networks using specialized process algebras. In particular, I will present the process calculus CBS# which has been applied to the specification and security analysis of mobile ad-hoc networks.

Keywords: Wireless networks, process calculi, static analysis

Formal Approaches to Information Protection

Catuscia Palamidessi (Ecole Polytechnique - Palaiseau, FR)

The protection of confidential information is an issue of increasing importance in the modern world.

In this talk, we will discuss the problem and give an overview of various approaches to its formalization: possibilistic, probabilistic, and information-theoretic. We will compare the various frameworks and their corresponding verification techniques, and point out the relations between them. We will illustrate these concepts with some concrete examples of information-hiding protocols and of their specification and verification.

Joint work of: Palamidessi, Catuscia; Chatzikokolakis, Konstantinos; Panangaden, Prakash

Assume-Guarantee Verification for Probabilistic Systems

David Parker (University of Oxford, GB)

Probabilistic automata are a natural formalism for modelling and verification of network protocols. Scalability of these techniques, however, represents a major challenge. We present a compositional verification technique for such systems, based on assume-guarantee reasoning. We model assumptions and guarantees as probabilistic safety properties, represented by finite automata. Our techniques are efficient and fully automated, based on a reduction to the problem of multi-objective probabilistic model checking.

Using a prototype tool, we illustrate probabilistic verification of several large case studies, including instances where conventional, non-compositional approaches are infeasible.

Keywords: Assume-guarantee reasoning, probabilistic model-checking

Policies for Self Tuning Home Networks

Dimosthenis Pediaditakis (Imperial College London, GB)

A home network (HN) is usually managed by a user who does not possess knowledge and skills required to perform management tasks. When abnormalities are detected, it is desirable to let the network tune itself under the direction of certain policies. However, self tuning tasks usually require coordination between several network components and most of the network management policies can only specify local tasks. In the context of my current work, I propose a state machine based policy framework to address the problem of fault and performance

management in the context of HN. Policies can be specified for complex management tasks as global state machines which incorporate global system behaviour monitoring and reactions. We demonstrate the policy framework through a case study in which policies are specified for dynamic selection of frequency channel in order to improve wireless link quality when interferences present.

Keywords: Home networks, policies, monitoring, state machines, link quality

Full Paper:
http://www.computer.org/portal/web/csdl/doi/10.1109/POLICY.2009.30

Probabilistic timed automata: an overview

Jeremy Sproston (University of Torino, IT)

I will give an overview of probabilistic timed automata, a formalism for modelling systems which exhibit nondeterministic, probabilistic and timed behaviour. Probabilistic timed automata have been used to model formally randomised, timed protocols. Model-checking techniques for probabilistic timed automata will be considered, and recent work will be surveyed.

Keywords: Probabilistic timed automata, model-checking

Spotlight Abstraction of Agents and Areas

Tobe Toben (OFFIS - Oldenburg, DE)

We present "spotlight abstraction" as a generic abstraction technique for the analysis of systems comprising an unbounded number of communicating agents.

The abstraction principle is heterogeneous in the sense that the behaviour of a finite number of agents is preserved while the others are only abstractly represented. The precision of the abstraction can be tuned by an iterative procedure based on the analysis of counterexamples.

Going beyond existing work, we show how to use the spotlight principle for analysing systems where the physical position of agents is relevant. To this end, we put the spotlight on areas rather than on fixed sets of agents.

Keywords: Spotlight Abstraction, Verification, Dynamic Communication Systems

Joint work of: Toben, Tobe; Westphal, Bernd; Rakow, Jan-Hendrik

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2517

Modelling Clock Synchronization in the Chess gMAC WSN Protocol

Frits Vaandrager (Radboud University Nijmegen, NL)

We present a detailled timed automata model of the clock synchronization algorithm that is currently being used in a wireless sensor network (WSN) that has been developed by the Dutch company Chess. Using the Uppaal model checker, we establish that in certain cases a static, fully synchronized network may eventually become unsynchronized if the current algorithm is used, even in a setting with infenitesimal clock drifts.

Keywords: Wireless sensor networks, clock synchronization, timed automata

Full Paper:
http://arxiv.org/abs/0912.1901v1

Optimisation and model checking applied to multipath routing

Nigel Walker (BT - Suffolk, GB)

Optimisation theory has a well established pedigree in networking, naturally capturing not only planning problems, but also control problems such as routing, flow control, overload control, and scheduling. The notion of optimisation decomposition, in which a top level problem is solved by a collection of interacting sub-problems, leads to distributed algorithms and protocols (such as the Bellman-Ford algorithm for shortest path routing). Optimisation theory augmented with different notions of decomposition is therefore a relevant formal framework for studying network behaviour and network-wide properties.

On the other hand, formal methods based on logic seem quite different in character, even though they are designed to answer detailed questions about system behaviour. We will report on our experiences of placing model checking alongside optimisation theory in analysing the properties of a simple multipath routing system, and will discuss some of the issues raised by this experiment.

Keywords: Optimisation, Network Protocols, Model Checking

Joint work of: Walker, Nigel; Wu, Peng; Lomuscio, Alessio

Automated Theorem Proving For Network Analysis

Christoph Weidenbach (MPI für Informatik - Saarbrücken, DE)

I will present three formalizations of increasing expressiveness to the analysis of networks based on first-order logic.

The first one abstracts from time and probability resulting in a "nice" fragment of first-order logic. We have successfuly used this approach to fully automatically analyze aspects of the overall network infrastructure of the Max Planck Institute for Informatics.

The second one extends the first one by considering time in addition.

For the resulting fragment of first-order logic extended by arithmetic we have developed sound and often complete automated theorem proving methods. We applied those to fully automatically analyze aspects of network protocols.

The third one extends the second one by considering probability in addition.

Here we have meanwhile developed a tool chain eventually resulting in sound and often complete methods for the analysis of network protocols.

Keywords: Theorem Proving, Network Analysis

Probabilistic Live Sequence Charts

Bernd Westphal (Universität Freiburg, DE)

Live Sequence Charts (LSC) are a formally rigorous variant of Message Sequence Charts which provide modalities to distinguish between possible and mandatory system behaviour within scenarios.

We propose to extend the given qualitative modalities to quantitative ones. Our extension is supposed to yield a probabilistic variant of LSCs called pLSC which shall allow for an intuitive visual specification of requirements for probabilistic systems.

We discuss design decisions regarding the formal semantics of pLSCs with respect to Markov Decision Processes and outline an approach for the automatic decision of the satisfaction relation between pLSC and MDP.

Keywords: Visual Formalism, LSC, MSC, Quantitative Requirements

Joint work of: Westphal, Bernd; Rakow, Jan; Toben, Tobe; Wischmeyer, Samuel

A Secure Key Establishment Protocol for ZigBee Wireless Sensor Networks

Ender Yueksel (Technical University of Denmark, DK)

ZigBee is a wireless sensor network standard that defines network and application layers on top of IEEE 802.15.4's physical and medium access control layers. In the latest version of ZigBee, enhancements are prescribed for the security sublayer but we show in this paper that problems persist. In particular we show that the End-to-End Application Key Establishment Protocol is flawed and we propose a secure protocol instead. We do so by using formal verification techniques based on static program analysis and process algebras. We present a way of using formal methods in wireless network security, and propose a secure key establishment protocol for ZigBee networks.

Keywords: Verification, static analysis, wireless sensor networks, key establishment

Joint work of: Yueksel, Ender; Nielson, Hanne Riis; Nielson, Flemming

Model Checking for Network Protocols

Jaco van de Pol (University of Twente, NL)

Methods for network protocol analysis should at least be able to deal with data, behaviour, time, and stochastics. In order for such methods to be effective, one needs an expressive and clean modelling language, and a scalable toolset. I will present a number of research projects that deal with these issues.

On the modelling side, my starting point will be mCRL, a formalism that incorporates data and behaviour, in the form of process algebra. Data and behaviour come together in the elegant linear process format. Recently, we proposed an extension of mCRL with probabilities, and invented a probabilistic linear format. We expect that this will enrich probabilistic analysis tools with the capability of reasoning about data.

On the tool side I will introduce LTSMIN, a new tool for high-performance model checking. It supports both symbolic and distributed model checking, for a variety of specification languages. I will discuss plans to extend its scope to probabilistic/stochastic analysis as well.

If time permits, I will illustrate the tools with previous case studies in network protocols, for instance the sliding window protocol, gossiping protocols, authentication protocols, leader election protocols, CAN bus applications, and blackboard coordination architectures.

Keywords: Model checking, data, behaviour, process algebra, network protocols, linear format, probabilistic