Software Modelling and Verification

Staff

- **Faculty**
  Prof. Dr. Ir. Joost-Pieter Katoen
  Prof. em. Dr. Klaus Indermark
  AOR Priv.-Doz. Dr. Thomas Noll

  http://moves.rwth-aachen.de

- **Secretary**
  Elke Ohlenforst

- **Technical Staff**
  Arnd Gehrmann

- **Research Assistants**
  Dr. Henrik Bohnenkamp (until May 2011)
  Hongfei Fu, M.Sc.
  Dipl.-Inform. Friedrich Gretz (since Feb. 2011)
  Dr. Tingting Han (until Feb. 2011)
  Dipl.-Inform. Jonathan Heinen
  Dipl.-Inform. Christina Jansen
  Dr. Etienne Lozes (until Jul. 2011)
  (Alexander-von-Humboldt Stipendium)
  Dr. Alexandru Mereacre (until Jan. 2011)
  Viet Yen Nguyen, M.Sc.
  Dipl.-Inform. Maximilian R. Odenbrett
  Arpit Sharma, M.Sc.
  Falak Sher, M.Sc.
  Dipl.-Inform. Sabrina von Styp
  Dipl.-Inform. Haidi Yue

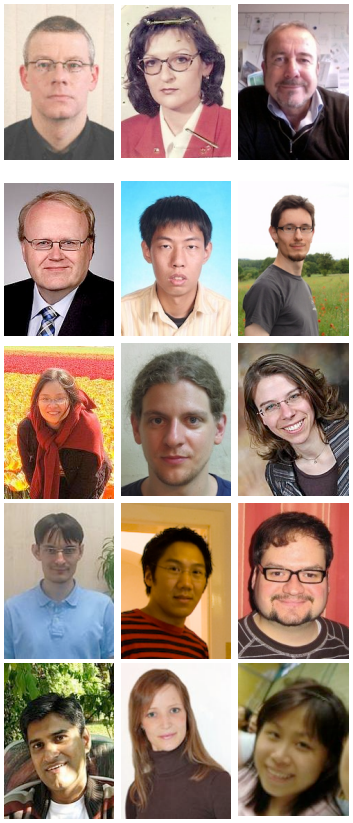

- **Diploma/Bachelor/Master Students**
  Markus Bals
  Henrik Barthels
  Lisa von Büttner
  Silvio de Carolis
  Jonas Dederichs
  Christian Dehnert
  Bernhard Ern
  Hussein Hamid Baagil
  Dirk Hauptmann
  Tobias Hoffmann
  Rafal Korzeniewski
  Gereon Kremer
  Bart Postma
  Hauke Schaper
  Stephen Wu

- **Student Researchers**
  Christian Dehnert
  Aamer Shah
  Hannah Spitzer
  Dennis Guck

- **Visiting Scientists**
  Alessandro Abate (TU Delft, NL)
  Albert Benveniste (IRISA Rennes, F)
  Dragan Bosnacki (TU Eindhoven, NL)
  Marco Bozzano (FBK Trento, I)
  Souymodip Chakraborty (IIT Delhi, India)
  Taolue Chen (University of Oxford, UK)
  Andreas Classen (University of Namur, B)
  Pedro D'Argenio (University of Cordoba, Argentina)
  Hubert Garavel (INRIA Rhone-Alpes, F)
  Holger Hermanns (Saarland University, D)
  Daniel Kroening (University of Oxford, UK)
  John Lygeros (ETH Zurich, CH)
  Hua Mao (Aalborg University, DK)
  Alexandru Mereacre (University of Oxford, UK)
  Vivek Sarkar (Rice University, USA)
  Arnaud Sangnier (LIAFA, F)
  Mani Swaminathan (Universty of Oldenburg, D)
  Mark Timmer (University of Twente, NL)
  Yuri Yushtein (ESTEC, NL)
  Anton J. Wijs (TU Eindhoven, NL)

Overview

Enormous progress has been booked in our research during 2011 which is reflected in the list of publications at the end of this section. In particular, the results of the ESA COMPASS project have had a large impact in the last year. The availability of a tool-set together with the successful application of the developed techniques by companies, have led to invited talks, invited tutorials (e.g., on one the largest European aerospace exhibitions), and formal inquiries from several (large) companies. This year the GUI will be finished, and an in-house case study for a new satellite system at ESA has been completed. It is great to see – and enormously stimulating! – that verification and dependability analysis techniques that are based on a firm theoretical basis indeed can have practical impact.

In September the Aachen Concurrency and Dependability Week attracted 300 attendants from all around the globe to our university. This week, organised by our chair, consisted of:

- the 22nd International Conference on Concurrency Theory (CONCUR);

- the 8th International Conference on Quantitative Evaluation of Systems (QEST);

- the 6th International Symposium on Trustworthy Global Computing (TGC).

In addition 9 workshops and 6 tutorials were held by world experts in the field. The social program consisted of a reception in the historic Ballsaal in Aachen (see picture), and a guided tour, a boat trip and a banquet (on the boat) in Maastricht, The Netherlands. The ability to visit 3 countries in (less than) a single day excited many of the attendants. To conclude, ACDW 2011 was a great success both scientifically as well as socially.

Other highlights during 2011:

- Joost-Pieter Katoen received the Educational Price 2010 for the course on Data Structures and Algorithms;
- Alexandru Mereacre finished his doctoral studies with honours;
- Falak Sher received an award for the best Master Thesis in the Department of Computer Science in 2010.

On the personnel side, Tingting Han (postdoc) and Alexandru Mereacre (PhD student) both left to join the group of Marta Kwiatkowska at the University of Oxford (UK), whereas Henrik Bohnenkamp continued his career in industry. In addition, our Alexander-von-Humboldt visiting professor, Etienne Lozes (ENS Cachan, F) left to join the group of Martin Lange in Kassel. As new member to the chair, we welcome Friedrich Gretz, who started early this year as PhD student in the Research Training Group AlgoSyn.

To conclude: 2011 was a busy, successful and exciting year, and we are looking forward to what 2012 will bring us.

Joost-Pieter Katoen.

Research Projects

Verification of Quantitative Properties of Embedded Software

T. Han, J.-P. Katoen

funded by the Dutch Research Council (NWO)

Embedded software typically executes on devices that, first and foremost, are not personal computers. Due to its embedded nature, its robustness is of prime importance, and timely reactions to stimuli from its -- mostly physical -- environment are essential. The aim of the QUPES project is to assess these quantitative aspects (e.g., timeliness and robustness) as an integral part of the embedded software validation phase.

To accomplish this, probabilistic model-checking techniques can be applied for models that are equipped with randomness and variants thereof which also exhibit nondeterminism. Based on efficient numerical methods and abstraction techniques, quantitative properties can be checked automatically even on large state spaces with millions of states using dedicated tools. Opposed to, amongst others, the essential feature of model checking, where evidences will be provided on a property refutation, counterexample generation in probabilistic model checking is almost not developed. We provide the theoretical and algorithmic foundations for counterexample generation in probabilistic model checking, in particular for discrete-time Markov chains. One of the key principles is the casting of the concepts of strongest evidence and smallest counterexample as (variants of) shortest path problems. This enabled the use of efficient and well-studied graph algorithms for counterexample generation. These results can be extended to Markov chains with rewards, to Markov decision processes (MDPs), to LTL model checking, and have been recently been adopted in probabilistic counterexample-guided abstraction-refinement (CEGAR) techniques for MDPs as well as in counterexample generation for continuous-time Markov chains (CTMC) and cpCTL logic. Compact representation of a counterexample by regular expressions are also studied.

Further, compositional reasoning is a key strategy in analysing complex systems as it allows the use of hierarchical and modular modelling formalisms like stochastic process algebras, stochastic activity networks or generalised stochastic Petri nets. Continuous-time Markov Decision processes (CTMDPs) are the nondeterministic counterpart of the aforementioned CTMCs and are well suited for compositional verification techniques. We define stochastic logics (like CSL) on CTMDPs and provide their measure-theoretic basis. Further, well-known equivalences like strong and weak bisimulation relations are adapted to CTMDPs which considerably reduce the state-space needed for quantitative analysis.

Verifying Pointer Programs with Unbounded Heap Structures

J. Heinen, C. Jansen, J.-P. Katoen, T. 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.

After considering (possibly cyclic) singly-linked list data structures, the approach was extended to analyse programs that handle more complex dynamic data structures. 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.

The central issues in this context are correctness, usability, and efficiency. The first essentially boils down to the requirement that a nonterminal can always be concretised to the data structure from which it was abstracted. To ensure his property, we defined a novel normal form for hyperedge replacement grammars that is inspired by the well-known Greibach normal form for string grammars. Moreover we developed an algorithm for constructing a normalised grammar from a given hyperedge replacement grammar with bounded degree.

To improve the usability of the overall approach, the idea is to adopt learning techniques to automatically infer abstraction grammars for data structures. More concretely this means that the heap configurations arising during the execution of the given pointer-manipulating program have to be inspected at runtime, and that hypergraph production rules for generating these structures have to be found. This proceeds in an incremental fashion as more and more heap structures are created during runtime. The automatic generation of corresponding rules then circumvents the complex and error-prone procedure of developing grammars manually.

The incremental construction of grammars also raises new challenges with regard to the Greibach normal form that was mentioned in the previous paragraph. Efficiency will be improved by avoiding the re-computation of the normal form for the whole grammar, instead using an incremental approach where the normal form of the extended grammar is obtained by adding new Greibach rules to the normal form of the previous grammar. A first step towards this direction was taken by developing an incremental algorithm for deriving the Greibach normal form of a string grammar.

Another step for improving the efficiency of the approach was taken by developing an automata-theoretic concept for finding hypergraph embeddings, which is required for implementing grammar-based heap abstractions.

COMPASS: Correctness, Modelling and Performance of Aerospace Systems

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

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)

In this project we develop a model-based approach to system-software co-engineering which is tailored to the specific characteristics of critical on-board systems for the space domain. The approach is supported by a System-Level Integrated Modelling (SLIM) Language in which engineers are provided with convenient ways to specify a.o. nominal hardware, as well as software operations, timed and hybrid behaviour, (probabilistic) faults and their propagation, error recovery and degraded modes of operation. This language is based on the Architecture Analysis and Design Language (AADL) and its Error Model Annex which allows for the modelling of error behaviour. A kernel of the SLIM Language is equipped with a formal semantics that provides the interpretation of SLIM specifications in a precise and unambiguous manner. Systems are considered as a hierarchy of (hardware and software) components which are defined by their type (interface) and implementation. Components interact via ports allowing for both message-oriented and continuous communication. The internal structure of a component implementation is specified by its decomposition into subcomponents, together with their HW/SW bindings and their interaction via connections over ports. Component behaviour is specified by a textual description of mode-transition diagrams. System reconfiguration is supported by mode-dependent presence of subcomponents and their connections. Error behaviour is described by probabilistic finite state machines, where error delays may be governed by continuous random variables.

Correctness properties, safety guarantees, and performance and dependability requirements are specified using requirement specification patterns which act as parametrized "templates" to the engineers and thus offer a comprehensible and easy-to-use framework for requirement specification.

The properties are checked on the SLIM specification using rigorous analysis methods. The precise character of these techniques together with the formal semantics of SLIM yield a trustworthy modelling and analysis framework for system and software engineers. The formal analysis is based on state-of-the-art model checking techniques such as bounded SAT-based and symbolic model checking, and extensions of model checking with numerical and simulative means to reason about quantitative requirements such as performance and dependability. The analysis facilities support, among others: automated derivation of dynamic (i.e., randomly timed) fault trees, Failure Modes and Effects Analysis (FMEA) tables, assessment of Fault Detection, Isolation, and Recovery (FDIR) measures, and observability requirements for effective diagnosability by FDIR.

An integrated platform on top of state-of-the-art tools with an accompanying graphical user interface is available, and has been evaluated by Thales Alenia Space using several case stud-

ies involving critical on-board computer-based systems from the satellite domain. Another outcome of the project is an extension of AADL's Error Model Annex and the corresponding semantics.

Present activities concentrate on improving both the applicability and efficiency of the toolset. Currently the SLIM models to be analysed have to be provided in a textual format that is based on the syntax of AADL. This requires users to perform text-based system modelling. The usability and adoption of the COMPASS toolset will greatly benefit from the possibility of graphical model ling, allowing engineers to specify systems in a graphical notation that eases the understanding of their architecture and behaviour. To this aim, a graphical editor for SLIM models is being under development.

Moreover we are addressing some shortcomings and limitations of the current dependability and performability analysis techniques of AADL models that been identified during evaluation. 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. Corresponding approaches are currently investigated in the framework of a PhD on "Extending and Improving Formal Methods for System/Software Co-Engineering", sponsored by both ESA and Thales Alenia Space.

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

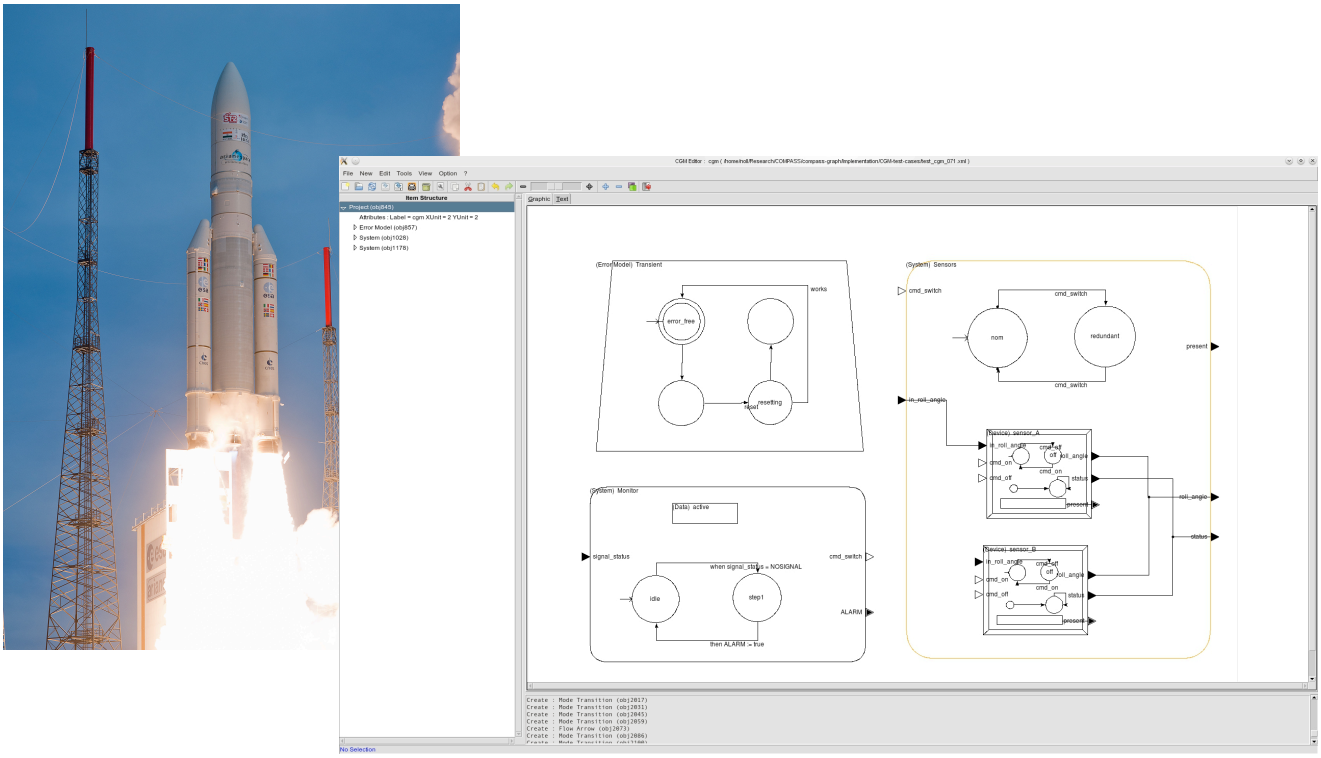

Quasimodo: Quantitive System Properties in Model-Driven Design of Embedded Systems

H. Bohnenkamp, H. Yue, J.-P. Katoen

EU FP7 project, coordinator: Aalborg University (K.G. Larsen)

The MOVES group participates in the European research project "Quasimodo", funded by the European Commission under the IST framework programme 7 for Information and Communication Technology, ICT. The objective of this project is to develop theory, techniques and tool components for handling quantitative (e.g. real-time, hybrid and stochastic) constraints in model-driven development of real-time embedded systems. Ultimate aim is to increase the competitiveness of European industrial companies which develop, implement and deploy embedded systems.

More specifically, the project aims are:

1. Improving the modelling of diverse quantitative aspects of embedded systems.
2. Providing a wide range of powerful techniques for analysing models with quantitative information and for establishing abstraction relations between them.
3. Generating code with predictable behaviour from quantitative models.
4. Improving the overall quality of testing by using suitable quantitative models as the basis for generating sound and complete test cases.
5. Applying the techniques to real-life case studies and disseminating the results to industry.

Project partners are universities, research institutes, and companies in Germany, The Netherlands, Denmark, Belgium, and France.

UMIC: Ultra High Speed Mobile Information and Communication

H. Yue, H. Bohnenkamp, J.-P. Katoen

The evaluation of the quality-of-service of Wireless Sensor Networks is mostly done by simulation. In the context of the Quasimodo and UMIC projects, 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 energy as low as possible.

The new GMAC protocol is modelled with a specific radio model: 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 networks. However, it also delayed message propagation. Furthermore, 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.

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

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.

SYRUP: SYmbolic RedUction of Probabilistic Models

J.-P. Katoen, C. Dehnert

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 techniques 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

J.-P. Katoen, A. Mereacre, 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 applied 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.

Synthesising of Model Based Testing for Process Control Engineering

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

G. Quirós, U. Epple (Chair of Process Control Engineering)

In Process Control Engineering controller for plants that are not correct with respect to their specification can cause fatal disasters, e.g. when tanks with acid run over, people get injured or too 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 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 in oder to reduce costs, time and therefore to automate and systematise the testing process.

In order 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 a simple controller, i.e. a motor controller, which allowed 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 behavior is an important feature of the plant controllers, e.g. the next action depends on the current filling level of a tank and it may be 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 semantic is defined. This semantic then is needed 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 applications 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.

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 formula 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.

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 check if two given system are equivalent under a pre-established equivalence relation. In model checking, the task is to check if 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 probabilistic setting, or instead discover new techniques if necessary.

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 to (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, Markov decision processes and non-probabilistic systems.

Efficient Multi-Core Model Checking

M. R. Odenbrett

joint work with dr. Dragan Bosnacki, dr.ing. Anton J. Wijs,
prof.dr. Mark G. J. van den Brand and prof.dr. Peter A. J. Hilbers
from Eindhoven University of Technology

funded by the Dutch Research Council (NWO)

Our project aims at developing new algorithms for model checking, including probabilistic and stochastic model checking, that can exploit the parallelism of multi-core systems. We consider multi-core CPU as well as many-core GPGPU (CUDA) hardware architectures.

The main motivation for our work arises, as for so many other projects on model checking, from the well-known state space explosion problem. It limits the practicability of model checking by two important factors: memory and run-time. While long run-times can be seen as an annoying necessity to prove correctness, the memory requirements are strict hardware-based limiting factors. However, since the shift from 32 to 64 bit word sizes, modern computer architectures can address 2^{32} times more memory than before. Consequently, the memory bottleneck is relieved, at least for the time being. This led us to the conclusion, that now the run-time itself should be addressed. Given the fact that Moore's Law (doubling of transistor density roughly every two years) does not imply a corresponding increase of processor clock rates anymore but now leads to multi- and many-core processors, shared memory parallelisation of the model checking problems seems to us to be the way to go.

We plan to develop prototype implementations of the new algorithms in model checkers, like Spin and its extensions, as well as the probabilistic model checker MRMC. The prototype implementations will be validated on case studies including models of biological systems.

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

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 communications 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

- Member 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)
- Member of the Steering Committee of QEST (Quantitative Evaluation of Systems).
- 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.
- Board Member of the Dutch Society on Theoretical Computer Science (NVTI).
- Senior Member of the Association of Computing Machinery (ACM).
- **Member of the Program Committee of the following events:**
  - Fundamental Approaches to Software Engineering (FASE 2011)
  - International Colloquium on Automata, Languages and Programming (ICALP 2011)
  - Perspectives of System Informatics (PSI 2011)
  - Concurrency Theory (CONCUR 2011) (co-chair)
  - Computer Aided Verification (CAV 2011).
- **Invited speaker at:**
  - LIAFA Workshop on Automata and Logic for Data Manipulating Programs, Paris, France.
  - ETAPS Workshop on Hybrid Autonomous Systems (HAS), Saarbrücken
  - 18th Summer School on Computer Science (RIO), Rio Cuarto, Argentina.
  - 4th Summer School on Verification Technology, Systems & Applications (VTSA), Liege, Belgium.
  - 5th Int. Workshop on Reachability Problems (RP), Genova, Italy.
  - 12th Int. Workshop on Formal Methods for Industrial Critical Systems (FMICS). Trento, Italy, 2011.
  - International Symposium on Interdisciplinary Modelling of Cyber-Physical Systems (IMCPS), Manchester, UK.
  - 4th Conference on Fundamentals of Software Engineering (FSEN), Tehran, Iran.
- Member of the IFIP Working Group 1.8 on Concurrency Theory.
- Member of the IFIP Working Group 2.2 on Programming Concepts

- Member of the EPSRC Review College (Engineering and Physical Sciences Research Council), UK.
- General Chair of the Aachen Concurrency and Dependability Week:
  - 22nd Int. Conference on Concurrency Theory (CONCUR)
  - 8th Int. Conference on Quantitative Evaluation of Systems (QEST)
  - 6th Int. Symposium on Trustworthy Global Computing (TGC)
- Member of several external international PhD committees.
- Chairman of Selection Committee of Full Professorship on Logics and Automata Theory.
- Chairman of Evaluation Committee of Junior-Professorship on Theory of Hybrid Systems.
- Chairman of the Examination Board of Department of Computer Science.

Thomas Noll

- Student advisor for the following applied subjects within CS:
  Electrical Engineering, Civil Engineering, and Medicine
- Member of CS Commission for Teaching Service
- **Member of the Program Committee of the following events:**
  - Software Engineering Track at the 27th Annual ACM Symposium on Applied Computing (SAC 2012)
  - 2nd Analytic Virtual Integration of Cyber-Physical Systems Workshop (AVICPS 2011) at the 32nd IEEE Real-Time Systems Symposium (RTSS 2011)
  - 5th IEEE International Conference on Secure System Integration and Reliability Improvement (SSIRI 2011)
  - Software Engineering Track at the 26th Annual ACM Symposium on Applied Computing (SAC 2011)

Talks and Publications

Talks

Hongfei Fu. Deciding Probabilistic Simulation between Probabilistic Pushdown Automata and Finite-State Systems. Talk, ROCKS workshop, Herrsching, 2011. Workshop Presentation.

Hongfei Fu. Deciding Probabilistic Simulation between Probabilistic Pushdown Automata and Finite-State Systems. Talk, YR-CONCUR, 2011. Workshop Presentation.

Hongfei Fu. Model Checking EGF on Basic Parallel Processes. Talk, ATVA, 2011. Conference Presentation

Friedrich Gretz. Reasoning about probabilistic programs. Talk, YR-CONCUR, 2011 Workshop Presentation.

Friedrich Gretz. Operational Semantics for a probabilistic guarded command language. Talk, ROCKS workshop, Herrsching, 2011. Workshop Presentation.

Jonathan Heinen. Juggrnaut - An Abstract JVM. Talk, 2nd International Conference on Formal Verification of Object-Oriented Software (FoVeOOS) Turin, Italy, 2011.

Jonathan Heinen. Verifying Pointer Programs Using Graph Grammars (Theory and Practice). Talk, D-CON, Münster, 2011.

Jonathan Heinen. Verificación de Programas Orientados a Objetos. Talk, Universidad Nacional de Colombia, 2011.

Jonathan Heinen. Verification of Object Oriented Programs. Talk, Universität Duisburg-Essen, 2011.

Christina Jansen. Heap Abstraction by means of Hyperedge Replacement Grammars. Talk, Summer School Marktoberdorf, 2011. Student presentation.

Christina Jansen. Heap Abstraction by means of Hyperedge Replacement Grammars. Talk, YR-CONCUR, 2011.

Christina Jansen. A Local Greibach Normal Form for Hyperedge Replacement Grammars. Talk, LATA, 2011.

Joost-Pieter Katoen. Can your CTMC keep up with your timed automaton? Talk, Quasimodo Meeting at Hydac, Saarbrücken, 2011. Project presentation.

Joost-Pieter Katoen. Efficient CTMC Model Checking of Linear Real-Time Objectives. Talk, TACAS Conference, 2011. Paper presentation.

Joost-Pieter Katoen. Can your CTMC keep up with your timed automaton?. Talk, Quasimodo Meeting at Hydac, Saarbrücken, 2011. Project presentation.

Joost-Pieter Katoen. Efficient CTMC Model Checking of Linear Real-Time Objectives. Talk, TACAS Conference, 2011. Paper presentation.

Joost-Pieter Katoen. Approximate model checking of stochastic hybrid systems. Talk, Hybrid Autonomous Systems (HAS) Workshop, 2011. Keynote presentation.

Joost-Pieter Katoen. Verifying Markov Chains. Five Lectures at RIO Summerschool, Rio Cuarto, Argentina, 2011. Invited Lecturer.

Joost-Pieter Katoen. One Can Do Much More with Model Checking Than You Think!. Talk, Foundations of Software Engineering (FSEN), Tehran, Iran, 2011. Invited Talk.

Joost-Pieter Katoen. Can your CTMC keep up with your timed automaton?. Talk, Technical University Delft, The Netherlands, 2011. Project meeting EU project MoVeS.

Joost-Pieter Katoen. Observing Markov Chains by Timed Automata. Talk, Interdisciplinary Modelling of Cyber-Physical Systems (IMCPS), Manchester, UK, 2011. Invited Talk.

Joost-Pieter Katoen. Probabilistic Model Checking. Two Lectures at Formal and Interdisciplinary models In Resilience Engineering (FIRE'11), Manchester, UK, 2011. Invited Lecturer.

Joost-Pieter Katoen. The Google Search Engine. Talk, Ringvorlesung Computer Science, RWTH Aachen University, 2011.

Joost-Pieter Katoen. Towards Trustworthy Aerospace Systems: An Experience Report. Talk, 11th Formal Methods for Industrial Critical Systems Workshop (FMICS), Trento, Italy, 2011. Invited Talk.

Joost-Pieter Katoen. Verifying Markov Chains. Two Lectures at Summerschool on Verification Technology, Systems and Applications (VTSA), Liege, Belgium, 2011. Invited Lecturer.

Joost-Pieter Katoen. Observing Stochastic Processes by Timed Automata. Talk, 5th International Workshop on Reachability Problems (RP), Genova, Italy, 2011. Invited Talk.

Joost-Pieter Katoen. Observing Stochastic Processes by Timed Automata. Talk, CNR/ISTI Research Labs, Pisa, Italy, 2011.

Joost-Pieter Katoen. Verifying Pointer Programs Using Graph Grammars. Talk, LIAFA Workshop on Automata and Logic for Data Manipulating Programs, 2010. Invited Talk.

Thomas Noll. The ESA COMPASS Project: Correctness, Safety and Fault Tolerance in Aerospace Systems. Talk, GI-Themenabend "Fehlerfreie Software – ein Widerspruch in sich?", Cologne, D, 2011.

Thomas Noll. Analyzing Reconfigurable Component-Based Systems Using Attribute Grammars. Talk, 8th International Symposium on Formal Aspects of Component Software, Oslo, Norway, 2011.

Maximilian R. Odenbrett, and Anton J. Wijs. Using GPGPUs for Bioinformatics. Talk, IPA Spring Days, Vlijmen, 2011.

Maximilian R. Odenbrett, and Anton J. Wijs. Efficient Reconstruction of Genetic Networks via Transitive Reduction on GPGPUs. Talk, CWI, Amsterdam, 2011.

Arpit Sharma. Weighted Lumpability on Markov Chains. Talk, ROCKS Workshop, ETAPS., 2011. Workshop presentation.

Arpit Sharma. Weighted Lumpability on Markov Chains. PSI Conference, Novosibirsk, Russia, 2011. Conference presentation.

Falak Sher. Abstraction Techniques for Markov Automata. Talk, Saarland University, 2011.

Sabrina von Styp. Model-Based Testing in Process Control Engineering. Talk, AlgoSyn in Dagstuhl, 2011.

Sabrina von Styp. Towards a Theory for Timed and Symbolic Testing. Talk, TAROT, 2011.

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

Publications

Alessandro Abate, Joost-Pieter Katoen, and Alexandru Mereacre. Quantitative Automata Model Checking of Autonomous Stochastic Hybrid Systems. In 14th ACM International Conference on Hybrid Systems: Computation and Control (HSCC). pages 83–92. ACM Press, 2011.

Alessandro Abate, Joost-Pieter Katoen, John Lygeros, and Maria Prandini. A two-step scheme for approximate model checking of stochastic hybrid systems. In Proceedings 18th IFAC World Congress 2011. Volume 18 of IFAC-PapersOnLine. Elsevier, 2011.

Benoit Barbot, Taolue Chen, Tingting Han, Joost-Pieter Katoen, and Alexandru Mereacre. Efficient CTMC Model Checking of Linear Real-Time Objectives. In Tools and Algorithms for the Construction and Analysis of Systems (TACAS). pages 128–142. Volume 6605 of LNCS. Springer-Verlag, 2011.

Benoit Delahaye, Joost-Pieter Katoen, Kim G. Larsen, Axel Legay, Mikkel Pedersen, Falak Sher, and Andrzej Wasowski. Abstract Probabilistic Automata. In Verification, Model Checking and Abstract Interpretation (VMCAI). pages 324–339. Volume 6538 of LNCS. Springer-Verlag, 2011.

Delahaye, Benoit, Katoen, Joost-Pieter, Larsen, Kim G., Legay, Axel, Pedersen, Mikkel, Sher, Falak and Wasowski, Andrzej, New Results on Abstract Probabilistic Automata, in: Applications of Concurrency to System Design (ACSD), pages 118-127, IEEE CS Press, 2011.

Hongfei Fu. Model Checking EGF on Basic Parallel Processes. In 9th International Symposium on Automated Technology for Verification and Analysis (ATVA). pages 120–134. Volume 6996 of LNCS. Springer-Verlag, 2011.

E. Moritz Hahn, Tingting Han, and Lijun Zhang. Synthesis for PCTL in Parametric Markov Decision Processes. In NASA Formal Methods - Third International Symposium (NFM). pages 146–161. Volume 6617 of LNCS. Springer-Verlag 2011.

Jonathan Heinen, and Christina Jansen. Juggrnaut - An Abstract JVM. In Papers presented at the 2nd International Conference on Formal Verification of Object-Oriented Software (FoVeOOS). pages 226–243. Karlsruhe Reports in Informatics 2011-26. Karlsruhe Institute of Technology, 2011.

Jonathan Heinen, Thomas Noll, and Stefan Rieger. Juggrnaut: Graph Grammar Abstraction for Unbounded Heap Structures. In Proc. 3rd Int. Workshop on Harnessing Theories for Tool Support in Software (TTSS). pages 93–107. Volume 266 of ENTCS. Elsevier, 2010.

Christina Jansen, Jonathan Heinen, Joost-Pieter Katoen, and Thomas Noll. A Local Greibach Normal Form for Hyperedge Replacement Grammars. In Proc. of 5th Int. Conf. on Language and Automata Theory and Applications (LATA). pages 323–335. Volume 6638 of LNCS. Springer-Verlag, 2011.

Joost-Pieter Katoen, Ivan S. Zapreev, E. Moritz Hahn, Holger Hermanns, and David N. Jansen. The Ins and Outs of the Probabilistic Model Checker MRMC. Performance Evaluation, 68(2):90–104, 2011.

Joost-Pieter Katoen. Towards Trustworthy Aerospace Systems: An Experience Report. In 16th International Workshop on Formal Methods for Industrial Critical Systems (FMICS). pages 1–4. Volume 6959 of LNCS. Springer-Verlag, 2011.

Joost-Pieter Katoen, and Thomas Noll. Trustworthy Aerospace Systems. Public Service Review: European Science and Technology, 11:204–205, 2011.

Joost-Pieter Katoen, and Barbara König, editors, 22nd Conference on Concurrency Theory (CONCUR), Volume 6901 of LNCS. Springer-Verlag, 2011.

Daniel Klink, Anne Remke, Boudewijn R. Haverkort, and Joost-Pieter Katoen. Time-Bounded Reachability in Tree-Structured QBDs by Abstraction. Performance Evaluation, 68(2):105–125, 2011.

Alexandru Mereacre. Verification of Continuous-Space Stochastic Systems. PhD Thesis, RWTH Aachen University, 2011.

Ralf Mitsching, Frank Fiedler, Henrik Bohnenkamp, Carsten Weise, and Stefan Kowalewski. TripleT: Improving Test Responsiveness for High Performance Embedded Systems. In Proc. 4th IEEE International Conference on Software Testing, Verification, and Validation (ICSTW). pages 67-74. IEEE CS Press 2011.

Bastian Schlich, Thomas Noll, Jörg Brauer, and Lucas Brutschy. Reduction of Interrupt Handler Executions for Model Checking Embedded Software. In Proc. of 5th Int. Haifa Verification Conference (HVC). pages 5–20. Volume 6405 of LNCS. Springer-Verlag, 2011.

Haidi Yue, Henrik Bohnenkamp, Malte Kampschulte, and Joost-Pieter Katoen. Analysing and Improving Energy Efficiency of Distributed Slotted Aloha. In 11th International Conference on Next Generation Wired/Wireless Advanced Networking (NEW2AN). pages 197–208. Volume 6869 of LNCS. Springer-Verlag, 2011.

Yuri Yushtein, Marco Bozzano, Alessandro Cimatti, Joost-Pieter Katoen, Viet Yen Nguyen, Thomas Noll, Xavier Olive, and Marco Roveri. System-Software Co-Engineering: Dependability and Safety Perspective. In 4th IEEE International Conference on Space Mission Challenges in Information Technology (SMC-IT). pages 18–25. IEEE CS Press, 2011.

Technical Reports

Jonathan Heinen, and Christina Jansen. Juggrnaut - An Abstract JVM. Technical Report AIB 2011-21 , RWTH Aachen University, Germany , 2011.

Christina Jansen, Jonathan Heinen, Joost-Pieter Katoen, and Thomas Noll. A Local Greibach Normal Form for Hyperedge Replacement Grammars. Technical Report AIB 2011-15, RWTH Aachen University, Germany , 2011.