Software Modelling and Verification



Staff

- **Faculty**
  Prof. Dr. Ir. Joost-Pieter Katoen
  Prof. em. Dr. Klaus Indermark
  apl. Prof. Dr. Thomas Noll

  http://moves.rwth-aachen.de

- **Secretary**
  Elke Ohlenforst
  Birgit Willms

- **Technical Staff**
  Arnd Gehrmann

- **Research Assistants**
  Souymodip Chakraborty, M.Sc.
  Dipl.-Inform. Christian Dehnert
  Hongfei Fu, M.Sc.
  Dipl.-Inform. Friedrich Gretz
  Dipl.-Inform. Jonathan Heinen
  Dipl.-Inform. Christina Jansen
  Viet Yen Nguyen, M.Sc.
  Arpit Sharma, M.Sc.
  Falak Sher, M.Sc.
  Dr. Balaguru Srivathsan (since July 2012)
  Dipl.-Inform. Sabrina von Styp
  Dipl.-Inform. Hao Wu
  Dipl.-Inform. Haidi Yue

- **Diploma/Bachelor/Master Students**
  Hussein Hamid Baagil
  David Clermont
  Bernhard Ern
  Florian Göbe
  Max Görtz
  Dennis Guck
  Jasper Heuser
  Tobias Hoffmann
  Jens Katelaan
  Tim Lange
  Bart Postma
  Tim Rohlfs
  Sebastian Staack
  Alexander Dominik Weinert
  Christoph Worreschk

- **Student Researchers**
  Henrik Barthels
  Philipp Berger
  Thomas Heinemann
  Sergey Sazonov
  Lisa von Büttner

- **Visiting Scientists**
  Pedro D'Argenio (University of Cordoba, ARG)
  Peter Höfner (NICTA Sydney, AUS)
  Ian Larson (Michigan University, USA)
  Prof. Dr. Chris J. Myers (University of Utah, USA)
  Nils Müllner (University of Oldenburg, GER)
  Johann Schuster (Universität der Bundeswehr München, GER)
  B. Srivathsan (LaBRI, Bordeaux, FRA)
  Daniel Stan (ENS Cachan, FRA)
  Gerjan Stokkink (University of Twente, NED)
  Mark Timmer (University of Twente, NED)

Overview

2012 has been an exciting year. New projects, several new group members, and several successes in terms of prizes and awards. Several changes occurred in the personnel. Maximilian Odenbrett left our group, whereas Souymodip Chakraborthy, Christian Dehnert and Hao Wu joined our group. Souy, Chris, and Stephen welcome to the chair! In July, Balaguru Srivathsan joined our group as a postdoctoral researcher. He received his PhD degree from LaBRI Bordeaux, France, for his dissertation on Timed Automata. Prior to this, he studied at IIT Bombay in India. In autumn of 2012, Sri gave a successful 10-hour lecture series on Timed Automata. Sri, welcome to our group!

Thomas Noll received this year the title of "außerplanmäßiger Professor" (shortly apl. Prof.). We congratulate Thomas with this great achievement! Dino Distefano received the prestigious Roger Needham Award for his work on software verification. This award is for a distinguished research contribution in computer science by a UK based researcher who has completed up to 10 years of post-doctoral research. Dino Distefano was one of the first PhD students supervised by Joost-Pieter Katoen and is currently professor at Queen Mary University, London. Marijn Jongerden received the German MMB (Messung, Modellierung und Bewertung von Rechensystemen) Award 2012 of the GI/ITG-Committee for Measurement, Modelling and Evaluation of Computing Systems for his PhD-thesis on "Model-based energy analysis of battery powered systems". His work was co-supervised by Boudewijn Haverkort (Twente). Finally, Joost-Pieter Katoen has been elected as Chairman of the QEST Steering Committee. QEST is an international annual conference on the Quantitative Evaluation of Systems.

We have successfully acquired two new European projects that started in autumn. SENSATION (Self Energy-Supporting Autonomous Computation), coordinated by the CISS Embedded System Institute in Aalborg (Denmark), aims at devising energy-centric modelling and optimisation tools for the design of resource-optimal reliable systems. The D-MILS (Distributed MILS for Dependable ICT) project, coordinated by The Open Group in Brussels, focuses on component-based modelling, as well as exploiting the compositional structure in the analysis.

We provided courses on Theoretical Foundations of the UML, Static Program Analysis, Advanced Model Checking, Compiler Construction, and Data Structures and Algorithms. In addition, we organised seminars on Foundations of Multi-Core Memory Models, Turing Award Winners, and Success Stories in Formal Methods. Numerous students completed their bachelor, master and diploma thesis at our chair in 2012.

Joost-Pieter Katoen

Research Projects

CARP: Correct and Efficient Accelerator Programming

J.-P. Katoen, C. Dehnert, F. Gretz, C. Jansen

EU FP7 project, coordinator: Imperial College London (A. Donaldson)

In recent years, massively parallel accelerator processors, primarily GPUs, have become widely available to end-users. Accelerators offer tremendous computing power at a low cost, and tasks such as media processing, simulation, medical imaging and eye-tracking can be accelerated to beat CPU performance by orders of magnitude. Performance is gained in energy efficiency and execution speed, allowing intensive media processing software to run in low-power consumer devices.

However, accelerators present a serious challenge for software developers. A system may contain one or more of the plethora of accelerators on the market, with many more products anticipated in the immediate future. Applications must exhibit portable correctness, operating correctly on any configuration of accelerators, and portable performance, exploiting processing power and energy efficiency offered by a wide range of devices.

The aim of CARP is to design techniques and tools for correct and efficient accelerator programming. In particular, a new portable programming language is to be developed, which can be both easily written and efficiently compiled to a lower level programming language for accelerator programming such as OpenCL and CUDA. With this new language, programmers can achieve portability of their applications, because the same source code may be compiled into efficient code for a number of accelerators, each having its own capabilities and limitations. Besides, programmers can focus on writing code without having to worry about hardware-specific effects of certain code optimisations.

Our main contribution in this EU-funded project is to provide the compiler backend with techniques to estimate quantitative aspects of accelerator programs. That is, given a program and an accelerator architecture, we aim for estimating the runtime and energy consumption of this program executing on this specific accelerator. The main motivation is that it is not at all obvious whether a program is well-suited for an accelerator, because of platform-dependent details, i.e. a program running well on one target platform might produce a poor performance on another. We aim for developing a (quantitative) formal model of an application executing on an accelerator, which is to be analysed using formal techniques. Integrating these techniques into the compiler will provide it with a means to evaluate the quality of generated low-level code with respect to multiple objectives and hence with a technique to select the most suited generated code for the target device and intended purpose.

Furthermore within the scope of the CARP project we aim at developing static verification techniques for automated analysis of software for accelerators. Due to a high degree of parallelism occurring especially in the case of highly optimised accelerator code, the error potential of this software is enormous. We focus on techniques allowing to analyse pointer manipulat-

ing behaviour of accelerator programs for both the portable programming language developed in CARP as well as low-level applications written in OpenCL.

Invariant Generation for Probabilistic Programs

F. Gretz, J.-P. Katoen,

A. McIver (Macquarie Univ, Sydney)

Verification of sequential programs rests typically on the pioneering work of Floyd, Hoare and Dijkstra in which annotations are associated with control points in the program. For probabilistic programs, quantitative annotations are needed to reason about probabilistic program correctness. We generalise the method of Floyd, Hoare and Dijkstra to probabilistic programs by making the annotations real- rather than Boolean-valued expressions in the program variables. The crucial annotations are those used for loops, the loop invariants. Thus in particular we focus on real-valued, quantitative invariants: they are random variables whose expected value is not decreased by iterations of the loop.

One way of finding annotations is to place them speculatively on the program, as parameterised formulae containing only first-order unknowns, and then to use a constraint solver to search for parameter instantiations that would make the associated "verification conditions" true. In this project, we aim to generalise and extend constraint-solving techniques for invariant generation to probabilistic programs. This allows for the verification of probabilistic programs that cannot be treated with currently available automated techniques such as abstraction refinement together with model checking. This work includes theory development as well as prototypical tool development to illustrate the feasibility.

Verifying Pointer Programs with Unbounded Heap Structures

J. Heinen, C. Jansen, J.-P. Katoen, Th. Noll

The incorrect use of pointers is one of the most common sources of software errors. Proving the correctness of pointer-manipulating programs with unbounded heap, let alone algorithmically, is a highly non-trivial task. This project attempts to develop automated verification techniques and accompanying tool support for programs with memory allocation that handle linked data structures which are potentially unbounded in their size.

We developed a novel abstraction framework that is based on graph grammars, more precisely context-free hyperedge replacement grammars, as an intuitive formalism for abstractly modelling dynamic data structures. The key idea is to use the replacement operations which are induced by the grammar rules in two directions. By a backward application of some rule, a subgraph of the heap can be condensed into a single nonterminal edge, thus obtaining an abstraction of the heap. By applying rules in forward direction, certain parts of the heap which have

been abstracted before can be concretised again. This avoids the necessity for explicitly defining the effect of pointer-manipulating operations on abstracted parts of the heap.

While in past years our focus was set on correctness and efficiency of the state space creation, we concentrated on the actual analysis this year. We developed an algorithm to check languages defined by graph grammars against MSO-formulae, allowing us to express complex properties of heap structures, such as reachability and shape properties. To express temporal properties we extended CTL* to quantified CTL*. While using CTL* allows to express temporal properties of the entire heap, quantified CTL* allows us to track objects over time and describe how their relation to other objects on the heap changes during the execution. We developed an on-the-fly algorithm for Quantified CTL*.

Another research topic was the investigation of the relations between our approach and Separation Logic. In Separation Logic heap structures are represented as formulae, which are strongly related to our abstract graph representation. We formalised the relation which will give us the possibility to carry over results from one approach to the other in the future.

Formal Models of Microcontroller Systems

Th. Noll,

J. Brauer (Chair of Software for Embedded Systems)

Embedded systems usually operate in uncertain environments, giving rise to a high degree of nondeterminism in the corresponding formal models. Moreover they generally handle data spaces whose sizes grow with the memory and the word length of the respective microcontroller architectures. This, together with other effects, leads to the well-known state-space explosion problem, meaning that the models of those systems grow exponentially in size as the parameter values increase. Careful handling of both nondeterminism and large data spaces is therefore crucial for obtaining efficient methods and tools for analysis and verification.

The goal of this project, carried out in close cooperation with the Embedded Software Laboratory of our department, is to develop abstraction techniques to tackle this problem. With regard to control structures, a technique for refining loops in microcontroller programs has been developed. It is based on abstract interpretation using octagons and affine equalities in order to identify infeasible sequences of loop iterations. Our approach naturally integrates wrap-around arithmetic during the generation of abstractions. Abstract interpreters operating on a refined control structure then typically derive strengthened program invariants without having to rely on complicated domain constructions.

With regard to data spaces, activities have been concentrating on static analysis methods for approximating the possible run-time values of data values. For this purpose, intervals have successfully been used for decades. Binary code on microcontroller platforms, however, is different from high-level code in that data is frequently altered using bit-wise operations and that the results of operations often depend on the hardware configuration. We therefore came up with a method that combines word- and bit-level interval analysis and that integrates a

hardware model by means of abstract interpretation in order to handle these peculiarities. Both techniques have successfully been applied to a suite of benchmark examples.

Model-Based Energy Optimization of Automotive Control Systems

J.-P. Katoen, Th. Noll, H. Wu,

joint project together with Th. Santen, D. Seifert (Microsoft Research, Advanced Technology Labs Europe)

funded by Microsoft Research and RWTH Aachen University Seed Fund

Reducing the energy consumption of controllers in vehicles requires sophisticated regulation mechanisms. Better power management can be enabled by allowing the controller to shut down sensors, actuators or embedded control units in a way that keeps the car safe and comfortable for the user, with the goal of optimising the (average or maximal) energy consumption. In this project, we develop an approach to systematically explore the design space of software-to-hardware mappings to determine energy-optimal deployments. It employs constraint-solving techniques for generating deployment candidates and probabilistic analyses for computing the expected energy consumption of the respective deployments. The feasibility and scalability of the method have been demonstrated by several case studies. Current efforts are concentrating on transferring the results to other application domains, such as the analysis of service-level requirements in cloud computing environments.

COMPASS: Correctness, Modelling and Performance of Aerospace Systems

C. Dehnert, J.-P. Katoen, V.Y. Nguyen, Th. Noll,

joint project together with the groups of Alessandro Cimatti (Fondazione Bruno Kessler, Centre for Scientific and Technological Research, Trento, Italy), and Xavier Olive (Thales Alenia Space, On Board Software Department, Cannes, France)

funded by European Space Agency (ESA), Thales Alenia Space and EU FP7 Programme

The size and complexity of software in spacecraft is increasing exponentially, and with that comes a strong need to validate it more effectively within the context of the overall (spacecraft) system. Current validation methods are labour-intensive as they rely on manual analysis, review and inspection. In the past four years we developed, with requirements from the European space industry, a novel modeling language and toolset to (semi-)automate validation activities. The AADL modeling language enables engineers to express the system, software and their reliability aspects, and the COMPASS toolset analyses them on correctness, safety, dependability and performance requirements using state-of-the-art model checking techniques and probabilistic versions thereof.

The results and experiences gained from this project have triggered follow-up activities to drive the further development of the overall technology. Several case studies have been performed with industrial involvement, of which two were conducted on a spacecraft in development at system-level, namely a satellite platform. These efforts were carried out in parallel with the conventional software development for this system. The whole effort revealed numerous inconsistencies in the early design documents, and the use of formal analyses provided additional insight on discrete system behaviour (comprising nearly 50 million states), on hybrid system behaviour involving discrete and continuous variables, and enabled the automated generation of large fault trees (66 nodes) for safety analysis that typically are constructed by hand. The model's size pushed the computational tractability of the algorithms underlying the formal analyses, and revealed bottlenecks for future theoretical research. Additionally, the effort led to newly learned practices from which subsequent formal modelling and analysis efforts shall benefit, especially when they are injected in the conventional software development lifecycle. The case demonstrates the feasibility of fully capturing a system-level design as a single comprehensive formal model and analyse it automatically using a toolset based on (probabilistic) model checkers.

A more research-oriented follow-up activity is a PhD project on "Extending and Improving Formal Methods for System/Software Co-Engineering" within ESA's Networking Partnering Initiative that started in 2010 and aims to extend the currently used formal methods to cover a broader class of models in terms of expressiveness and scalability. This particularly applies to systems that exhibit complex characteristics in terms of non-determinism, timed and hybrid behaviour, and discrete and continuous-time probabilistic errors. Our goal is to specify formal semantics for such models, and to develop novel and advanced techniques to support their analysis. In particular, compositional techniques exploiting the hierarchical structure of component-based models will be required to combat the state-space explosion problem.

Another closely related project named Distributed MILS (Multiple Independent Layers of Security), funded by the European Union in the framework of the FP7 programme, has recently been started. It aims to to investigate and develop models and certification techniques for security architectures in heterogeneous, networked, service and computing environments. The AADL modelling language is also a cornerstone in this project, and it is aimed to leverage (parts of) the COMPASS toolset to provide analyses that support the assessment and certification of security aspects.

More information on COMPASS is available at http://compass.informatik.rwth-aachen.de.

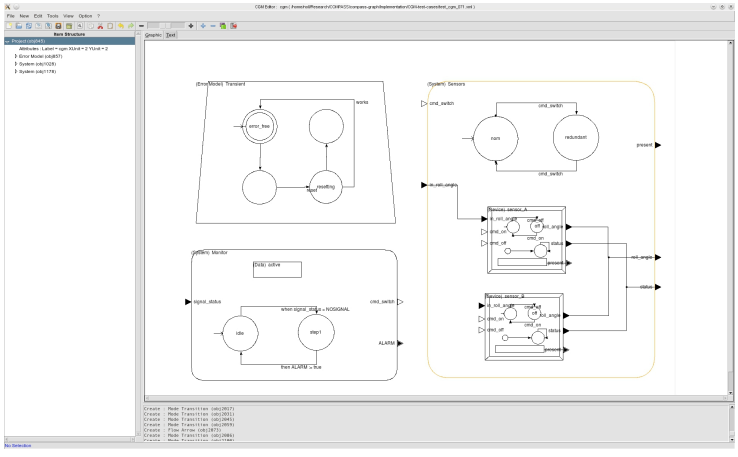

Efficient Reachability Analysis for Probabilistic Timed Automata

J.-P. Katoen, B. Srivathsan

Soon after the success of conventional model-checking in the 80's where qualitative temporal properties of systems could be automatically verified, people started considering real-time properties. This gave birth to the timed automaton (TA) model in the early 90's. Timed automata are finite automata extended with clocks whose values can be compared with constants and can be set to zero during transitions. They model systems where events happen according to certain timing constraints, for instance a car-brake controller system which has to accentuate the hardware within a fixed number of seconds.

A timed automaton expresses a system model purely in terms of non-determinism. However, in some situations, it is desirable to express the probability of a system exhibiting a certain behaviour. To model such behaviour in systems, the model of probabilistic timed automata (PTA) has been introduced and has been implemented in the tool PRISM. Advances in timed automata technology have a direct influence on probabilistic timed automata.

A fundamental property in verification is state reachability. This problem is known to be decidable both for these automaton models and has been an active subject of research for over two decades now. Currently, there are mature tools implementing efficient reachability algorithms and have been used for industrial case studies.

The standard solution to the reachability problems for both TA and PTA involve the so-called abstractions of zones. A very recent work on TA has proposed a new reachability algorithm that involves non-convex abstractions and an on-the-fly learning of parameters for abstraction. Prototype implementation has shown a three-fold gain in some standard benchmarks.

The goal of this project is to understand to what extent the new results on abstractions for timed automata apply to probabilistic timed automata. The work would involve theory development as well as implementation to validate the gains.

Infinite-State Probabilistic Systems

H. Fu, J.-P. Katoen,

funded by China Scholarship Council (CSC)

Verification of infinite structures has been extensively studied in the past two decades. The motivation of this study is that (i) typical system components are often infinite-state (e.g., counters, buffers), which cannot be modelled by a finite-state system; and (ii) adding timed information to a finite-state system will also induce an infinite-state system.

The difference between finite-state and infinite-state verification lies in the fact that exhaustive traversal of the state space which is effective on finite-state systems cannot be applied to

infinite-state systems. Thus new techniques should be developed. Currently, the study of infinite-state verifications is divided into two sub-areas: equivalence checking and model checking. In equivalence checking, the task is to decide if two given system are equivalent under a pre-established equivalence relation. In model checking, the task is to verify that a given system satisfies a certain property encoded by a logical formula.

In the non-probabilistic setting, namely on labeled transition systems, verification of infinite structures has been well studied. Various results have been obtained on infinite-state models such as Pushdown Automata, Petri Nets, etc. The aim of this project is to investigate infinite-state verification in a probabilistic setting. Probability is a mechanism to model uncertainty, which can be caused by randomised algorithms, unpredictable errors, or simply underspecification in system design.

Our main work is to study probabilistic model checking and probabilistic equivalence checking on probabilistic infinite-state systems. To do so, we may extend existing techniques on discrete infinite-state system to the probabilistic setting, or instead discover new techniques if necessary.

UMIC: Ultra High Speed Mobile Information and Communication

J.-P. Katoen, H. Yue

The evaluation of the quality-of-service of Wireless Sensor Networks is mostly done by simulation. In the context of the UMIC project, and in cooperation with the company CHESS, Haarlem, NL, we evaluated the second generation of a gossiping MAC protocol (GMAC), a TDMA protocol for completely unconfigured wireless networks, which aims to maintain message propagation with minimal energy demand.

The new GMAC protocol is modelled with a specific radio model, the SINR model of Gupta Kumar. And as network topology, we considered three different node arrangements: grid, uniformed distribution and Gaussian distribution.

Simulations, carried out with the MoDeST/Möbius tool set, show that comparing to the simple slotted Aloha protocol, the new GMAC has indeed a significant beneficial influence on energy consumption for all the three types of networks. However, it also delayed message propagation. Besides that, we discovered that GMAC with constant sending power may cause a lot of interference in dense area, and in sparse area, nodes may disconnect to each other. Hence we proposed a dynamic energy scheduling schema, so that a node can regulate its sending power with respect to its current number of neighbours. Simulation results show that for Gaussian networks, our dynamic power management not only improved energy consumption, but also accelerated message propagation. Furthermore, we observed the existence of the Pareto Principle in the system, i.e., 20% of the nodes consumes 80% of the energy of the system.

Our current research aims at gathering evidence that the SINR model is in fact realistic enough to allow the derivation of reliable measures for wireless sensor networks using simu-

lation. For that we try to explain measurement data from CHESS by means of the SINR model.

In cooperation with the UMIC Mobile Network Performance Group, we modelled and examined a newly developed Wireless Token Ring Protocol (WTRP), which aims to establish robust wireless communication for real-time applications, such as the braking system in vehicles. Due to the high requirement on correctness from the application layer, the protocol should be ultra reliable and robust to all possible failures. Traditional evaluation techniques like simulation provide essential characteristics of the WTRP, and model checking gives evidence on the correctness of the protocol.

We model the WTRP as a PTA (probabilistic timed automaton) in PRISM, a tool for formal modelling and analysis of systems that exhibit random, probabilistic, or timed behaviour. The WTRP has several components, and our starting point is the core part, the normal ring module, and the ring consists of five stations. With PRISM, we not only verified the "eventually" path property of the protocol under reliable channel assumption, but also calculated the probability of token missing when the channel is unreliable.

Since the latter situation is more common and almost inevitable in real applications, our future research will focus on the relation between channel failure rate and the number of necessary retransmissions of messages, to obtain a wireless communication as reliable as possible.

SYRUP: SYmbolic RedUction of Probabilistic Models

J.-P. Katoen

M. Timmer, M. Stoelinga, J. van de Pol (all three from University of Twente, NL)

funded by the Dutch Research Council (NWO)

Efficient model-checking algorithms exist for qualitative and quantitative properties for a range of probabilistic models. Their popularity is due to the presence of powerful software tools, and their wide applicability; security, distributed algorithms, systems biology, dependability and performance analysis, to mention a few. The main deficiencies of probabilistic model checking are the state explosion problem and the restricted treatment of data.

The state space grows exponentially in the size of system components and data domains. Whereas most abstraction techniques obtain smaller models by collapsing sets of concrete states at the model level, this project takes a radically different approach. We will develop and implement symbolic reduction techniques for probabilistic models. These techniques aim to reduce models by model transformations at the language level in order to minimise state spaces prior to their generation while preserving functional and quantitative properties. Our symbolic reductions will support data as first-class citizens, i.e., we will develop techniques to symbolically reduce formalisms for modelling probabilistic systems that are equipped with rich data types, allowing, e.g., probabilistic choices parameterised with data.

Our approach is based on successful symbolic transformation techniques in the traditional and timed setting, viz. linear process equations (LPEs). We will generalise and extend these tech-

niques to probabilistic automata (PA), a model akin to Markov Decision Processes that is tailored to compositional modelling. The LPE technique is applicable to large or even infinite systems, and will be equipped with symbolic transformations such as confluence reduction, bisimulation minimisation and static analysis for PA.

MoVeS: Modeling, Verification and Control of Complex Systems

S. Chakraborty, J.-P. Katoen, F. Sher

EU FP7 project, coordinator: ETH Zurich (J. Lygeros)

In the context of the EU FP7 project "Modelling, verification and control of complex systems: From foundations to power network applications" (partners: ETH Zurich, TU Delft, University of Oldenburg, Politecnico Milano, and Honeywell), we propose novel methods for modelling, analysis and control of complex, large scale systems. Fundamental research is motivated by application problems in power networks. We adopt the framework of stochastic hybrid systems (SHS), which allows one to capture the interaction between continuous dynamics, discrete dynamics and probabilistic uncertainty. In the context of power networks, SHS arise naturally: continuous dynamics model the evolution of voltages, frequencies, etc. Discrete dynamics reflect changes in network topology, and probability represents the uncertainty about power demand and (with the advent of renewables) power supply. More generally, because of their versatility, SHS are recognised as an ideal framework for capturing the intricacies of complex, large scale systems.

Motivated by this, considerable research effort has been devoted to the development of modelling, analysis and control methods for SHS, in both computer science (giving rise to theorem proving and model checking methods) and in control engineering (giving rise to optimal control and randomised methods). Despite several success stories, however, none of the methods currently available is powerful enough to deal with real life large scale applications. We feel that a key reason for this is that the methods have been developed by different communities in relative isolation, motivated by different applications. As a consequence, synergies between them have never been fully explored.

In this project, we systematically explore such synergies. Our multi-disciplinary team, which brings together experts on all the state of the art SHS methods, will establish links between model checking, theorem proving, optimal control and randomised methods. Leveraging on their complementary strengths we will develop combined strategies and tools to enable novel applications to complex, large scale systems. Common power networks case studies will provide a testing ground for the fundamental developments, motivate them, and keep them focused.

Minimisation of Markov Models

J.-P. Katoen, A. Sharma

funded by the India4EU Programme

Markov chains are widely used for the evaluation of performance and dependability of information processing systems. Extending Markov chains with rewards results in Markov reward models which are useful for analysing the average behaviour of executions in Markov chains. Equivalence relations are used to reduce the state space of Markov chains, by combining equivalent states into a single state. The reduced state space obtained under an equivalence relation called a quotient can then be used for analysis provided it preserves a rich class of properties of interest. Various branching-time relations on Markov chains have been defined such as weak and strong variants of bisimulation equivalence and simulation pre-orders. Their compatibility with (fragments of) stochastic variants of CTL has been thoroughly investigated. Stochastic model checking tools such as PRISM and MRMC have been used to model check interesting properties on Markov chains and Markov reward models, respectively.

The goal of this project is to explore and investigate the linear-time equivalence relations and interesting properties that are preserved under these equivalences for Markov chains. During the course of this project we also plan to study and explore if these linear-time equivalences are compatible with compositional modelling of systems. Next step would involve developing quotienting algorithms and implementing tools for computing these equivalences. Finally, we plan to extend the minimisation techniques developed for Markov chains to other more expressive models, for example Markov automata, Interactive Markov chains and Markov decision processes and also to non-probabilistic systems.

Synthesising of Model Based Testing for Process Control Engineering

S. von Styp, J.-P. Katoen,

L. Yu, U. Epple (Chair of Process Control Engineering)

In Process Control Engineering, controller for plants that are not correct to their specification can cause fatal disasters, i. e. when tanks with acid run over people get injured or to high pressure leads to explosions. Therefore an intensive testing of the controller is crucial but it consumes a lot of time and money.

Model based testing is one promising technique allowing the automatic generation of test-cases from a given formal model. In Model based testing the specification is given by a transition system. A conformance relation formally defines under which circumstances an implementation is correct with respect to the specification. Based on this relation, test-cases are derived automatically and are used to for testing the real implementation. In this project, which is a cooperation with the institute for process control engineering, we apply the methods of

model based testing on the plant controller to automate and systematise the testing process in order to reduce costs and time.

In oder to apply model based testing, first the specifications of controllers, given as sequential function charts, are translated to transition systems. Hence first rules for the systematic translation had to be developed. We started testing using simple controller, i.e. a motor controller, which allow us to use the already existing theories such as ioco and sioco including the testing tool JTorX. This test-cases are still in an early stage and therefore only allow testing for a restricted set of programs. Future research shall loosen this restriction.

Data-dependent control flow in combination with real-time behaviour is an important feature of the plant controllers, i.e. the next action depends on the current filling level of a tank and it may crucial that certain actions are executed within a certain time. Therefore this project looks at extending the existing test theory to allow real-time behaviour together with data-dependent control flow. We start by giving a formal definition in form of a transition system for representing systems that allow data-dependent control flow for inputs and outputs and real-time behaviour. Afterwards a symbolic trace semantics is defined. This semantics then is employed to define the conformance relation, which describes under which conditions an implementation is correct with respect to a given specification. Future steps will include to look at the application such as on the fly testing. This then shall be implemented in the test tool JTorX and finally be used to test controllers in process control engineering.

Model Based Testing for AADL Models

S. von Styp, S. Sazonov, J.-P. Katoen

The COMPASS toolset allows model checking the specification given as an AADL model. To verify whether the actual implementation conforms to this specification model based testing is needed. So far model based testing using an AADL model for specification is not possible. Therefore a translation from AADL to transition systems with data and data-dependent control flow is developed. This transition system then is used by the test tool JTorX which finally will be integrated into the COMPASS toolset. Furthermore a theory is developed, allowing testing single components and its composition providing an method of simply adding components without regenerating the hole transition system again.

ROCKS: Rigorous Dependability Analysis using Model Checking Techniques for Stochastic Systems

S. Chakraborty, H. Fu, F. Gretz, J.-P. Katoen, A. Sharma, F. Sher

funded by the Dutch Research Council (NWO) and DFG

Today's society relies increasingly on the correct and timely functioning of a large variety of information and communication technology systems. Can this reliance be justified? Dependability analysis answers this question. Rigorous and systematic dependability analysis must then play an important role in the design of such systems. Since many dependability properties are stochastic in nature, stochastic analysis techniques are crucial in developing reliable and safe computer systems.

The ROCKS project focuses on two system classes which are gaining prominence in the world of computing but which are not amenable to classic stochastic analysis techniques. Large scale homogeneous systems, such as wireless sensor networks and gossiping protocols, provide a challenge because of the sheer size of the systems involved. Safety-critical heterogeneous systems, such as production plants and automotive control systems, on the other hand consist of a number of very different components. The challenge here is to handle the diversity of system modalities.

Within ROCKS we further study how, given a system configuration or parameter set, the optimal design can be synthesised automatically. Attention will also be given to the study of architectural description languages which are increasingly being used to describe complex systems, but for which analysis techniques are often lacking. The members of the ROCKS project cooperate in four different research areas: Modelling, analysis, synthesis and case studies. In modelling we study how complex systems can be represented concisely, accurately and hierarchically. Analysis techniques to study the properties of such models are developed as well as synthesis techniques in order to automatically generate optimal models. Finally the applicability of the newly developed models and techniques is studied in a number of industrial case studies.

Other Activities

Joost-Pieter Katoen

- Deputy Chair of the Steering Committee of ETAPS (European Joint Conferences on Theory and Practice of Software).
- Member of the Steering Committee of FORMATS (Formal Methods and Analysis of Timed Systems)
- Chair of the Steering Committee of QEST (Quantitative Evaluation of Systems), since September 2012.
- Member of the Steering Committee of TACAS (Tools and Algorithms for the Construction and Analysis of Systems).
- Member of the Editorial Board of the Journal on Software Tools for Technology Transfer (STTT), Springer Verlag.
- Member editorial board The Scientific World Journal, Hindawi Publishers.
- Board Member of the Dutch Society on Theoretical Computer Science (NVTI).
- Senior member of the Association of Computing Machinery (ACM)
- **Member of the Program Committees of the following events:**
  - 15th International Conference on Foundations of Software Science and Computation Structures (FoSSaCS 2012)
  - 17th International Workshop on Formal Methods for Industrial Critical Systems (FMICS 2012)
  - 17th International Conference on Implementation and Application of Automata (CIAA 2012)
  - 9th International Conference on Quantitative Evaluation of SysTems (QEST 2012)
  - 10th International Symposium on Automated Technology for Verification and Analysis (ATVA 2012)
  - 10th Summer School on Modelling and Verifying Parallel Processes (MOVEP 2012)
- **Invited speaker at:**
  - Workshop on the 15th Anniversary of LSV, ENS Cachan, France
  - Tutorial at Joint Conference iFM & ABZ 2012, Pisa, Italy
  - 33rd International Joint Conference Petri Nets & ACSD 2012, Hamburg, Germany
  - NATO Summer School on Engineering Dependable Software Systems, Marktoberdorf 2012
  - Tutorial at Automated Software Engineering (ASE), Essen, Germany
  - ROCKS Autumn School, Vahrn, Italy

- Member of the IFIP Working Group 1.8 on Concurrency Theory.
- Member of the IFIP Working Group 2.2 on Programming Concepts.
- Member of several external international PhD committees.
- Member Selection Committee of Professorship at Free University of Brussels, Belgium.
- Member Selection Commitee of Associate Professor DTU Lyngby, Denmark.
- Chairman of Selection Committee of Full Professorship on Logics and Automata Theory.
- Chairman of the Examination Board of Department of Computer Science (until April 2012).
- Head of Computer Science Department (since April 2012).

Thomas Noll

- Student advisor for the following applied subjects within CS:
  Electrical Engineering, Civil Engineering, and Medicine
- Member of the examination board for Computer Science
- Member of CS Commission for Teaching Service
- **Member of the Program Committees of the following events:**
  - Software Engineering Track at the 28th Annual ACM Symposium on Applied Computing (SAC 2013)
  - 3rd Analytic Virtual Integration of Cyber-Physical Systems Workshop (AVICPS 2012) at the 33rd IEEE Real-Time Systems Symposium (RTSS 2012)
  - 1st International Workshop on Formal Techniques for Safety-Critical Systems (FTSCS 2012) at the 14th International Conference on Formal Engineering Methods (ICFEM 2012)
  - 6th International Conference on Software Security and Reliability (SERE 2012)

Talks and Publications

Talks

Marco Bozzano, and Joost-Pieter Katoen. Safety, Dependability and Performance Analysis of Extended AADL Models. Tutorial at 9th Int. Conf. on Integrated Formal Methods (iFM), 2012.

Marco Bozzano, Joost-Pieter Katoen, and Thomas Noll. Safety, Dependability and Performance Analysis of Extended AADL Models. Tutorial at 27th IEEE/ACM Int. Conf. on Automated Software Engineering (ASE), 2012.

Pierre Dissaux, Viet Yen Nguyen, and Thomas Noll. Tutorial on COMPASS Toolset. AADL Standards Meeting at SAE AeroTech Congress & Exhibition, Toulouse, France, 2011.

Hongfei Fu. Computing Game Metrics on Markov Decision Processes. Talk, ICALP, Birmingham, 2012.

Hongfei Fu. Computing Game Metrics on Markov Decision Processes. Talk, Oxford University, Oxford, 2012.

Friedrich Gretz. Prinsys - A Software Tool for the Synthesis of Probabilistic Invariants. Talk, WING workshop, Manchester, 2012.

Friedrich Gretz. PRINSYS – Probabilistic INvariant SYnthesiS . Talk, GRK Workshop, Dagstuhl, 2012. Gemeinsamer Workshop der DFG Informatik-Graduiertenkollegs.

Friedrich Gretz. Towards Analysis of Sequential Probabilistic Programs. Talk, Oxford, 2012.

Friedrich Gretz. Average Runtime Estimation. Talk, CARP, Cambridge, 2012. CARP Project Meeting at ARM .

Friedrich Gretz. Operational versus Weakest Precondition Semantics for the Probabilistic Guarded Command Language. Talk, QEST, London, 2012.

Jonathan Heinen. Analysis and Verification of Heap-Manipulating Programs. Talk at D-CON, Kaiserslautern, 2012.

Jonathan Heinen. Juggrnaut: The Analysis of Object Oriented Programs. Talk at mac, Uni Bamberg, 2012.

Christina Jansen. Verifying Pointer Programs Using Hyperedge Replacement Grammars. Talk, ULB Bruxelles, 2012. Invited Presentation.

Joost-Pieter Katoen. Towards Trustworthy Aerospace Design Using Formal Methods: An Experience Report. Talk at PUMA Graduiertenkolleg, TU Munich,, 2012.

Joost-Pieter Katoen. Analyzing Probabilistic Programs: Pushing the Limits of Automation. Invited Talk at Workshop on 15 Years of LSV, ENS Cachan, 2012.

Joost-Pieter Katoen. One Can Do Much More With Model Checking Than You Think! . Invited Talk at Opening of the FUNDP Research Centre Fundamentals of Computer Science (FOCUS). Namur, Belgium, 2012.

Joost-Pieter Katoen. Towards Trustworthy Aerospace Design Using Formal Methods: An Experience Report. Talk at Embedded Systems Institute, Eindhoven, NL, 2012.

Joost-Pieter Katoen. Revisiting GSPNs: New Semantics and Analysis Algorithms. nvited Talk at 33rd International Conference on Application and Theory of Petri Nets and Concurrency (ICATPN) and 12th International Conference on Application of Concurrency to System Design (ACSD), Hamburg, Germany, 2012.

Joost-Pieter Katoen. Performance Analysis by Model Checking. Five Lectures at NATO Summer School on Engineering Dependable Software Systems. Marktoberdorf, Germany, 2012.

Joost-Pieter Katoen. Revisiting GSPNs: New Semantics and Analysis Algorithms. Talk at IFIP WG 2.2 Meeting, CWI Amsterdam, NL, 2012.

Joost-Pieter Katoen. Revisiting GSPNs: New Semantics and Analysis Algorithms. Kolloquium at IMT Lucca, Italy, 2012.

Joost-Pieter Katoen. Towards Trustworthy Aerospace Design Using Formal Methods: An Experience Report. Invited lecture at ROCKS Autumn School, Vahrn, Italy, 2012.

Viet Yen Nguyen. COMPASS Graphical Modeller. Talk at ESTEC Final Presentation Days, 2012.

Viet Yen Nguyen. Formal Correctness, Safety, Dependability and Performance of a Satellite. Talk at ICSE 2012, 2012.

Viet Yen Nguyen. Satellite Platform Case Study With SLIM and COMPASS. Talk, Dagstuhl Seminar on Architecture-Driven Semantic Analysis of Embedded Systems, 2012.

Thomas Noll. Correctness, Safety and Fault Tolerance in Aerospace Systems: The ESA COMPASS Project. Talk, RWTH Aachen University, 2012. CS Colloquium.

Thomas Noll. Correctness, Safety and Fault Tolerance in Aerospace Systems: The ESA COMPASS Project. Talk, Dagstuhl Seminar on Architecture-Driven Semantic Analysis of Embedded Systems, 2012.

Arpit Sharma. Weighted Probabilistic Equivalence Preserves ω-Regular Properties. MMB/DFT Conference, Kaiserslautern, Germany, 2012.

Falak Sher. Compositional Abstraction Techniques for Probabilistic Automata. Talk, TCS, Amsterdam, 2012.

Publications

Christel Baier, E. Moritz Hahn, Boudewijn R. Haverkort, Holger Hermanns, and Joost-Pieter Katoen. Model Checking for Performability. Mathematical Structures in Computer Science. pages 1–42. 2013.

Kamal Barakat, and Thomas Noll. A Native Approach to Modeling Timed Behavior in the Pi-Calculus (short paper). In Proc. of 6th IEEE Int. Symp. on Theoretical Aspects of Software Engineering (TASE). pages 253–256. IEEE, 2012.

Marco Bozzano, Alessandro Cimatti, Joost-Pieter Katoen, Viet Yen Nguyen, Thomas Noll, and Marco Roveri. Safety, Dependability, and Performance Analysis of Extended AADL Models. The Computer Journal, 54(5):754–775, 2011.

Alessandro D'Innocenzo, Alessandro Abate, and Joost-Pieter Katoen. Robust PCTL Model Checking. In Hybrid Systems: Computation and Control (HSCC). pages 275–286. ACM Press, 2012.

Marie-Aude Esteve, Joost-Pieter Katoen, Viet Yen Nguyen, Bart Postma, and Yuri Yushtein. Formal Correctness, Safety, Dependability and Performance Analysis of a Satellite. In 34th International Conference on Software Engineering (ICSE). pages 1022–1031. ACM and IEEE CS Press, 2012.

Hongfei Fu. Computing Game Metrics on Markov Decision Processes. In 39th International Colloquium on Automata, Languages and Programming (ICALP). pages 227–238. Volume 7392 of LNCS. Springer-Verlag, 2012.

Hongfei Fu, and Joost-Pieter Katoen. Deciding Probabilistic Simulation between Probabilistic Pushdown Automata and Finite-State Systems. In Foundations of Software Technology and Theoretical Computer Science (FSTTCS). LIPIcs. Schloss Dagstuhl, 2011.

Friedrich Gretz, Joost-Pieter Katoen, and Annabelle McIver. Operational versus Weakest Pre-condition Semantics for the Probabilistic Guarded Command Language. In Quantitative Evaluation of Systems (QEST). IEEE CS Press, 2012.

Dennis Guck, Tingting Han, Joost-Pieter Katoen, and Martin R. Neuhäußer. Quantitative Timed Analysis of Interactive Markov Chains. In NASA Formal Methods Symposium (NFM). pages 8–23. Volume 7226 of LNCS. Springer-Verlag, 2012.

Jonathan Heinen, Christina Jansen, and Henrik Barthels. Juggrnaut - An Abstract JVM. In 2nd Int. Conf. on Formal Verification of Object-Oriented Software (FoVeOOS). Volume 7421 of LNCS. Springer-Verlag, 2012.

Frédéric Herbreteau, B Srivathsan, and Igor Walukiewicz. Efficient emptiness check for timed Büchi automata. Formal Methods in System Design, 40(2):122–146, 2012.

Frédéric Herbreteau, B Srivathsan, and Igor Walukiewicz. Better abstractions for timed automata. In LICS. pages 375 –384. IEEE, 2012.

Nils Jansen, Erika Abraham, Matthias Volk, Ralf Wimmer, Joost-Pieter Katoen, and Bernd Becker. The COMICS Tool - Computing Minimal Counterexamples for DTMCs. In Proc. of the 10th Int. Symp. on Automated Technology for Verification and Analysis (ATVA). pages 349–353. Volume 7561 of LNCS. Springer-Verlag, 2012.

Joost-Pieter Katoen, Jaco van de Pol, Marielle Stoelinga, and Mark Timmer. A linear process-algebraic format with data for probabilistic automata. Theoretical Computer Science, 413:36–57, 2012.

Joost-Pieter Katoen, Daniel Klink, Martin Leucker, and Verena Wolf. Three-Valued Abstraction for Probabilistic Systems. Journal on Logic and Algebraic Programming, 81(4):356–389, 2012.

Joost-Pieter Katoen. Model Checking: One Can Do Much More Than You Think!. In Fundamentals of Software Engineering (FSEN). pages 1–14. Volume 7141 of LNCS. Springer-Verlag, 2012.

Joost-Pieter Katoen. GSPNs Revisited: Simple Semantics and New Analysis Algorithms. In Applications of Concurrency to System Design (ACSD). pages 6–12. IEEE CS Press, 2012.

Etienne Lozes, Florent Jacquemard, Jules Villard, and Ralf Treinen. Multiple Congruence Relations, First-Order Theories on Terms, and the Frames of the Applied Pi-Calculus. In Theory of Security and Applications (TOSCA 2011). pages 166–185. Volume 6993 of LNCS. Springer-Verlag, 2012.

Thomas Noll. Correctness, Safety and Fault Tolerance in Aerospace Systems: The ESA COMPASS Project (Abstract). In Architecture-Driven Semantic Analysis of Embedded Systems (Dagstuhl Seminar 12272). pages 42–42. Volume 2 of Dagstuhl Reports. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2012.

Arpit Sharma, and Joost-Pieter Katoen. Weighted Lumpability on Markov Chains. In 8th Ershov Informatics Conference (PSI) . pages 322–339. Volume 7162 of LNCS. Springer-Verlag, 2012.

Arpit Sharma. Weighted Probabilistic Equivalence Preserves ω-Regular Properties. In 16th Measurement, Modeling, and Evaluation of Computing Systems and Dependability and Fault Tolerance Conference (MMB/DFT). pages 121–135. Volume 7201 of LNCS. Springer Verlag, 2012.

Falak Sher, and Joost-Pieter Katoen. Compositional Abstraction Techniques for Probabilistic Automata. In IFIP Conference on Theoretical Computer Science (TCS). pages 325–341. Volume 7604 of LNCS. Springer-Verlag, 2012.

B Srivathsan, and Igor Walukiewicz. An alternate proof of Statman's finite completeness theorem. Information Processing Letters, 112(14-15):612–616, 2012.

Mani Swaminathan, Joost-Pieter Katoen, and Ernst-Rüdiger Olderog. Layered Reasoning for Randomized Distributed Algorithms. Formal Aspects of Computing, 24(4–6):477–496, 2012.

Bart Theelen, Joost-Pieter Katoen, and Hao Wu. Model Checking of Scenario-Aware Dataflow with CADP. In Design, Automation, and Test in Europe (DATE). pages 653–658. IEEE CS Press, 2012.

Mark Timmer, Joost-Pieter Katoen, Jaco van de Pol, and Marielle Stoelinga. Efficient Modelling and Generation of Markov Automata. In Concurrency Theory (CONCUR). pages 364–379. Volume 7454 of LNCS. Springer-Verlag, 2012.

Ralf Wimmer, Nils Jansen, Erika Abraham, Joost-Pieter Katoen, and Bernd Becker. Minimal Critical Subsystems as Counterexamples for ω-Regular DTMC Properties. In Methoden und Beschreibungssprachen zur Modellierung und Verifikation von Schaltungen und Systemen (MBMV). pages 169–180. Verlag Dr. Kovac, 2012.

Ralf Wimmer, Nils Jansen, Erika Abraham, Bernd Becker, and Joost-Pieter Katoen. Minimal Critical Subsystems for Discrete-Time Markov Models. In Tools and Algorithms for the Construction and Analysis of Systems (TACAS). pages 299–314. Volume 7214 of LNCS . Springer-Verlag, 2012.

Technical Reports

Ralf Wimmer, Nils Jansen, Erika Abraham, Joost-Pieter Katoen, and Bernd Becker. Minimal Counterexamples for Refuting ω-Regular Properties of Markov Decision Processes. Technical Report 88, Reports of SFB/TR 14 AVACS, 2012. ISSN: 1860-9821.