07101 Abstracts Collection
Quantitative Aspects of Embedded Systems
— Dagstuhl Seminar —

Boudewijn Haverkort[1], Joost-Pieter Katoen[2] and Lothar Thiele[3]

[1] Univ. of Twente, NL
brh@cs.utwente.nl
[2] RWTH Aachen, DE
katoen@cs.rwth-aachen.de
[3] ETH Zürich, CH
thiele@tik.ee.ethz.ch

Abstract. From March 5 to March 9, 2007, the Dagstuhl Seminar 07101 "Quantitative Aspects of Embedded Systems" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Embedded systems, quantitative analysis, model checking, real-time calculus, validation, verification, model-based and model-driven design

1 07101 Executive Summary – Quantitative Aspects of Embedded Systems

Embedded software systems

Embedded software controls the core functionality of many systems: it controls telephone switches and satellites, drives the climate control in our offices and cars, runs pacemakers, is at the heart of our power plants, and makes our cars and TVs work properly. As such systems are massively encroaching on daily life, our reliance on embedded software is growing rapidly. But, how justifiable is this reliance?

Whereas traditional software has a rather transformational nature mapping input data onto output data, *embedded software is different* in many respects. Most importantly, embedded software is subject to complex and permanent interactions with their—mostly physical—environment via sensors and actuators. Typically, software in embedded systems does not terminate and interaction usually takes place with multiple concurrent processes at the same time. Reactions

to the stimuli provided by the environment should be prompt (timeliness or responsiveness), i.e., the software has to "keep up" with the speed of the processes with which it interacts.

Furthermore, characteristic for embedded systems is that they have to meet a multitude of quantitative constraints. These constraints involve the resources that a system may use (computation resources, power consumption, memory usage, communication bandwidth, costs, etc.), assumptions about the environment in which it operates (task arrival rates, task sizes), requirements on the services that the system has to provide (timing constraints, performance, response time) and requirements of the continuity with which these services are delivered (availability, dependability, fault tolerancy, etc.).

The observed difference between traditional software and embedded software has recently led to initiatives for a variety of dedicated international conferences and journals. Also, in various countries research institutes on embedded systems have been set up with a strong industrial cooperation, for instance, in Denmark (CISS), the Netherlands (ESI), and the US (CHESS).

A lack of quantitative assessment

Despite the importance of the quantitative constraints for the well-operation of embedded software systems, the proper assessment of cost, resources, performance, dependability, robustness, etc., often comes as an afterthought. It is rather common for embedded software to be fully designed and functionally tested before any attempt is undertaken to determine its performance, dependability or resource-usage characteristics. One of the main reasons for this situation is that well-developed and rigorous evaluation techniques for non-functional, i.e., quantitative system aspects have not become an integral part of standard software engineering practice. This undesirable situation has led to the increased interest by embedded software researchers to extend the usual functional specification and properties with a set of "performance indices", e.g., stated in terms of costs, timeliness, speed and the like, and constraints on these indices. Also in industry, a growing interest in assessing non-functional aspects of embedded systems as early as possible in the system design life cycle can be witnessed.

Where are we going?

Model-Driven Development (MDD) is a new software development technique in which the primary software artifact is a model. Ideally, the MDD technique allows engineers to (graphically) model the requirements, behaviour and functionality of computer-based systems. The design is iteratively analysed, validated, and tested throughout the development process, and automatically generated production-quality code can be output in a variety of languages.

Existing MDD tools for embedded systems are rather sophisticated in handling functional requirements but their treatment of quantitative constraints is still in development. Although methods for verification of real-time system designs, using for instance timed automata, are being developed, these methods

are not yet mature enough for dealing with larger industrial embedded systems. Hence, MDD will not realise its full potential in the embedded systems area unless the ability to handle quantitative properties is drastically improved.

In contrast to the situation in the design of embedded software systems, in the design of computer-communication systems, **quantitative methods** to determine the quality of the system, expressed in terms of throughput or response time, have been used for a long time. Next to methods from classical queueing theory and discrete-event simulation, recently the use of analytical/numerical methods for evaluating complex systems has become more widespread. This has lead to methods and techniques to specify complex system behaviour using formal methods (like Petri nets, process algebra) enhanced with time and probabilities. Subsequently, appropriate models are generated from these high-level models, which, after numerical analysis, provide detailed insight in performance and dependability measures of interest.

Over the last, say, 5 years, very good progress has been made in pairing the above quantitative techniques to techniques known from the verification area, esp. model checking of properties specified in logics like CSL (an extension of CTL with stochastic time). This has led to model checking algorithms and tools for Markovian models of system. The state-of-the-art, however, is still such that expert knowledge is required to use these techniques, hence, large-scale application in embedded software system design and implementation is still a dream rather than a reality.

Of course, also known quantitative techniques from the area of real-time systems (classical ones, such as EDF, or more advanced compositional ones), or methods known from network calculus (originally developed for dimensioning communication networks at a high level of abstraction), data flow graphs, and so on, can, and probably should be used as part of the embedded system design.

What is clear, though, is that all of the above techniques can only be used during the design of embedded systems after appropriate adaptation and embedding in a design trajectory, e.g., based on MMD.

Furthermore, were each of the above mentioned approaches has its strengths and weaknesses, an important first task is to map these strength and weaknesses (applicability, scope, modelling power, costs of evaluation, etc.). A second and more challenging question is then how to combine or integrate these methods. Such questions can only be answered when key researchers for these various approaches come together and exchange and discuss their ideas.

Seminar goal

Given the above considerations, the goal of this Dagstuhl seminar has been to bring together experts in the areas of embedded software design and implementation, model-based analysis of quantitative system aspects, and researchers working on extending all kinds of formal (design and analysis) methods with quantitative system aspects. These three areas are clearly well-related in the context of embedded systems, but have not been addressed as such in the past, as they have been worked upon in different communities. Thus, the seminar will

lay bridges between these three areas, so that knowledge and experience can be shared, transferred and, ultimately, be generated.

Seminar results

The results of the seminar can be classified in four areas:

- Three tutorials have been presented, that helped in bringing together the variety of disciplines involved in model-based embedded (software) system design:
  - "Worst-case estimation methods" by Kim Larsen and Reinhard Wilhelm;
  - "Stochastic model checking" by Holger Hermanns;
  - "Model-based design for embedded systems" by Pieter Mosterman.
- 24 regular talks across the whole spectrum of topics were presented.
- A discussion session was held about which ingedients are necessary for a (master) curriculum on embedded systems. Moderated by Holger Hermanns, Reinhard Wilhelm reported about the ARTIST embedded system curriculum design, and Gerard Smit reported about the initiative at the 3 technical universities in the Netherlands to start a joint embedded system curriculum between the three electrical engineering and the three computer science departments.
- At the first day of the seminar working groups were formed, that partly later split and joined, but that led to a number of peer working group meetings and working group reports. At the end of the seminars, three working groups reported on their findings about "worst-case versus stochastic methods", "how good are formal methods for model-based design", and "performance measures other than time".

In all of the four mentioned areas, the sharing and transfer of ideas between researchers with different background and viewpoints played a foreground role. This role, as such, can be regarded an important result in itself (the process, rather than the product). Cooperations have been started, and mutual interest in each others areas, conferences and scientific communities has been aroused.

Slides of the Program

These slides contain the programme of the seminar, as executed during the week of March 5–9, 2007. At March 12, 2007, the seminar was presented by participant Holger Hermanns (University Saarland) at a meeting of the Dagstuhl executive board. The slides used for that presentation (with contributions from Holger Hermanns) have been added as well.

Joint work of: Haverkort, Boudewijn; Katoen, Joost-Pieter; Thiele, Lothar

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2007/1137

2 Abstracts of regular presentations

Minimizing Cache Conflicts by Optimal Task Placement

Sebastian Altmeyer (Universität des Saarlandes, D)

In non-preemptive embedded systems, precise timing guarantees can be computed and the overall performance of the system simply consists of the performances of all tasks, viewed independent from each other.

In preemptive systems, the situation gets more complex. The overall performance is also determined by the interference of the tasks during preemption. Hereby especially the interference on the cache matters.

We propose a new method to analyze and optimize the task interference on the cache already during compile time. Therefor, we optimize the memory placement of the systems' tasks. On the hand hand, we minimize the number of cache misses to optimize the performance and on the other hand, we try to ensure the persistence of cache-lines of hard real-time tasks during a single run of a task. If it is not possible to ensure cache persistence, we can classify cache-lines as safe or unsafe and use this information to derive timing guarantees also valid in case of preemption.

Dynamic fault tree analysis using I/O interactive Markov chains

Hichem Boudali (University of Twente, NL)

Dynamic fault trees (DFTs) are a versatile and common formalism to model and analyze the reliability of computer-based systems. This talk presents a formal semantics of DFTs in terms of input/output interactive Markov chains (I/O-IMCs), which extend continuous-time Markov chains with discrete input, output and internal actions.

This semantics provides a rigorous basis for the analysis of DFTs.

Our semantics is fully compositional, that is, the semantics of a DFT is expressed in terms of the semantics of its elements (i.e. basic events and gates). This enables an efficient analysis of DFTs through compositional aggregation, which helps to alleviate the state-space explosion problem by incrementally building the DFT state space.

Keywords: Reliabilty, dynamic fault trees, interactive Markov chains

Computing Battery Lifetime Distributions

Lucia Cloth (University of Twente, NL)

The usage of mobile devices like cell phones, navigation systems, or laptop computers, is limited by the lifetime of the included batteries.

This lifetimes depends naturally on the rate at which energy is consumed, however, it also depends on the usage pattern of the battery. Continuous drawing of a high current results in an excessive drop of residual capacity. However, during intervals with no or very small currents, batteries do recover to a certain extend.

We model this complex behaviour with an inhomogeneous Markov reward model, thereby following the approach of the so-called Kinetic battery Model (KiBaM).

The state-dependent reward rates thereby correspond to the power consumption of the attached device and to the available charge, respectively. We develop a tailored numerical algorithm for the computation of the distribution of the consumed energy and show how different workload patterns influence the overall lifetime of a battery.

Joint work of: Cloth, Lucia; Jongerden, Marijn; Haverkort, Boudewijn

Precision-Timed (PRET) Machines

Stephen A. Edwards (Columbia University, USA)

I argue that we need processors that provide predictable timing for real-time embedded systems.

Keywords: Embedded systems, processor architecture, real-time

Probabilistic Model Checking of Contention Resolution in the IEEE 802.15.4 Low-Rate Wireless Personal Area Network Protocol

Matthias Fruth (University of Birmingham, GB)

The international standard IEEE 802.15.4 defines low-rate wireless personal area networks, a central communication infrastructure of pervasive computing. In order to avoid conflicts caused by multiple devices transmitting at the same time, it uses a contention resolution algorithm based on randomised exponential backoff that is similar to the ones used in IEEE 802.3 for Ethernet and IEEE 802.11 for Wireless LAN.

We model the protocol using probabilistic timed automata, a formalism in which both nondeterministic and probabilistic choice can be represented. The probabilistic timed automaton is transformed into a finite-state Markov decision process via a property-preserving integral-time semantics. Using the probabilistic model checker PRISM, we verify correctness properties, compare different operation modes of the protocol, and analyse performance and accuracy of different model abstractions.

Keywords: Wireless networks, Contention resolution, Performance analysis, Probabilistic model checking, Probabilistic timed automata

Full Paper:
http://www.cs.bham.ac.uk/~mxf/docs/isola_2006.pdf

Measuring and Modeling Quantitative Aspects of Networked Embedded Systems

Reinhard German (Universität Erlangen, D)

In this talk I will report on some efforts in measuring and modeling networked embedded systems. For the measurements, three approaches will be illustrated: measuring one-way network delays, interrupt system response times, and internal communication delays between componants in a system. Afterwards I will illustrate a modeling approach in which UML-based modeling, discrete-event simulation, and test automation will be integrated. In the last part I will present two automotive applications in which various kinds of modeling and also measurements are necessary: transmission of safety-relevant sensor data and time-triggered architectures.

Bisimulation minimisation mostly speeds up probabilistic model checking

David N. Jansen (RWTH Aachen, D)

We studied the effect of bisimulation minimisation in model checking of monolithic discrete-time and continuous-time Markov chains. Our results show thatas for traditional model checkingenormous state space reductions (up to logarithmic savings) may be obtained. In contrast to traditional model checking, in many cases, the verification time of the original Markov chain exceeds the quotienting time plus the verification time of the quotient. We consider probabilistic bisimulation as well as versions thereof that are tailored to the property to be checked.

Keywords: Markov chain, bisimulation, model checking

Joint work of: Katoen, Joost-Pieter; Kemna, Tim; Zapreev, Ivan; Jansen, David N.

Virtual Memory Environments for Instruction Scratchpad Memory Management

Jaejin Lee (Seoul Nat. University, ROK)

In this work, we propose virtual memory environments for instruction scratchpad memory (SPM). Our approach consists of two mechanisms.

One is for the embedded systems with no memory management unit (MMU) and the other is for the systems with an MMU. Both are based on demand paging and postpass optimization techniques.

For the systems with no MMU, the virtual memory is implemented purely by software. An SPM manager routine loads required code segments into the SPM on demand at runtime. The code mapping is determined by solving mixed integer linear programming formulation that approximates our demand paging technique. We increase the effectiveness of demand paging by extracting from functions natural loops that are smaller in size and have a higher instruction fetch count. The postpass optimizer analyzes the object files of an application and transforms them into an application binary image that enables demand paging to the SPM.

On the other hand, for the systems with an MMU, we propose a dynamic scratchpad memory allocation strategy targeting a horizontally partitioned memory subsystem. The memory subsystem is equipped with an MMU, and the physically addressed SPM is mapped into the virtual address space. A small instruction minicache is added to further reduce energy consumption and improve performance. Using the MMUs page fault exception mechanism, we track page accesses and copy frequently executed code sections into the SPM before they are executed. Because the minimal transfer unit between the external memory and the SPM is a single memory page, good code placement is of great importance for the success of our method. We discuss code grouping techniques and also analyze the effect of the minicache on execution time and energy consumption.

We evaluate our techniques on several embedded benchmark applications and compare the results to those of a processor core with an instruction cache. The experimental results show that our techniques are quite effective in reducing energy consumption and improving performance, and the SPM can replace the instruction cache.

MOVES - A Tool for Modelling and Verification of Embedded Systems

Jan Madsen (Technical University of Denmark, DK)

One of the major challenges in designing an embedded system is to find a mapping of the application onto the execution platform which effectively fulfills the non-functional requirements of the embedded system such as timing, memory usage, energy consumption, and other cost. A particular challenge is to model and analyse cross-layer dependencies, where the change of a property in one part of the system, e.g. scheduling policy, may impact the performance of another part of the system, e.g. deadline miss on another processor, and hence, the overall system performance.

MOVES is a tool which supports formal analysis of non-functional properties of an embedded system, covering the system layers of an application mapped on

an execution platform, consisting of a heterogeneous multiprocessor architecture where each processor may run a real-time operating system, and where all processors are connected through a network.

Model-Based Embedded System Design

Pieter J. Mosterman (The MathWorks Inc. - Natick, USA)

Model elaboration aspects of Model-Based Design are presented. It is shown how an edge detection filter can be systematically moved through different design phases. To this end, a reference algorithm can be compared to an implementation that is increasingly closer to an implementation. Automatic program synthesis allows the generation of C code or a representation in a hardware description language (HDL). The HDL emulation can then be co-simulated in the system context to study behavior of an implementation at the cycle level. This facilitates analysis of system characteristics with detailed component implementation models and mitigates the need for extensive testbench design.

Keywords: Model-Based Design, image processing, code generation, simulation

Model-Based Design–An overview and research

Pieter J. Mosterman (The MathWorks Inc. - Natick, USA)

Model-Based Design is increasingly adopted in industry. This has resulted in a need for supporting an enterprise-wide use of models. As such, it has to be dealt with computational models that are designed in many different formalisms, that have very different computational semantics, and that capture the same parts at very different levels of detail or in different stages of design. A Computer Automated Multiparadigm Modeling (CAMPaM) framework is outlined that can be used to support and facilitate Model-Based Design. In addition to recent CAMPaM research, hybrid dynamic systems are presented as a mathematical framework for execution. Paradigmatic and pathological behaviors are illustrated by means of a multiple phase space transition ontology.

Keywords: Model-Based Design, Computer Automated Multiparadigm Modeling, hybrid systems, model coverage, testing

Experience in Formal Model Verification for Mission-Critical Domains

Vladimir Okulevich (OOO Siemens CT - St. Petersburg, RUS)

Nowadays the distributed software-intensive systems are widespread in many areas including mission-critical domains such as train, medical, automotive and aviation. Engineering of dependable software to be used in this kind of domains is the difficult task.

Such advanced methods and techniques as Formal Model Verification have to be applied to provide required software quality as it is requested by field standards.

This report contains information about related experience of Dependability Engineering group at Siemens CT Software Engineering. Formal Model Verification approach based on SPIN tool developed in Bell Labs group is shown. This approach was applied for verification of activity arbitrage algorithm between hot-reserved processing units, driver for UNIX-compatible operation system and component interaction in safety-related control system.

The report discusses focus for analysis which could be set on different software artifacts: requirements, design specification, source code, etc. Different fault types identified in carried-out projects are shown and links with used software artifacts are given. Examples of fault types are deadlocks, livelocks, race conditions, safety violations. Benefits for software dependability depending on findings identified are discussed.

Hints are provided and Challenges are identified for further research in this area.

Algorithms for Omega-regular games of Incomplete Information.

Jean-Francois Raskin (Université Libre de Bruxelles, B)

We study observation-based strategies for two-player turn-based games on graphs with omega-regular objectives. An observation-based strategy relies on imperfect information about the history of a play, namely, on the past sequence of observations. Such games occur in the synthesis of a controller that does not see the private state of the plant. Our main results are twofold. First, we give a fixed-point algorithm for computing the set of states from which a player can win with a deterministic observation-based strategy for any omega-regular objective. The fixed point is computed in the lattice of antichains of state sets. This algorithm has the advantages of being directed by the objective and of avoiding an explicit subset construction on the game graph. Second, we give an algorithm for computing the set of states from which a player can win with probability 1 with a randomized observation-based strategy for a Büchi objective. This set is of interest because in the absence of perfect information, randomized strategies are more powerful than deterministic ones. We show that our algorithms are optimal by proving matching lower bounds.

Keywords: Games of imperfect information, controller synthesis.

Full Paper:
http://www.ulb.ac.be/di/ssd/cfv/TechReps/TechRep_CFV_2006_68.pdf

See also: Khrishnendu Charterjee, Laurent Doyen, Thomas A. Henzinger and Jean-Francois Raskin. Algorithms for Omega-regular games of Incomplete Information. In CSL'06, Lecture Notes in Computer Science, 4207, 287-302, 2006.

Timing-Predictability of Cache Replacement Policies

Jan Reineke (Universität des Saarlandes, D)

Hard real-time systems must obey strict timing constraints.

Therefore, one needs to derive guarantees on the worst-case execution times of the systems tasks. In this context, predictable behavior of system components is crucial for the derivation of tight and thus useful bounds.

We present results about the predictability of common cache replacement policies. To this end, we introduce two metrics that capture aspects of cache-state predictability. A thorough analysis of the LRU, FIFO, MRU, and PLRU policies yields the respective values under these metrics. To the best of our knowledge, this work presents the first quantitative, analytical results for the predictability of replacement policies. They support empirical evidence in static cache analysis.

Keywords: Cache analysis, replacement policies, predictability, lru, plru, fifo, mru

Joint work of: Reineke, Jan; Grund, Daniel; Berg, Christoph; Wilhelm, Reinhard

Full Paper:
http://www.avacs.org/Publikationen/Open/avacs_technical_report_009.pdf

See also: Jan Reineke, Daniel Grund, Christoph Berg, and Reinhard Wilhelm. Predictability of cache replacement policies. AVACS Technical Report No. 9, SFB/TR 14 AVACS, September 2006. ISSN: 1860-9821, http://www.avacs.org.

Bottleneck analysis in IEEE 802.11e

Anne Remke (University of Twente, NL)

In a 2-hop IEEE 802.11-based wireless LAN, the distributed coordination function (DCF) tends to equally share the available capacity among the contending stations. Recently alternative capacity sharing strategies have been made possible through the QoS enhancement IEEE 802.11e. We propose a versatile infinite-state Markov reward model to study the bottleneck node in a 2-hop IEEE 802.11-based ad hoc network for different adaptive capacity sharing strategies. We use infinite-state stochastic Petri nets (iSPNs) to specify our model, from which the underlying QBD-type Markov-reward models are automatically derived. The impact of the different capacity sharing strategies is analyzed by CSRL model checking of the underlying infinite-state QBD. Our modeling approach helps in deciding under wich circumstances which adaptive capacity sharing strategy is most appropriate. Furthermore we analyze how the QoS parameters from IEEE 802.11e fit into our model.

Keywords: Model checking, infinite-state Markov chains

Joint work of: Remke, Anne; Haverkort, Boudewijn; Cloth, Lucia; Heijenk, Geert

A Hybrid Approach for Performance Simulation of Distributed Embedded Software

Wolfgang Rosenstiel (Universität Tübingen, D)

For the development of software for networked embedded systems it is important to have a cycle-accurate simulation of these systems. One important problem of this simulation is its long execution time. In this talk it will be shown, how this problem can be solved using a combination of a simulation-based approach with an analytical approach. This hybrid approach combines the advantages of both approaches.

Using this hybrid approach, in a first step the binary code of the program running on the embedded system is analyzed. For the second step there are two possibilities to get a cycle-accurate SystemC program out of the software that will run on the embedded system. Both of them will be shown in the talk. The first possibility is to transform the C source code of the program into a SystemC program and annotate it with cycle information gained from the analysis. The second possibility is to take the binary code of the program and use it to generate annotated SystemC code. The biggest benefit of our approach is, that both generated codes can be annotated by additional dynamic correction code that considers effects of the branch prediction and the caches during run-time and improves the cycle-accuracy.

Keywords: Performance Simulation, Distributed Embedded Software

Applying Model Checking to Real Microcontroller Code

Bastian Schlich (RWTH Aachen, D)

This talk addresses model checking of microcontroller programs. Preliminary to the actual talk a short overview over model checking and microcontrollers is given. The talk starts with a motivation for model checking microcontroller source code. It is described why assembly code instead of C code was chosen for model checking. After that, the model checking tool [mc]square is detailed. The features are itemized, the procedure used is presented, and several details of the architecture are depicted. Next, some features used during model checking are explained in detail (e.g. handling of nondeterminism). Subsequently, two programs are introduced which are used in the following live presentation of the tool. In the end a conclusion is drawn and some potentials for future improvements are described.

Demonstration of Run-time Spatial Mapping of Streaming Applications to a Heterogeneous Multi-Processor System-on-Chip (MPSOC)

Gerard Smit (University of Twente, NL)

In this paper, the problem of spatial mapping is defined. Reasons are presented to show why performing spatial mappings at run-time is both necessary and desirable and criteria for the qualitative comparison of spatial mappings are introduced. An algorithm is described that implements a preliminary spatial mapper. The methods used in the algorithm are demonstrated with an illustrative example.

Joint work of: Hölzenspies, Philip K. F.; Kuper, Jan; Smit, Gerard J. M.; Hurink, Johann

Keywords: Run-time spatial mapping, streaming applications, MPSoC

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2007/1138

Evaluation and Optimization of Reliability of Embedded Systems during Design Space Exploration

Jürgen Teich (Universität Erlangen, D)

Increasing reliability is one of the most important design goals for current and future embedded systems. In this talk, we will put focus on the design phase in which reliability constitutes one of several competing design objectives. Existing approaches considered the simultaneous optimization of reliability with other objectives to be too extensive. Hence, they firstly design a system, secondly analyze the system for reliability and finally exchange critical parts or introduce redundancy in order to satisfy given reliability constraints or optimize reliability. Unfortunately, this may lead to suboptimal designs concerning design objectives.

We will present a) a novel approach that considers reliability with all other design objectives simultaneously, b) an evaluation technique that is able to perform a quantitative analysis in reasonable time even for real-world applications, and c) experimental results showing the effectiveness of our approach.

Interface-based Composition of Embedded Systems

Lothar Thiele (ETH Zürich, CH)

Interface-based design is now considered to be one of the keys to tackling the increasing complexity of modern embedded systems.

The central idea is that different components comprising such systems can be developed independently and a system designer can connect them together only if their interfaces match, without knowing the details of their internals. We use the concept of *rate interfaces* for compositional (correct-by-construction) design of embedded systems whose components communicate through data streams. Using the associated rate interface algebra, two components can be connected together if the *output rate* of one component is "compatible" with the *input rate* of the other component.

We formalize this notion of compatibility and show that such an algebra is non-trivial because it has to accurately model the burstiness in the arrival rates of such data streams and the variability in their processing requirements. We discuss how rate interfaces simplify compositional design and at the same time help in functional and performance verification which would be difficult to address otherwise. Finally, we illustrate these advantages through a realistic case study involving a component-based design of a multiprocessor architecture running a picture-in-picture application.

Joint work of: Thiele, Lothar; Samarjit Chakraborty

Design for Timing Predictability

Reinhard Wilhelm (Universität des Saarlandes, D)

A large part of safety-critical embedded systems has to satisfy hard real-time constraints. These need sound methods and tools to derive reliable run-time guarantees. The guaranteed run times should not only be reliable, but also precise. The achievable precision highly depends on characteristics of the target architecture and system layers of the software. Trends in hardware and software design run contrary to predictability. This article describes threats to timing predictability of systems and proposes design principles that support timing predictability. The ultimate goal is to design performant systems with sharp upper and lower bounds on execution times.

Keywords: Hard real-time, embedded system, run-time guarantees, predictability, timing analysis

Joint work of: Thiele, Lothar; Wilhelm, Reinhard

Full Paper:
http://www.springerlink.com/content/h6650898wp670500/

Abstraction for continuous-time Markov chains

Daniel Willems (RWTH Aachen, D)

A novel abstraction technique for continuous-time Markov chains (CTMCs) will be presented.

The technique fits within the realm of three-valued abstraction methods that have been used successfully for traditional model checking.

The key idea is to apply abstraction on uniform CTMCs that are readily obtained from general CTMCs, and to abstract transition probabilities by intervals. It is shown that this provides a conservative abstraction for both true and false for a three-valued semantics of the branching-time logic CSL (Continuous Stochastic Logic). Experiments on an infinite-state CTMC indicate the feasibility of this abstraction technique.

Implementation of High Performance Flash Memory Based Storage

Jin Hyuk Yoon (Seoul Nat. University, ROK)

In this talk, I present basic materials on flash memories focusing on NAND type flash memory which is mainly used as storage media in various storage devices. For performance evaluation, low level and high level performance metrics for general storage devices are described. Then, a brief sketch of the project recently done by our group is provided with performance results. In conclusion, I summarize characteristics of flash memory based storage device which now starts gaining interests in non-volatile storage hierarchy.

Predicate abstraction for probabilistic models

Lijun Zhang (Universität des Saarlandes, D)

Probabilistic models are widely used to analyze embedded, networked, and more recently biological systems. Existing analysis techniques are limited to finite-state models and suffer from the state explosion problem. We propose predicate abstraction for probabilistic models to tackle the state explosion problem and admit infinite-state models. We have implemented PASS, a predicate abstraction model checker for probabilistic models supporting unbounded integer and real arithmetic, as well as bit vectors.

We demonstrate the feasibility of our approach by verifying properties of the Bounded Retransmission Protocol for any file length and any number of retransmissions.

To the best of our knowledge, this is the first time that properties of probabilistic infinite-state models have been verified at this level of automation.

3 Working group reports

Performance Measures Other Than Time

This presentation shows a few possible performance measures that might be interesting and possible evaluation methods.

Joint work of: Cloth, Lucia; Crouzen, Pepijn; Fruth, Matthias; Han, Tingting; Jansen, David N.; Kattenbelt, Mark; Smit, Gerard; Zhang, Lijun

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2007/1139

Working Group: Stochastic vs. Worst-Case Methods

Lothar Thiele (ETH Zürich, CH)

The slides present the results of the working group.