IST-214373 ArtistDesign
Network of Excellence
on Design for Embedded Systems

Activity Progress Report for Year 4

Validation

Cluster:

Modeling and Validation

Activity Leader:

Professor Kim G. Larsen (CISS, Aalborg University)

http://www.cs.aau.dk/~kgl

Policy Objective (abstract)
The objective is to address the growth in complexity of future embedded products while reducing time and cost to market requires methods allowing for early exploration and assessment of alternative design solutions as well as efficient methods for verifying final implementations. This calls for a range of model-based validation techniques ranging from simulation, testing, model-checking, compositional techniques, refinement as well as abstract interpretation. The challenge will be in designing scalable techniques allowing for efficient and accurate analysis of performance and dependability issues with respect to the various types of (quantitative) models considered. The activity brings together the leading teams in Europe in the area of model-based validation.

Versions

number	comment	date
1.0	First version delivered to the reviewers	Feb 15, 2012

Table of Contents

1. Overview of the Activity (2008-2011)

1.1 ArtistDesign participants and their role within the Activity

Dr. Jan Tretmans (ESI - Netherlands);
Testing, performance analysis, predictability..

Prof. Werner Damm (OFFIS - Germany);
formal analysis techniques, mainly on compositional techniques regarding safety and real, and deployment synthesis.

Prof. Tom Henzinger (IST, Austria);
Rich interface theory for component-based design. Algorithms for checking quantitative reliability measures of implementations. Compositional code generation for time-triggered architectures. Algorithms for stochastic and timed games.

Prof. Thierry Jéron, Bertrand Jeannet (INRIA - France);
Models with data and time for model-based test selection and coverage criteri. qualitative *and quantitative verification, control synthesis.*

Prof. Christoph Kirsch (Salzburg - Austria)*;*
Compositional Compositional timing and reliability validation in Giotto-inspired languages and systems

Prof. Kim Larsen (CISS, Aalborg - Denmark);
Quantitative verification, synthesis, performance evaluation and model-based testing for timed automata and games with priced and probabilitistic extensions.

Alberto Sangiovanni-Vincentelli, University of Trento, Italy.
Platform-Based Design, the Metropolis and COSI frameworks, distributed sense and control systems, industrial applications and international activities.

Roberto Passerone, University of Trento (Italy)
Formal analysis of heterogeneous composition, abstract algebra, and meta-modelling..

Prof. Joseph Sifakis – VERIMAG (France)
Contributions of his team: component-based design, the BIP framework, platform-aware implementation of embedded systems, structural verification

Prof. Saddek Bensalem – VERIMAG (France)
Contributions of her team: structural analysis.

Prof. Oded Maler – VERIMAG (France)
Contribution *of his team: timing analysis, scheduling and hybrid systems*

Prof. Martin Törngren, Prof. Axel Jantsch, KTH, Stockholm, Sweden
Integrated models supporting cross-layer validation. Methods for validation of self-configuring systems.Compositional validation of integrated models/components..

Prof., Wang Yi (Uppsala - Sweden);
Scheduling and Verification (UPPAAL and TIMES), Combination of State-Based and Analytical Analysis Techniques (CATS tool)

Prof. Christophe Gaston
compositional validation, CEA symbolic execution of models of heterogeneous systems as a basis for testing or model checking activities, also taking compositionality into account.

.

-- Changes wrt Y3 deliverable --

No changes

1.2 Affiliated participants and their role within the Activity

Prof. Yiannis Papadopolis, Univ. Of Hull (UK)
Compositional safety analysis and design optimization w.r.t. safety.

Prof. Ahmed Bouajjani - LIAFA (France)
Real-time and hybrid model checking

Stavros Tripakis – University of California at Berkeley (USA)
Monitoring and test of real-time properties

Prof. Pierre Wolper and Prof. Jean-Francois Raskin (CVF – Belgium);
Efficient Model-checking of linear-time properties.
Verification and synthesis for reactive systems. Timed and hybrid automata.

Joost-Pieter Katoen (Aachen – Germany)
Model checking of quantitative system properties. Verification of (continuous-time) probabilistic and stochastic systems.

Prof. Dr. Holger Hermanns (Saarland U – Germany);
Probabilistic and stochastic model checking.

Prof. Christel Baier (Dresden – Germany);
Probabilistic and stochastic model checking

Dr. Patricia Bouyer, Dr. Nicola Markey and Dr. Phillippe Schnoebelen (LSV Cachan – France),
Decidability and algorithms for priced timed automata and games.
Algorithms for solving games of imperfect information

Prof. Roderick Bloem (TU Graz)
Algorithms for controller synthesis

Prof. dr. ir. Wil van der Aalst, professor at Eindhoven University of Technology, The Netherlands.
Information System. Affiliated participant in the ESI Octopus project.

Prof. dr. Mehmet Aksit, professor at Twente University, The Netherlands.
Software Engineering. Affiliated participant in the ESI Darwin project.

Prof. dr. Sandro Etalle, professor at Eindhoven University of Technology, The Netherlands.
Security. Affiliated participant in the ESI Darwin project.

Prof. dr. Arjen van Gemund, professor at Delft University of Technology, The Netherlands. Embedded Software Laboratory.
Affiliated participant in the ESI projects Trader and Octopus.

Prof. dr. Frits Vaandrager, professor at Radboud University, The Netherlands.
Formal methods. Affiliated participant in the ESI Octopus project.

Prof. dr. Hans van Vliet, professor at Vrije Universiteit Amsterdam, Software Engineering.
Affiliated participant in the ESI Darwin project.

Prof. dr. Jack van Wijk. professor at Eindhoven University of Technology, The Netherlands.
Visualization. Affiliated participant in the ESI Poseidon project.

Prof. Peter Habermehl – LIAFA (France)
verification of programs with arrays and dynamic data structures

-- Changes wrt Y3 deliverable --

No changes.

1.3 Starting Date, and Expected Ending Date

Starting date: January 1st 2008

Expected ending date: the activity is intended to continue beyond the end of the project (December 2011). The needs for new techniques (algorithms and data structures) for verifying and analysing system models that incorporate both functional and quantitative aspects (such as safety requirements, timing, resource constraints, reliability, etc.) are expected to continue increase in the next decade. Moreover, the feedback from the concrete applications should give to this activity new directions to investigate for researchers, most likely beyond the duration of the project.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.4 Policy Objective

The objective is to address the growth in complexity of future embedded products while reducing time and cost to market requires methods allowing for early exploration and assessment of alternative design solutions as well as efficient methods for verifying final implementations. This calls for a range of model-based validation techniques ranging from simulation, testing, model-checking, compositional techniques, refinement as well as abstract

interpretation. The challenge will be in designing scalable techniques allowing for efficient and accurate analysis of performance and dependability issues with respect to the various types of (quantitative) models considered. The activity brings together the leading teams in Europe in the area of model-based validation.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.5 Background

By far the most common validation technique applied in embedded industrial today is based on rather ad-hoc and manual (hence quite error-prone) testing. Given that some 30-50% of the overall development time and cost are related to testing activities it is clear that the impact of improved validation technologies is substantial. Given this current industrial practice the academic state-of-the-art has a lot to offer. In particular the cluster combines the efforts and skills on of the individual leading researchers in Europe into a world-class virtual team for advancing the state-of-the-art and industrial take-up of model-based validation techniques.

Whereas validation techniques for assessing functional correctness have reached a certain level of maturity and industrial acceptance, there is a need for mature validation techniques addressing quantitative aspects (e.g. real-time, stochastic and hybrid phenomena) being accessible from within industrial tool-chains. Thus, particular effort should be made to transfer of validation methods and tools to industry, including integration of the techniques developed into existing tools.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

1.6 Technical Description: Joint Research

The joint research falls into the following three sub-activities:

A Compositional validation:
The complexity of a given analysis method is not only determined by its accuracy (and issues addressed) but mainly by the sheer size of the model analysed measure in number of components, tasks, variables, etc. In order to achieve methods which scale to the need of industry *compositionality* is paramount. That is, it should be possible for composite models to be interrelated and properties to be inferred only by consideration of the components of the models and their interfaces. In the presence of composite models with heterogeneous components – in particular involving components where quantitative aspects are considered – this is a challenge that has not yet been dealt with satisfactory.

B Quantitative validation:
Whereas functional validation addresses issues concerning logical correctness with respect to stated temporal specifications, quantitative validation takes the quantitative aspects into account. For embedded systems applied in safety-critical applications hard real-time guarantees are often imperative. For embedded systems in less critical applications

performance and QoS are often more important properties: in this case the quantitative validation should return a value as to the "quality" of the model with respect to a given relevant metric, e.g. expected energy consumption pr time-unit. The quantitative aspects to be dealt with involve real-time, stochastic and hybrid phenomena. Also joint work on software verification, and more particularly on modelling and verification of quantitative properties of programs using integer arrays has been made, as well as work joint work on the evaluation of performance properties by connecting the DOL performance analysis and BIP

C Cross-layer validation
During the design trajectory, the software engineer will create, refine and make use of several models of the same system focusing on different aspects and varying in terms of particular to transfer properties established of one (early) model to properties guaranteed to hold of other (later) models without any additional effort.

Techniques for validating the conformance between design models and executing code (on particular platforms) are particular important. This includes considerations of (robust) methods for automatic code generation as well as methods for synthesizing controllers from plant models and control objectives.

In order for validation methods to be industrial applicable it is essential that existing (or thirdparty) code may be dealt with. Here software verification techniques (combining static analysis and model checking) need to be extended to involve quantitative aspects.

-- Changes wrt Y3 deliverable --

No changes with respect to Year 3.

2. Work Achieved in the NoE

2.1 Synthesis View of the Main Overall Achievements

We maintain the division of the validation activities into the three sub-activities:

- **A.** "Compositional Validation" where we focus on exploiting the compositional structure of designs in their verification.
- **B.** "Quantitative Validation" where we provide methods, algorithms and tools for analyzing non-functional properties, e.g. energy consumption.
- **C.** ""Cross Layer Validation" where we provide methods for verifying conformance system descriptions appearing at different stages of the design process(requirements, design models and code).

The research within the *Validation Activity* have progressed substantially within the four years of the project oin terms of methods and validation techniques developed, and with significant synergy with proposed modeling formalisms proposed of the *Modeling Activity* .

Within the sub-activity *Compositional Validation* the main focus has been on methods for deriving non-functional properties from properties of their components, with the purpose of developing scalable compositional techniques for performance analysis and verification. Also validation methods based on abstractions and refinements for quantitative models have been developed.

Within sub-activity *Quantitative Validation.* the focus was on design frameworks for quantitative modeling, in particular Markov models, timed automata, priced timed automata, memory models involving stacks and queue and linear hybrid. A main achievement has been the wealth of algorithmic techniques allowing for efficient and scalable validation of formalism whose expressive power was previously out of reach. A particular scalable technique which has emerged is that of statistical model checking which allows several performance properties of very rich models to be established on the basis of simulation *up to a desired level of confidence.*

Within the sub-activity *Cross-Layer Validation* a substantial line of results have been obtained with respect to improved schedulability analysis and WCET analysis supporting multiprocessor and multi-core applications. The methods include WCET analysis and schedulability analysis addressing mixed-criticallity systems including tool implementation using model checking, as well as introduction new task models (e.g. Digraph based) allowing for more scalable and efficient schedulability analysis. Main results within *Cross-Layer Validation* concerns automatic controller synthesis from various rich game models (timed and probabilistic) with possible partial obsevability, and with a number of industrial successful application already having been achieved (e.g. the automatic synthesis of climate control in pig-stable, and synthesis of optimal control of hydraulic pumps). This shows that the distance from fundamental theoretical breakthroughs to industrial impact may be very short. Also, a number of results have been obtained with respect to conformance testing of non-functional properties based on quantitative model. Finally, within the theory of timed automata substantial effort has been made towards the analysis of their robustness: i.e. to what extent does the realization of the model on a non-perfect platform preserve properties already established.

-- The above is new material, not present in the Y3 deliverable --

2.2 Work achieved in Year 1 *(Jan-Dec 2008)*

The following provide a high-level description of the work achieved in Year 1:

Within the sub–activity A "Compositional Validation", we focused on methods for deriving functional as well as non-functional properties of composite systems from properties of their components. In particular compositional approaches dealing with timing properties as well as safety, failure and reliability was addressed. Also, validation methods based on abstractions and refinements were developed.

Within the sub-activity B "Quantitative Validation", we provided (un)decidability results as well as efficient datastructures and algorithms supporting the validation of a number of non-functional models (e.g. Markov chains, timed automata, priced timed automata, memory models involving stacks and queues, linear hybrid systems) as well as their interrelation.

Within the sub-acitivity C "Cross-layer Validation", main effort was made towards controller synthesis from rich game models as well as conformance testing of non-functional propeties.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in section 1.7.

2.3 Work achieved in Year 2 *(Jan-Dec 2009)*

Within the sub–activity A "Compositional Validation", we have worked on combining state-based and analytical methods to develop scalable compositional techniques for performance analysis and verification of timed systems. Also a number of compositional development and verification frameworks for timed and stochastic systems have been put forward allowing to infer in a compositional manner that programs exhibit predictable behaviour. Development of symbolic execution of heterogeneous systems and a symbolic execution framework devoted to system models defined recursively by interconnecting heterogeneous component models has been made. Finally work has continued its work on deadlock detection/verification and its implementation in the D-Finder tool by checking incrementally deadlock-freedom of component-based systems described as the composition of interacting components is proposed.

Within the sub-activity B "Quantitative Validation", a substantial amount of work from different partners has been made on schedulability and execution time analysis for multiprocessor platforms with pipelines and shared caches. New tools supporting verification of quantitative models combining both timing and stochastic properties have been developed. We have applied three-valued abstraction techniques for probabilistic systems showing that certain abstractions provide rather tight bounds. We have developed methods for verification of programs with arrays and dynamic data structures, investigated improved widening techniques for the abstract interpretation of numerical programs with polyhedra with the purpose of analysing Linear Hybrid Systems, and developed extendable tools for verification of hybrid systems.

Within the sub-acitivity C "Cross-layer Validation", we have continued the effort on controller synthesis from rich game models and from models with partial observability. Work conformance testing of non-functional properties has also been continued. New effort has been made on model learnability from experimentation. Also tools for establishing refinement between specification at different abstraction levels have been developed. Work on translations from real-time temporal logics to deterministic timed automata in the context of synthesis of real-time controllers as well as work on verifying real-time models with respect to scenario-based specifications constitutes contributions to cross-layer validation.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in section 1.8.

2.4 Work achieved in Year 3 *(Jan-Dec 2010)*

We maintain the division of the validation activities in the following three subactivities:

A. Compositional Validation: aiming at developing validation techniques for establishing properties of composite heterogeneous systems from properties of its components.

B. Quantitative Validation: aiming at developing validation techniques for quantitative system properties including time, resource, hybrid as well as stochastic properties.

C. Cross-layer Validation: aiming at developing methods validating the conformance between designs at different levels of abstraction as well as conformance of executable code and designs.

Within the sub-activity A “Compositional Validation” several partners have worked substantially and collaboratively on compositional design and verification methodologies for functional, timing and stochastic aspect. This methods span assume/guarantee reasoning, interface automata as well as modal transition systems for rich models. Also, theoretical foundations and coordination languages has been developed for heterogeneous systems. Finally, a framework for tool integration based on meta-models and model-transformations has been established.

Within sub-activity B “Quantitative Validation” substantial work has been made on improved schedulabilty analyses supporting multiprocessor and multi-core applications. The improved methods include taking scheduler overhead into account, being power-aware – i.e. exploiting slacks in the system of processes to reduce power consumption while insuring deadlines are met. This work includes combination of abstract interpretation and model-checking for timing and interference analysis of parallel programs on multi-core, and schedulability analysis of Safety Critical Java applications on FPGAs,. Symbolic validation methods for timed automata based formalisms extended with cost/energy information as well as stochastic interpretations has been developed. Also, the work within this sub-activity includes exploitation and implementation of statistical model-checking techniques within the BIP tool.

Within the subactivity C “Cross-Layer Validation” substantial work has been made on improved methods for model-based testing. This work includes incremental testing of composite systems, off-line test generation from timed automata models, model-based test generation for data-intensive systems, as well as runtime monitoring. Work on controller synthesis has continued, focusing on synthesis techniques for timed games, modal specifications, and scenario-based specifications with the purpose of addressing timing properties, multi-objective optimization as well as fault-tolerance. Validation related to executable implementation includes run-time programming, optimal implementation of communication for time-constrained synchronous reactive modules, as well as modular WCET analysis of C-code executing on ARM-processors.

-- No changes wrt Y3 deliverable --

This section was already presented in the Y3 deliverable, in sections 1.9.

2.5 Work achieved in Year 4 *(Jan-Dec 2011)*

Within the sub-activity *Compositional Validation* several partners have worked substantially and collaboratively on compositional design and verification methodologies for functional, timing and other non-functional aspects. These methods span assume/guarantee reasoning, interface automata as well as modal transition systems for rich models. In particular, composition frameworks have been proposed, as well as frameworks addressing design for integratability, maintainability, as well as methods for component adaptation (e. g. in the case of protocol mismatches). Also, theoretical foundations and coordination languages for heterogeneous systems have been further developed. Moreover, frameworks for tool integration based on meta-models and model-transformations have been consolidated and applied to case studies

Within sub-activity *Quantitative Validation* work from previous years on improved schedulability analysis and WCET analysis supporting multiprocessor and multi-core applications has been made. The methods include WCET analysis and schedulability analysis addressing mixed-criticallity systems as well as introduction new task models (e.g. Digraph based) allowing for more scalable and efficient schedulability analysis. Substantial work has been made with timed automata as based, including frequency analysis and off-line test slection, analysis of parametric quantitative models, analysis of resource consumption using energy- and price-extensions of timed auotmata, as well as highly scalable statistical model checking of performance properties of timed automata models. Finally, notions of metrics (providing notions of approximate correctness) and robustness for timed automata models have been substantiated and refined.

Within the sub-activity *Cross-Layer Validation* substantial work has been made on further improved methods for model-based testing. This work on conformance testing of real-time systems using time- and data abstraction, asynchronous testing and test-case generation for embedded Simulink, includes incremental testing of composite systems, off-line test generation from timed automata models, model-based test generation for data-intensive systems, as well as runtime monitoring. Closely related to that of testing is work on the learning of (probabilistic) automata.

In terms of contributions of the partners of the activity the following work has been achieved:

INRIA has continued to develop foundations for quantitative verification, model-based testing, controller synthesis and monitoring in the context of timed systems, infinite states systems and probabilistic systems.

UPPSALA Timing analysis of Non-LRU caches. Non LRU caches are widely used in practice. However there has been little work on the analysis of non LRU caches because they are in general considered as non predictable. A recent work by the Uppsala team reveals that the MRU replacement strategy used in the Nehalem architecture is as predictable as LRU, and thus it may change the view that non-LRU caches are non analyzable.

UPPSALA Multicore scheduling. The first objective is to evaluate the existing multiprocessor scheduling algorithms and implementation overheads in operating systems. The second is to further improve the precision of schedulability test based on utilization bounds.

UPPSALA Mixed-criticality systems. The goal is to deploy and integrate different types of applications e.g. hard and soft real-time applications on the same hardware platforms. The applications may have different levels of criticality and thus different requirements of computation resources.¨

UPPSALA Expressiveness vs. analysis efficiency of models for timed systems. The objective is to find the optimal trade-off to develop models that are expressive enough for modeling of realistic systems, and also tractable in the sense they can be efficiently analyzed automatically. We also establish hardness results on the analysis of scheduling problems.

UPPSALA The Matlab/Simulink language has become the standard formalism for modeling and implementing control software in areas like avionics, automotive, railway, and process automation. Such software is often safety critical, and bugs have potentially disastrous consequences for people and material involved. We define a verification methodology to assess the correctness of Simulink programs by means of automated test-case generation. In the style of fault- and mutation-based testing, the coverage of a Simulink program by a test suite is defined in terms of the detection of injected faults.

UPPSALA Craig interpolation has become a versatile tool in formal verification, for instance to generate intermediate assertions for safety analysis of programs. We investigated the decision and interpolation problem for a number of logics and theories relevant for verification, including Presburger arithmetic and arrays, and presented new calculi, algorithms, and implementations.

UPPSALA Modern multicore processors, such as the Cell Broadband Engine,½achieve high performance by equipping accelerator cores with small½"scratch-pad" memories. The price for increased performance is higher programming complexity -- the programmer must manually orchestrate½data movement using direct memory access (DMA) operations.½Programming using asynchronous DMA operations is error-prone, and DMA½races can lead to nondeterministic bugs which are hard to reproduce and fix. We present a method for DMA race analysis in C programs.

KTH, Volvo and other affiliated partners worked on extending the support for safety modeling as part of the EAST-ADL architecture description language. Closely related an investigation in modeling and robustness analysis through model-level fault-injection in behavior models was concluded with a thesis (Technical Achievement 30).

OFFIS, KTH and Volvo developed a common meta-modeling approach supporting interoperability of models and tools.

Salzburg has been working on analytical methods for isolating concurrent processes in terms of their power consumption. The main challenges in providing power isolation are to find a relationship between the power consumption of a system and the contribution of a single process to this power consumption as well as understanding the trade-off between quality and cost of power isolation. Power isolation may enable per-process cost accounting based on power consumption.

CEA LIST has defined testing techniques to address incremental testing of component oriented systems. Two approaches where addressed. One is based on specifications of intended interactions between components given in the form of sequence diagrams with real time constraints. The second consists in generating test cases for components of orchestrated systems, from the specification of orchestrator components.

TRENTO (in collaboration with VERIMAG) has worked on the integration of BIP functional models into the MetroII environment, in order to study the performance of the system when mapped onto different architecture platforms. The tool integration enriches the capabilities of the two tools and preserve the structure and the semantics of the original model. The method has been tested on a distributed sorting algorithm case study.

TRENTO (in collaboration with ETHZ) has worked on the parametric analysis and validation of embedded systems specified in real-time calculus. For this, the partners have developed an integration flow between Modular Performance Analysis (MPA) developed at ETHZ and the Parametric Schedulability Analysis (PSA) developed at TRENTO. The tool allows the feasibility of a system to be studied in a range of parameters. The method has been tested on simple

cases, and on a more advanced system that uses state based components which are difficult to model in MPA.

CISS has worked on specification formalisms for consumption of resource, e.g. energy. This includes full study and presentation of priced timed automata, introduction of energy automata (allowing for positive and negative weigh rates), weighted extensions of modal transition systems, as well as studying cost-constrained (by interval) behavior of multi-weighted transition systems.

CISS (in collaboration with INRIA) has worked on compositional specification formalisms for probabilistic systems, ranging from Interval Markov Chains, Constraint Markov Chains and Abstract Probabilistic Automata, and has worked on Logicl and compositional reasoning for continuous Markovian systems.

CISS has worked statistical model checking for timed automata, and even more generally energy timed automata. This line of research provides a "natural" stochastic interpretation of timed automata (in terms of distributions of delays), as well as efficient implementation within the UPPAAL tool suite. The implementation also contains a distributed implementation of sequential testing algorithms.

CISS has worked on metrics on quantitative behaviours resulting in several papers and in several directions. In particular, this line of research contains novel contributions to the understanding of robustness for timed automata.

CISS has worked on introduction of and complexity results for parametric modal transition systems and modal transitions with data, which provides a useful extension of the well-established specification theory of modal transition systems.

CISS has provided methods for value-set analysis of low-level code, providing vital information for improved WCET estimations.

CISS has worked on methods for learning probabilistic automata.

CISS has worked on development around UPPAAL, e.g. the prototype tool OPAAL (in Phyton) allowing for rapid prototyping of new algorithmic and datastructuring ideas, and TAPAAL a real-time verification tool for Timed-Arc Petri Nets using part of the UPPAAL model checking engine.

-- The above is new material, not present in the Y3 deliverable --

3. Detailed view of the progress in Year 4 *(Jan-Dec 2011)*

3.1 Technical Achievements

Sub-activity A: Compositional Validation

Adding precise semantics to the EAST-ADL2 architecture description language to support formal analysis (**KTH + Volvo**)
KTH, in cooperation with Volvo has developed a behavior extension to the EAST-ADL2 language. This extension enhances the behavior modeling capability of EAST-ADL2, so that the model is precise and susceptible to the SPIN model checker. An algorithm was provided to convert (transform) an EAST-ADL2 behavior model to a SPIN model.
Further work in this direction has assessed the need for such behavior modeling capabilities in the context of future electrical vehicles. Future fully electrical vehicles (FEV) are safety critical and have particular complexity in operations, design and integration of control functionalities, power efficiency and harness. To develop appropriate language support, there is a need to clarify the specific engineering and quality concerns and thereby to specify necessary language support in regard to both methodology and modeling artifacts. In particular, according to the derived requirements, the expected modeling support for behavior description ranges from the definitions of operation and boundary conditions, to the specifications of functional constraints, mode and error behaviors, and to the physical dynamics of plants, harness and power. While providing an enhanced support for requirements engineering, functional safety, architectural composition, refinement, and contract specification, the proposed native behavior extension of EAST-ADL would also constitute the basis for integrating external methods and tools for early estimations and V&V of performance and dependability [MD11]. The proposed language extension consists of three categories of behavior constraints:

- Attribute Quantification Constraint – relating to the declarations of value attributes and the related acausal quantifications (e.g., U=I*R).
- Temporal Constraint – relating to the declarations of behavior constraints where the history of behaviors on a timeline is taken into consideration.
- Computation Constraint – relating to the declarations of cause-effect dependencies of data in terms of logical transformations (for data assignments) and logical paths.

Each of these behavior constraints can be associated to time conditions given in terms of logical time, of which the exact semantics can be given by the existing EAST-ADL support for execution events (e.g. the triggering, and port data sending & receiving events of a function). The meta-model integration is done in a modular way such that no existing EAST-ADL constructs are modified by the extension.
Model transformations for the purpose of EAST-ADL analysis using Uppaal have also been investigated, first by assessing different mappings and then by transformation experiments, [NC11a]. Mappings from EAST-ADL concepts to Autosar have also been investigated. Three case studies, of a position control, fuel control and a brake-by-wire system, have been used to support and validate the work. The resulting mapping scheme provides a basis for automated architecture refinements and synthesis, [NC11b].

Behavioural constraqints on Components (CEA LIST) In early design phases, system models can be characterized as intended interactions between black box components. Moreover, when dealing with embedded systems, it is usual that interactions are constrained by timing issues. CEA ([BGS11]) has proposed to represent such system models as structured scenarios by using UML sequence diagrams specialized with the MARTE profile to handle timing constraints. By using symbolic execution techniques, CEA has shown how to analyze

these system models and how to extract behavioral constraints concerning components. Those constraints can be used as unitary test purposes to select components of the system.

Orchastration of Components (CEA LIST) Orchestrations are distributed systems deployed on a network where there is a central component (called orchestrator) coordinating other components (called services). Services are off-the-shelves components pre-existing to the orchestration design phase. CEA ([EGL11]) has proposed an approach to select candidate services (to be chosen to implement an orchestration) by taking into account the intended behaviors of those services as they can be inferred from the orchestrator specifications. Services are tested with respect to those behaviors to decide whether or not they can be selected. Specifications of orchestrators are Timed Input/Output Symbolic Transition Systems. Service intended behaviors are elicited by means of symbolic execution and projection techniques. Those behaviors can be used as test purposes for a timed symbolic conformance testing algorithm.

Performance Analysis of BIP functional models using the MetroII environment (TRENTO, VERIMAG). The BIP and MetroII frameworks provide substantial complementary features, the first more oriented towards formal analysis, while the second more towards performance estimation. For this reason, we have developed a modeling and validation flow that integrates the two environments to provide substantial benefits to the designer. A BIP model is imported into MetroII by taking advantage of the distributed code generation available in BIP. Once the BIP model is imported in MetroII, it is relatively easy to explore the performance of the system under different mappings. The architecture model is developed using the MetroII infrastructure, which includes ways of estimating performance and characterizing platform components. The separation of the the functional and architectural models is essential to take advantage of strong analysis techniques, which exploit the particular characteristics of the models. From the point of view of the BIP model, the architecture model and the MetroII semantic can be seen as additional constraints on the execution of the BIP model. Thus the traces of the mapped model are a subset of all the possible traces of the original BIP model. Therefore the good properties in the BIP model are retained in the mapped model. For example, if the architecture can make sure that all the processes in the BIP model eventually have a chance to run, the mapped model will remain deadlock-free as long as the original BIP model is deadlock-free. The semantic preservation is achieved by reproducing in MetroII the BIP interactions through appropriate communication channels, derived from the original functional specification. These model transformation rules have been implemented as an automatic BIP to MetroII ANTLR-based converter which takes as input the BIP functional model and produces a corresponding, structurally equivalent, MetroII functional model. The results have been submitted to publication.

Compositional specification formalisms for probabilistic systems (CISS with INRIA)
Interval Markov Chains (IMC) are the base of a classic probabilistic specification theory by Larsen and Jonsson in 1991. They are also a popular abstraction for probabilistic systems. In [DLLPW11] we study complexity of several problems for this abstraction, that stem from compositional modeling methodologies. In particular we close the complexity gap for thorough refinement of two IMCs and for deciding the existence of a common implementation for an unbounded number of IMCs, showing that these problems are EXPTIME-complete. We also prove that deciding consistency of an IMC is polynomial and discuss suitable notions of determinism for such specifications.

Notions of specification, implementation, satisfaction, and refinement, together with operators supporting stepwise design, constitute a specification theory. In [CDLLPW11] We construct such a theory for Markov Chains (MCs) employing a new abstraction of a Constraint MC. A Constraint MCs permit rich constraints on probability distributions and thus generalize prior

abstractions such as Interval MCs. Linear (polynomial) constraints suffice for closure under conjunction (respectively parallel composition). This is the first specification theory for MCs with such closure properties. We discuss its relation to simpler operators for known languages such as probabilistic process algebra. Despite the generality, all operators and relations are computable.

Probabilistic Automata (PAs) are a widely-recognized mathematical framework for the specification and analysis of systems with non-deterministic and stochastic behaviors. This paper [KSDLLPW11] proposes Abstract Probabilistic Automata (APAs), that is a novel abstraction model for PAs. In APAs uncertainty of the non-deterministic choices is modeled by may/must modalities on transitions while uncertainty of the stochastic behaviour is expressed by (underspecified) stochastic constraints. We have developed a complete abstraction theory for PAs, and also propose the first specification theory for them. Our theory supports both satisfaction and refinement operators, together with classical stepwise design operators. In addition, we study the link between specification theories and abstraction in avoiding the state-space explosion problem. In [KSDLLPW11b we discuss APAs over dissimilar alphabets, a determinisation operator, conjunction of non-deterministic APAs, and an APA-embedding of Interface Automata. We conclude introducing a tool for automatic manipulation of APAs.
[DLLPW11] provides and reports on the implementation of the approach in the Abstract Probabilistic Automata Checker toolset, exploiting the Z3 SMT solver of Microsoft.

Metrics and Robustness (CISS) In [FTL11] a general framework is developed for reasoning about distances between transition systems with quantitative information. Taking as starting point an arbitrary distance on system traces, we show how this leads to natural definitions of a linear and a branching distance on states of such a transition system. We show that our framework generalizes and unifies a large variety of previously considered system distances, and we develop some general properties of our distances. We also show that if the trace distance admits a recursive characterization, then the corresponding branching distance can be obtained as a least fixed point to a similar recursive characterization. The central tool in our work is a theory of infinite path-building games with quantitative objectives.

Simulation distances are essentially an approximation of simulation which provide a measure of the extent by which behaviors in systems are inequivalent. In [LFT11], we consider the general quantitative model of weighted transition systems, where transitions are labeled with elements of a finite metric space. We study the socalled point-wise and accumulating simulation distances which provide extensions to the well-know Boolean notion of simulation on labeled transition systems. We introduce weighted process algebras for finite and regular behavior and offer sound and (approximate) complete inference systems for the proposed simulation distances. We also settle the algorithmic complexity of computing the simulation distances.

[FLeT11] present a distance-agnostic approach to quantitative verification. Taking as input an unspecified distance on system traces, or executions, we develop a game-based framework which allows us to define a spectrum of different interesting system distances corresponding to the given trace distance. Thus we extend the classic linear-time–branching-time spectrum to a quantitative setti ng, parametrized by trace distance. We also prove a general transfer principle which allows us to transfer counterexamples from the qualitative to the quantitative setting, showing that all system distances are mutually topologically inequivalent.

Timed automata follow a mathematical semantics, which assumes perfect precision and synchrony of clocks. Since this hypothesis does not hold in digital systems, properties proven formally on a timed automaton may be lost at implementation. In order to ensure implementability, several approaches have been considered, corresponding to different

hypotheses on the implementation pla□orm. In [BLMST11] we address two of these: A timed automaton is samplable if its semantics is preserved under a discretization of time; it is robust if its semantics is preserved when all timing constraints are relaxed by some small positive parameter. We propose a construction which makes timed automata implementable in the above sense: From any timed automaton A, we build a timed automaton A' that exhibits the same behaviour as A, and moreover A' is both robust and samplable by construction.

Specification theories for real-time systems allow to reason about interfaces and their implementation models, using a set of operators that includes satisfaction, refinement, logical and parallel composition. To make such theories applicable throughout the entire design process from an abstract specification to an implementation, we need to be able to reason about possibility to effectively implement the theoretical specifications on physical systems. In the literature, this implementation problem has already been linked to the robustness problem for Timed Automata, where small perturbations in the timings of the models are introduced. [LLTW11] addresses the problem of robust implementations in timed specification theories. Our contribu tions include the analysis of robust timed games and the study of robustness with respect to the operators of the theory.

Sub-activity B: Quantitative Validation

Frequency analysis for timed automata (INRIA) In [BBBS11] we propose a natural quantitative semantics for timed automata based on the so-called frequency, which measures the proportion of time spent in the accepting states. We study various properties of timed languages accepted with positive frequency, and in particular the emptiness and universality problems.

Off-line Test selection for non-deterministic timed automata (INRIA) In [BJSK11], we propose novel off-line test generation techniques for non-deterministic timed automata with inputs and outputs (TAIOs) in the tioco conformance theory. The underlying determinization problem is solved in [BSJK11] thanks to an approximate determinization using a game approach. We adapt this procedure here to over- and under-approximation, in order to guarantee the soundness of generated test cases. Test selection guided by test purposes is performed by a symbolic co-reachabiity analysis

Runtime Enforcement Monitoring (INRIA) [FFM11] and [FMFR11] present a unified view of runtime verification and enforcement of properties in the Safety-Progress classification. The properties which can be verified (monitorable properties) and enforced (enforceable properties) at runtime are characterized. Furthermore, we propose a systematic technique to produce a monitor from the automaton recognizing a given safety, guarantee, obligation or response property. Finally, we show that this notion of enforcement monitors is more amenable to implementation and encompasses previous runtime enforcement mechanisms.

Controllers for probabilistic systems (INRIA) Partially Observable Markov Decision Processes (POMDP for short) are a formalism to model systems in which users do not have a full access to the information. In [BG11] we tackle the problem of the minimal information a user needs at runtime to reach an objective with probability one, where, at each step the user can either choose to use the partial information, or pay a fixed cost and receive the full

information. This optimization question gives rise to two different problems, whether we consider to minimize the worst case cost, or the average cost. We show that efficient techniques from the model checking community can be adapted to compute the optimal worst case cost and give optimal strategies for the users.
On the other hand, we show that the optimal average price cannot be computed in general, nor can it be approximated in polynomial time even up to a large approximation factor.

Symbolic Supervisory Control of Infinite Transition Systems (INRIA) We address control problems for infinite state discrete event systems modelled by Symbolic Transition Systems In [KGMM11], we propose algorithms for the synthesis of state-feedback controllers with partial observation. We provide models of safe memoryless controllers both for potentially deadlocking and deadlock free controlled systems. In [KGMM11b], we consider the computation of safe decentralized controllers ensuring the avoidance of a set of forbidden states and then extend this result to the deadlock free case. In both cases, abstract interpretation techniques ensure termination at the price of an overapproximation of the transitions to disable. Our tool SMACS gives an empirical validation of our methods by showing their feasibility, usability and efficiency.

Control of Distributed Systems (INRIA) In [KGMM11d], we consider the control of distributed systems communicating asynchronously in the context of the state avoidance problem. Local controllers can only observe the behavior locally and use the FIFO queues to communicate by piggybacking timestamps and state estimates.
We provide an algorithm that computes at runtime an estimate of the current global state of the distributed system. Abstract interpretation is used to overapproximations queue contents of global states. An implementation of our algorithms provides an empirical evaluation of our method [KGMM11c].

WCET Analysis with MRU Caches: Challenging LRU for Predictability (Uppsala) Most previous work in cache analysis for WCET estimation assumes a particular replacement policy called LRU. In contrast, much less work has been done for non-LRU policies, since they are generally considered to be unpredictable. However, most commercial processors are actually equipped with these non-LRU policies, since they are more efficient in terms of hardware cost, power consumption and thermal output, but still maintaining almost as good average-case performance as LRU. In this work, we study the analysis of MRU, a non-LRU replacement policy employed in mainstream processor architectures like Intel Nehalem. Our work shows that the predictability of MRU has been significantly underestimated before, mainly because the existing cache analysis techniques and metrics, originally designed for LRU, do not match MRU well. As our main technical contribution, we propose a new cache hit/miss classification, k-Miss, to better capture the MRU behavior, and develop formal conditions and efficient techniques to decide the k-Miss memory accesses. A remarkable feature of our analysis is that the k-Miss classifications under MRU are derived by the analysis result of the same program under LRU. Therefore, our approach inherits all the advantages in efficiency, precision and composability of the state-of-the-art LRU analysis techniques based on abstract interpretation. Experiments with benchmarks show that the estimated WCET by our proposed MRU analysis is rather close to (5% - 20% more than) that obtained by the state-of-the-art LRU analysis, which indicates that MRU is also a good candidate for the cache replacement policy in real-time systems.

Scheduling of Certifiable Mixed-Criticality Task Systems (Uppsala) An increasing trend in embedded system design is to integrate components with different levels of criticality into a shared hardware platform for better cost and power efficiency. Such mixed-criticality systems are subject to certifications at different levels of rigorousness, for validating the correctness of

different subsystems on various confidence levels. The realtime scheduling of certifiable mixed-criticality systems has been recognized to be a challenging problem, where using traditional scheduling techniques may result in unacceptable resource waste. In this work, we present an algorithm called PLRS to schedule certifiable mixed-criticality sporadic tasks systems. PLRS uses fixed-job-priority scheduling, and assigns job priorities by exploring and balancing the asymmetric effects between the workload on different criticality levels. Comparing with the state-of-the-art algorithm by Li and Baruah for such systems, which we refer to as LB, PLRS is both more effective and more efficient: (i) The schedulability test of PLRS not only theoretically dominates, but also on average significantly outperforms LBs. (ii) The run-time complexity of PLRS is polynomial (quadratic in the number of tasks), which is much more efficient than the pseudo-polynomial run-time complexity of LB.

Hardness results on the Analysis of Digraph-Based Task Models (Uppsala) In formal analysis of real-time systems, a major concern is the analysis efficiency. As the expressiveness of models grows, so grows the complexity of their analysis. A recently proposed model, the digraph real-time task model (DRT), offers high expressiveness well beyond traditional periodic task models. Still, the associated feasibility problem on preemptive uniprocessors remains tractable. It is an open question to what extent the expressiveness of the model can be further increased before the feasibility problem becomes intractable. In this work, we study that tractability border. We show that system models with the need for global timing constraints make feasibility analysis intractable. However, our second technical result shows that it remains tractable if the number of global constraints is bounded by a constant. Thus, this paper establishes a precise borderline between tractability and intractability. A recent work along this line shows that the feasibility problem of fixed-priority scheduling is NP-hard.

Interpolation in Extensions of Presburger Arithmetic (Uppsala) Craig interpolation has emerged as an effective means of generating candidate program invariants. We presented interpolation procedures for the theories of Presburger arithmetic combined with (i) uninterpreted predicates (QPA+UP), (ii) uninterpreted functions (QPA+UF) and (iii) extensional arrays (QPA+AR). We prove that none of these combinations can be effectively interpolated without the use of quantifiers, even if the input formulae are quantifier-free. We go on to identify fragments of QPA+UP and QPA+UF with restricted forms of guarded quantification that are closed under interpolation. Formluae in these fragments can easily be mapped to quantifier-free expressions with integer division. For QPA+AR, we formulated a sound interpolation procedure that potentially produces interpolants with unrestricted quantifiers.

Automatic detection of DMA races in multicore software (Uppsala) We present a method for DMA race analysis in C programs. Our method works by automatically instrumenting a program with assertions modeling the semantics of a memory flow controller. The instrumented program can then be analyzed using state-of-the-art software model checkers. We show that bounded model checking is effective for detecting DMA races in buggy programs. To enable automatic verification of the correctness of instrumented programs, we present a new formulation of k-induction geared towards software, as a proof rule operating on loops. Our techniques are implemented as a tool, Scratch, which we apply to a large set of programs supplied with the IBM Cell SDK, in which we discover a previously unknown bug. Our experimental results indicate that our k-induction method performs extremely well on this problem class. To our knowledge, this marks both the first application of k-induction to software verification, and the first example of software model checking in the context of heterogeneous multicore processors.

Power Isolation (Salzburg) Salzburg has proposed the concept of power isolation for EDF-scheduled hard real-time systems running periodic software tasks. A task is power-isolated if there exist lower and upper bounds on its power consumption independent of any other

concurrently running tasks. The variance between lower and upper bounds and the power consumption overhead determine quality and cost of power isolation, respectively. So far, lower and upper bounds on the power consumption of an EDF-scheduled, periodic task have been identified as functions of task utilization, frequency scaling, and power model.

Development of parametric quantitative models based on timed automata (TRENTO, ETHZ). In this work, we propose a rigorously formal and compositional style for obtaining key performance and/or interface metrics of systems with real-time constraints. We propose a hierarchical approach that couples the independent and different by nature frameworks of Modular Performance Analysis with Real-time Calculus (MPA-RTC) and Parametric Feasibility Analysis (PFA). Recent work on Real-time Calculus (RTC) has established an embedding of state-based component models into RTC-driven performance analysis for dealing with more expressive component models. However, with the obtained analysis infrastructure it is possible to analyze components only for a fixed set of parameters, e. g., fixed CPU speeds, fixed buffer sizes etc., such that a big space of parameters remains unstudied. In this work, we overcome this limitation by integrating the method of parametric feasibility analysis in an RTC-based modeling environment. Using the PFA tool-flow, we are able to find regions for component parameters that maintain feasibility and worst-case properties. As a result, the proposed analysis infrastructure produces a broader range of valid design candidates, and allows the designer to reason about the system robustness. The parametric analysis technique, which was enhanced in this work, uses a combination of the NuSMV and the Uppaal tool to improve the efficiency of the parameter space exploration. The translator takes care of the semantics adaptation, and in particular updates the model with the required error states necessary to represent the properties of interest. The results are reported in conference proceedings [SRPL11].

Specification formalisms for consumption (CISS) The problems of time-dependent behavior in general, and dynamic resource allocation in particular, pervade many aspects of modern life. Prominent examples range from reliability of efficient use of communication resources in a telecommunication network to allocation of tracks in a continental railway network, from scheduling the usage of computational resources on a chip for durations of nano-seconds to weekly, monthly or longer-range reactive planning in a factory or supply chains. The invited CACM paper [BFLM11] provides a full account of *priced timed automata*, which is an extension of timed automata with additional continuous cost variables, observer variables growing with positive – but possibly varying – rate. This model is particularly useful for modeling additional consumption of resource, e.g. energy. A number of decision problems related to priced timed automata and the pricinple underlying how they are settle is provided.

Energy games have recently abstracted a lot of attention. These are games played on finite weighted automata and concern the existence of infinite runs subject to boundary constraints on the accumulated weight, allowing e.g. only for behaviours where a resource is always available (nonnegative accumulated weight), yet does not exceed a given maximum capacity. In [FJLS11] we extend energy games to a multiweighted and parameterized setting, allowing us to model systems with multiple quantitative aspects. We present reductions between Petri nets and multiweighted automata and among different types of multiweighted automata and identify new complexity and (un)decidability results for both one- and two-player games. We also investigate the tractability of an extension of multiweighted energy games in the setting of timed automata.

Development of and around UPPAAL (CISS) Uppaal is a tool suitable for model checking real-time systems described as networks of timed automata communicating by channel synchronisations and extended with integer variables. Its first version was released in 1995

and its development is still very active. It now features an advanced modelling language, a userfriendly graphical interface, and a performant model checker engine. In addition, several flavours of the tool have matured in recent years. In [BDLPW11] we present how we managed to maintain the tool during 15 years, its current architecture with its challenges, and we give future directions of the tool.

In [DHJLOOS11] we present a new open source model checker, opaal, for automatic verification of models using lattice automata. La□ce automata allow the users to incorporate abstractions of a model into the model itself. This provides an efficient verification procedure, while giving the user fine-grained control of the level of abstraction by using a method similar to Counter-Example Guided Abstraction Refinement. The opaal engine supports a subset of the UPPAAL timed automata language extended with lattice features. We report on the status of the first public release of opaal, and demonstrate how opaal can be used for efficient verification on examples from domains such as database programs, lossy communication protocols and cache analysis.

Timed-Arc Petri Nets (TAPN) are an extension of the classical P/T nets withccontinuous time. Tokens in TAPN carry an age and arcs between places and transitions are labelled with time intervals restricting the age of tokens available for transition firing. The TAPN model posses a number of interesting theoretical properties distinguishing them from other time extensions of Petri nets. [JJMS11] give an overview of the recent theory developed in the verification of TAPN extended with features like read/transport arcs, timed inhibitor arcs and age invariants. We will examine in detail the boundaries of automatic verification and the connections between TAPN and the model of timed automata. Finally, we will mention the tool TAPAAL that supports modelling, simulation and verification of TAPN and discuss a small case study of alternating bit protocol

Data and Parametric Extensions of Modal Transition Systems (CISS) Specification theories as a tool in the development process of component based software systems have recently abstracted a considerable attention. Current specification theories are however qualitative in nature and hence fragile and unsuited for modern software systems. In [BFJLLT11] we propose the first specification theory which allows to capture quantitative aspects during the refinement and implementation process.

[BJLLS11] introduce a novel formalism of label-structured modal transition systems that combines the classical may/must modalities on transitions with structured labels that represent quantitative aspects of the model. On the one hand, the specificition formalism is general enough to include models like weighted modal transition systems and allows the system developers to employ even more complex label refinement than in previously studied theories. On the other hand, the formalism maintains the desirable properties required by any specification theory supporting compositional reasoning. In particular, we study modal and thorough refinement, determinization, parallel composition, conjunction, quotient, and logical characterization of label-structured modal transition systems.

Modal transition systems (MTS) is a well-studied specification formalism of reactive systems supporting a step-wise refinement methodology. Despite its many advantages, the formalism as well as its currently known extensions are incapable of expressing some practically needed aspects in the refinement process like exclusive, conditional and persistent choices. [BKLMS11] introduce a new model called parametric modal transition systems (PMTS) together with a general modal refinement notion that overcome many of the limitations and we investigate the computational complexity of modal refinement checking.

Modal Specifications (MSs) is a well-known and widely used abstraction theory for transition systems. MSs are transition systems equipped with two types of transitions: must transitions that are mandatory to any implementation, and may transitions that are optional. The duality of transitions allows to develop a unique approach for both logical and structural compositions, and ease the step-wise refinement process for building implementations. [BLLNW11] propose Modal Specifications with Data (MSDs), the first modal specification theory with explicit representation of data. Our new theory includes all the essential ingredients of a specification theory. As MSDs are by nature potentially infinite-state systems, we also propose symbolic representations based on effective predicates and show equivalence with the semantic definitions. Our theory serves as a new abstraction-based formalism for transition systems with data.

Modular Markovian Logic (CISS) In [MCL11] and [CLM11] we introduce Modular Markovian Logic (MML) for compositional continuous-time and continuous-space Markov processes. MML combines operators specific to stochastic logics with operators reflecting the modular structure of the models, similar to those used by spatial and separation logics. We present a complete Hilbert-style axiomatization for MML, prove the small model property and analyze the relation between stochastic bisimulation and logical equivalence.

Statistical Model Checking for Timed Automata (CISS with INRIA) The paper [DLLMPVW11] offers a natural stochastic semantics of Networks of Priced Timed Automata (NPTA) based on races between components. The semantics provides the basis for satisfaction of Probabilistic Weighted CTL properties (PWCTL), conservatively extending the classical satisfaction of timed automata with respect to TCTL. In particular the extension allows for hard real-time properties of timed automata expressible in TCTL to be refined by performance properties, e.g. in terms of probabilistic guarantees of time- and cost-bounded properties. A second contribution is the application of Statistical Model Checking (SMC) to efficiently estimate the correctness of non-nested PWCTL model checking problems with a desired level of confidence, based on a number of independent runs of the NPTA. In addition to applying classical SMC algorithms, we also offer an extension that allows to efficiently compare performance properties of NPTAs in a parametric setting. The third contribution is an efficient tool implementation within UPPAAL of our result and applications to several case studies [DLLMW11].

Sub-activity C: Cross-layer Validation

.

Abstracting time and data for conformance testing of real-time systems (INRIA) The paper [AMJM11] with Federal University Campina Grande (Brasil) proposes an extension of timed-automata with data on infinite domains. We adapt the tioco conformance testing theory to deal with this model and describe a test case generation process based on a combination of symbolic execution and constraint solving for the data part and symbolic analysis for timed aspects.

Asynchronous Testing (INRIA) We addressed the issue of testing a component of an asynchronously communicating distributed system. Testing a system which communicates asynchronously (i.e., through some medium) with its environment is more difficult than testing a system which communicates synchronously (i.e., directly without any medium): the actual behavior of the implementation under test (IUT) appears distorted and infinite to the tester. To this end, the paper [B11] proposes a tagging protocol which when implemented by the

asynchronously communicating distributed system will make the problem of generating a complete test suite, from the specification of any of its component, feasible.

Model-based Safety Engineering of Interdependent Functions (KTH and Volvo)
For systems where functions are distributed but share support for computation, communication, environment sensing, and actuation, it is essential to understand how such functions can affect each other. Preliminary Hazard Analysis (PHA) is the task through which safety requirements are established. This is usually a document-based process where each system function is analyzed alone, making it difficult to reason about the commonalities of related functional concepts and the distribution of safety mechanisms across a system-of-systems. This work explored a model-based approach to PHA with the EAST-ADL2 language and in accordance with the ISO/DIS 26262 functional safety standard.
The language explicitly supports the definition and handling of requirements, functions and technical solutions, and their various relations and constraints as a coherent whole with multiple views. We have shown in particular the engineering needs for a systematic approach to PHA and the related language features for precise modeling of requirements, management of functions and their interdependencies, and the reasoning of safety mechanisms.
As part of collaborative research in the Maenad project, further work has been carried out to extend the support for modeling and analysis of safety critical automotive embedded systems. The work has included analysis of the ISO26262 safety lifecycle (ISO26262 is the new automotive standard for functional safety), including the key tasks and artefacts, and in providing an overall assessment of what the EAST-ADL architecture description language provides to support the ISO26262 requirements, [TT11]. Recent advances of EAST-ADL support for the design of functional safety and its integration with system architecture design is presented in [CJ11]. The work complements and consolidates our earlier work, [SC11], [CJ08], by introducing in detail the adopted (meta) modeling pattern as well as the language alignment with ISO26262 in regard to the safety life-cycle, safety requirements, architectural and analytical artifacts.

Implementation and Empirical Comparison of Multi-core Scheduling Algorithms (Uppsala) Recent theoretical studies have shown that partitioning-based scheduling has better real-time performance than other scheduling paradigms like global scheduling on multi-cores. Especially, a class of partitioning-based scheduling algorithms (called semi-partitioned scheduling), which allow to split a small number of tasks among different cores, offer very high resource utilization. The major concern about the semi-partitioned scheduling is that due to the task splitting, some tasks will migrate from one core to another at run time, which incurs higher context switch overhead. So one would suspect whether the extra overhead caused by task splitting would counteract the theoretical performance gain of semi-partitioned scheduling. In this work, we implement a semi-partitioned scheduler in the Linux operating system, and run experiments on an Intel Core-i7 4-cores machine to measure the real overhead in both partitioned scheduling and semi-partitioned scheduling. Then we integrate the measured overhead into the state-of-theart partitioned scheduling and semi-partitioned scheduling algorithms, and conduct empirical comparisons of their realtime performance. Our results show that the extra overhead caused by task splitting in semi-partitioned scheduling is very low, and its effect on the system schedulability is very small. Semi-partitioned scheduling indeed outperforms partitioned scheduling in realistic systems.

Resource Sharing Protocols for Task Graph Systems (Uppsala) Previous works on real-time task graph models have ignored the crucial resource sharing problem. Due to the nondeterministic branching behavior, resource sharing in graphbased task models is

significantly more difficult than in the simple periodic or sporadic task models. In this work we address this problem with several different scheduling strategies, and quantitatively evaluate their performance. We first show that a direct application of the well-known EDF+SRP strategy to graph-based task models leads to an unbounded speedup factor. By slightly modifying EDF+SRP, we obtain a new scheduling strategy, called EDF+saSRP, which has a speedup factor of 2. Then we propose a novel resource sharing protocol, called ACP, to better manage resource sharing in the presence of branching structures. The scheduling strategy EDF+ACP, which applies ACP to EDF, can achieve a speedup factor of 1.618, which is the golden ratio.

Test-case generation for embedded Simulink via formal concept analysis (Uppsala) Mutation testing suffers from the high computational cost of automated test-vector generation, due to the large number of mutants that can be derived from programs and the cost of generating test-cases in a white-box manner. We propose a novel algorithm for mutation-based test-case generation for Simulink models that combines white-box testing with formal concept analysis. By exploiting similarity measures on mutants, we are able to effectively generate small sets of short test-cases that achieve high coverage on a collection of Simulink models from the automotive domain. Experiments show that our algorithm performs significantly better than random testing or simpler mutation-testing approaches.

Learning Probabilistic Automata (CISS) Obtaining accurate system models for verification is a hard and time consuming process, which is seen by industry as a hindrance to adopt otherwise powerful model-driven development techniques and tools. In [CMJNLN11] we pursue an alternative approach where an accurate high-level model can be automatically constructed from observations of a given black-box embedded system. We adapt algorithms for learning finite probabilistic automata from observed system behaviors. We prove that in the limit of large sample sizes the learned model will be an accurate representation of the data-generating system. In particular, in the large sample limit, the learned model and the original system will define the same probabilities for linear temporal logic (LTL) properties. Thus, we can perform PLTL model-checking on the learned model to infer properties of the system. We perform experiments learning models from system observations at different levels of abstraction. The experimental results show the learned models provide very good approximations for relevant properties of the original system.

-- The above is new material, not present in the Y3 deliverable --

3.2 Individual Publications Resulting from these Achievements

INRIA

[BG11] N. Bertrand, B. Genest. Minimal Disclosure in Partially Observable Markov Decision Processes. In IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, Bombay, Inde, December 2011.

[B11] P. Bhateja. A Tagging Protocol for Asynchronous Testing. In 5th IEEE International Conference on Theoretical Aspects of Software Engineering, August 2011.

[AMJM11] W. L. Andrade, P. D. L. Machado, T. Jéron, H. Marchand. Abstracting Time and Data for Conformance Testing of Real-Time Systems, in "7th Workshop on Advances in Model Based Testing A-MOST 2011", Berlin, Germany, 2011, http://hal.inria.fr/hal-00646089/en.

[BJSK11] N. Bertrand, T. JÃ©ron, A. Stainer, M. Krichen. Off-line Test Selection with Test Purposes for Non-Deterministic Timed Automata. In 17th International Conference on Tools and Algorithms for the Construction And Analysis of Systems (TACAS), LNCS, Volume 6605, Pages 96-111, Saarbücken, Germany, April 2011.

[BSJK11] N. Bertrand, A. Stainer, T. Jéron, M. Krichen. A game approach to determinize timed automata. In 14th International Conference on Foundations of Software Science and Computation Structures (FOSSACS), LNCS, Volume 6604, Pages 245-259, SaarbrÃ¼cken, Germany, April 2011.

UPPSALA

[GLGY12] Nan Guan, Mingsong Lv, Yu Ge and Wang Yi. WCET Analysis with MRU Caches: Challenging LRU for Predictability. To appear in the proc. of RTAS 2012.

[GSGY12] Nan Guan, Martin Stigge, Yu Ge, and Wang Yi. Parametric Utilization Bounds for Fixed-Priority Multiprocessor Scheduling. In the proc. of the 26th IEEE International Parallel and Distributed Processing Symposium. May 21-25, 2012, Shanghai, China.

[GESY11a] Nan Guan, Pontus Ekberg, Martin Stigge and Wang Yi. Effective and Efficient Scheduling of Certifiable Mixed-Criticality Sporadic Task Systems. In the proc. of IEEE RTSS 2011, Vienna, Austria. Nov 30 - Dec 2, 2011.

[LGDYY11] Mingsong Lv, Nan Guan, Qingxu Deng, Ge Yu, Wang Yi: McAiT - A Timing Analyzer for Multicore Real-Time Software. In the proc. Of ATVA 2011: 414-417.

[ZGXY11] Yi Zhang, Nan Guan, Yanbin Xiao, Wang Yi. Implementation and Empirical Comparison of Partitioning-based Multi-core Scheduling. In the proc. of the 6th IEEE International Symposium on Industrial Embedded Systems (SIES11), Vaesteraas, Sweden, June 15th - 17th, 2011.

[SEGY11a] Martin Stigge, Pontus Ekberg, Nan Guan, and Wang Yi. On the Tractability of Digraph-Based Task Models. In the proc of the 23rd Euromicro Conference on Real-Time Systems, Porto, Portugal July 6th - 8th, 2011.

[JGDY11] Xi Jin, Nan Guan, Qingxu Deng, Wang Yi. Memory Aware Mapping for Network-on-Chips, In the proc. of the 17th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA 2011), August 28-31, Toyama, Japan.

[GESY11b] Nan Guan, Pontus Ekberg, Martin Stigge and Wang Yi. Resource Sharing Protocols for Real-Time Task Graph Systems. In the proc. of the 23rd Euromicro Conference on Real-Time Systems. Porto, Portugal July 6th - 8th, 2011. (pdf)

[SEGY11b] Martin Stigge, Pontus Ekberg, Nan Guan and Wang Yi. The Digraph Real-Time Task Model. In the proc. of the 17th IEEE Real-Time and Embedded Technology and Applications Symposium, Chicago, IL, USA April 11 - 14, 2011. Best Paper Nomination.

[GYDGY11] Nan Guan, Wang Yi, Qingxu Deng, Zonghua Gu, Ge Yu. Schedulability analysis for non-preemptive fixed-priority multiprocessor scheduling. Journal of Systems Architecture – Embedded Systems Design 57(5): 536-546 (2011).

[KYD11] Fanxin Kong, Wang Yi, Qingxu Deng. Energy-efficient scheduling of real-time tasks on cluster-based multicores. In the proc. of DATE 2011: 1135-1140.

[KGDY11] Fanxin Kong, Nan Guan, Qingxu Deng, Wang Yi. Energy-efficient scheduling for parallel real-time tasks based on level-packing. In the proc of ACM SAC 2011: 635-640.

[Ruemmer12] Philipp Reummer. E-Matching with Free Variables. To appear in 18th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Marida, Venezuela, March 2012

[DHRK12] Alastair F. Donaldson, Nannan He, Philipp Reummer, Daniel Kroening. Tightening test coverage metrics: A case study in equivalence checking using k-induction. 9th International Symposium on Formal Methods for Components and Objects (FMCO), 2010, Graz, Austria. Revised papers

[BKRW11-2] Angelo Brillout, Daniel Kroening, Philipp RÃeummer, Thomas Wahl. An Interpolating Sequent Calculus for Quantifier-Free Presburger Arithmetic. Journal of Automated Reasoning, Volume 47 / 2011, Issue 4, Pages 341-367

[DKR11-2] Alastair F. Donaldson, Daniel Kroening, Philipp Reummer. Automatic Analysis of DMA Races Using Model Checking and k-induction. International Journal on Formal Methods in System Design (FMSD), Volume 39 / 2011, Number 1, pages 83-113

[DHKR11] Alastair F. Donaldson, Leopold Haller, Daniel Kroening, Philipp Reummer. Software Verification Using k-Induction. 18th International Static Analysis Symposium (SAS), Venice, Italy, September 2011

[HRK11] Nannan He, Philipp Reummer, Daniel Kroening. Test-case generation for embedded Simulink via formal concept analysis. Design Automation Conference (DAC), San Diego, USA, June 2011

[DKR11] Alastair F. Donaldson, Daniel Kroening, Philipp Reummer. SCRATCH: a tool for automatic analysis of DMA races. Poster at16th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP), San Antonio, TX, USA, 2011

[BKRW11] Angelo Brillout, Daniel Kroening, Philipp Reummer, Thomas Wahl. Beyond Quantifier-Free Interpolation in Extensions of Presburger Arithmetic. 12th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI), Austin, Texas, January 2011.

Saltzburg

[CKS11] S.S. Craciunas, C.M. Kirsch, and A. Sokolova. The Power of Isolation. Technical Report 2011-02. Department of Computer Sciences, University of Salzburg, July, 2011.

Trento

[RPMP11] Tizar Rizano, Roberto Passerone, David Macii and Luigi Palopoli, Model-Based Design of Embedded Control Software for Hybrid Vehicles. In Proceedings of the 6th IEEE International Symposium on Industrial Embedded Systems (SIES11), Västerås, Sweden, June 15-17, 2011.

CEA LIST

[BGS11] BANNOUR B., GASTON C., SERVAT D., "Eliciting unitary constraints from timed Sequence Diagram with symbolic techniques: application to testing" to appear in Proceedings of Asian-Pacific Software Engineering Conference APSEC 2011

[EGL11] ESCOBEDO J.P., GASTON C., Le GALL P., "Timed Conformance Testing for Orchestrated Service Discovery", 8th International Symposium on Formal Aspects of Component Software: FACS 2011, September 14-16 2011

CISS (Aalborg)

[BFLM11] Patricia Bouyer, Ulrich Fahrenberg, Kim G. Larsen, and Nicolas Markey. Quantitative modelling and analysis of embedded systems. Communications of the ACM, 2011. Invited paper.

[FTL11] Uli Fahrenberg, Claus Thrane, and Kim G. Larsen. Distances for weighted transition systems: Games and properties. In In Proceedings of Ninth Workshop on Quantitative Aspects of Programmaming Languages (QAPL'11), Electronic Proceedings in Theoretical Computer Science, 2011.

[BDLPW11] Gerd Behrmann, Alexandere David, Kim G. Larsen, Paul Pettersson, and Wang Yi. Developing UPPAAL over 15 years. Software - Practice and Experience, 41(2):133–142, 2011.

[LFT11] Kim G. Larsen, Uli Fahrenberg, and Claus Thrane. Metrics for weighted transition systems: Axiomatization and complexity. Theor. Comput. Sci., 412(28):3358–3369, 2011

[CMJNLN11] Yingke Chen, Hua Mao, Manfred Jaeger, Thomas D. Nielsen, Kim G. Larsen, and Brian Nielsen. Learning probabilistic automata for model checking. In The Eighth International Conference on Quantitative Evaluation of SysTems (QEST 2011). Accepted., Springer LNCS, 2011.

[MCL11] Radu Mardare, Luca Cardelli, and Kim G. Larsen. Continuous markovian logic - from complete axiomatization to the metric space of formulas. In Computer Science Logic (CSL), 25th International Workshop / 20th Annual Conference of the EACSL, CSL 2011, September 12-15, 2011, Bergen, Norway, Proceedings, LIPIcs, pages 144–158. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik,2011.

[CLM11] Luca Cardelli, Kim G. Larsen, and Radu Mardare. Modular markovian logic. In Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II, pages 380–391, 2011.

[FJLS11] Uli Fahrenberg, Line Juhl, Kim G. Larsen, and Jiri Srba. Energy games in multiweighted automata. In Proceedings of Theoretical Aspects of Computing (ICTAC'11) - 8th

International Colloquium,Johannesburg, South Africa, volume 6916 of Lecture Notes in Computer Science, pages 95–115. Springer, 2011.

[BLMST11] Patricia Bouyer, Kim G. Larsen, Nicolas Markey, Ocan Sankur, and Claus Thrane. Timed automata can always be made implementable. In CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings, volume 6901 of Lecture Notes in Computer Science, pages 76–91, 2011.

DHJLOOS11] Andreas Engelbredt Dalsgaard, René Rydhof Hansen, Kenneth Yrke Jørgensen, Kim Guldstrand Larsen, Mads Chr. Olesen, Petur Olsen, and Jirí Srba. opaal: A lattice model checker. In NASA Formal Methods - Third International Symposium, NFM 2011, Pasadena, CA, USA, April 18-20, 2011. Proceedings, pages 487–493, 2011.

[BKLMS11] Beneš, J. Křesinský, K.G. Larsen, M.H. Møller, and J. Srba. Parametric modal transition systems. In Proceedings of the 9th International Symposium on Automated Technology for Verification and Analysis (ATVA'11), LNCS. Springer-Verlag, 2011.

[BHKLO11] Jörg Brauer, René Rydhof Hansen, Stefan Kowalewski, Kim G. Larsen, and Mads Chr. Olesen.Adaptable value-set analysis for low-level code. In Proceedings of The 6th International Workshop on Systems Software Verification (SSV2011), 2011.

[FLeT11] Uli Fahrenberg, Axel Legay, and Claus Thrane. The quantitative linear-time–branching-time spectrum. In Proceedings of IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), 2011.

[JJMS11] L. Jacobsen, M. Jacobsen, M.H. Møller, and J. Srba. Verification of timed-arc Petri nets. In Proceedings of the 37th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM'11), volume 6543 of LNCS, pages 46–72. Springer-Verlag, 2011.

.

-- The above are new references, not present in the Y3 deliverable --

3.3 Interaction and Building Excellence between Partners

(1) Uppsala is collaborating with Absint, ETH at Zurich, TU braunschweig, and Verimag on Mixed Criticality Systems (MCS).

(2) Uppsala is collaborating with ETH at Zurich to combine analytic methods with model checking for efficient timing analysis.

(3) The work of Saltzburg on power isolation is part of a recent initiative in rigorous systems engineering (RiSE) with nine partners in Austria including IST Austria.

(4) Uppsala and CISS has continued collaboration on the distribution, maintainance and further dissemination of UPPAAL

(5) CISS and INRIA has collaborated on development compositional specification theories for probabilistic systems as well as the development of statistical model checking for networks of timed automata.

(6) CISS and LSV has collaborated on the development of priced timed automata, energy automata, energy games and robustness for timed automata.

(7) CISS, INRIA and RWTH has collaborated on a specification theory based on abstract probabilistic automata.

-- Changes wrt Y3 deliverable --

This section is completely new with respect to Y3 deliverable.

3.4 Joint Publications Resulting from these Achievements

INRIA and Verimag

[FFM11] Y. Falcone, J-C Fernandez, L. Mounier. What can you verify and enforce at runtime?. International Journal on Software Tools for Technology Transfer (STTT), 2011.

[FMFR11] Y. Falcone, L. Mounier, Fernandez J.-C, J.-L. Richier. Runtime enforcement monitors: composition, synthesis, and enforcement abilities. Formal Methods in System Design, 2011.

INRIA and ULB

[KGMM11] G. Kalyon, T. Le Gall, H. Marchand, T. Massart. Symbolic Supervisory Control of Infinite Transition Systems under Partial Observation using Abstract Interpretation. Discrete Event Dynamic System: Theory and Applications, 2011.

[KGMM11b] G. Kalyon, T. Le Gall, H. Marchand, T. Massart. Decentralized Control of Infinite Systems. Discrete Event Dynamic Systems : Theory and Applications, 21(3):359-393, September 2011.

[KGMM11c] G. Kalyon, T. Le Gall, H. Marchand, T. Massart. Synthesis of Communicating Controllers for Distributed Systems. In 50th IEEE Conference on Decision and Control and European Control Conference, Pages 198-212, Orlando, USA, December 2011.

[KGMM11d] G. Kalyon, T. Le Gall, H. Marchand, T. Massart. Global State Estimates for Distributed Systems. In 31th IFIP International Conference on FORmal TEchniques for Networked and Distributed Systems, FORTE, LNCS, Volume 6722, Pages 198-212, Reykjavik, Iceland, June 2011.

INRIA and LSV:

[BBBS11] N. Bertrand, P Bouyer, Th. Brihaye, A. Stainer. Emptiness and Universality Problems in Timed Automata with Positive Frequency. In Proceedings of the 38th International Colloquium on Automata, Languages and Programming (ICALP'11), LNCS, Pages 246-257, ZÃ¼rich, Switzerland, July 2011.

CISS and Uppsala

[BDLPY11] Gerd Behrmann, Alexandre David, Kim Guldstrand Larsen, Paul Pettersson and Wang Yi. Developing UPPAAL over 15 years. In Journal: Software - Practice and Experience, 41(2): 133-142 (2011).

TRENTO and Rennes

[RBBC11] Jean-Baptiste Raclet, Eric Badouel, Albert Benveniste, Benoît Caillaud, Axel Legay and Roberto Passerone. A Modal Interface Theory for Component-based Design. Fundamenta Informaticae, 108(1-2):119-149, 2011.

Trento and ETHZ

SRPL11] Alena Simalatsar, Yusi Ramadian, Roberto Passerone, Kai Lampka, Simon Perathoner and Lothar Thiele. Enabling Parametric Feasibility Analysis in Real-time Calculus Driven Performance Evaluation. In Proceedings of the International Conference on Compilers, Architectures and Synthesis of Embedded Systems (CASES11), Taipei, Taiwan, October 9-14, 2011.

CISS and INRIA

[DLLPW11] Benoit Delahaye, Kim G. Larsen, Axel Legay, Mikkel Larsen Pedersen, and Andrzej Wasowski. Decision problems for interval markov chains. In Proceedings of the 5th International Conference on Language and Automata Theory and Applications (LATA), 2011.

[CDLLPW11] Benoît Caillaud, Benoît Delahaye, Kim G. Larsen, Axel Legay, Mikkel L. Pedersen, and Andrzej Wasowski. Constraint markov chains. Theoretical Computer Science (TCS), 412(34):4373 – 4404, 2011.

[BFJLLT11] Sebastian S. Bauer, Uli Fahrenberg, Line Juhl, Kim G. Larsen, Axel Legay, and Claus Thrane. Quantitative refinement for weighted modal transition systems. In Mathematical Foundations of Computer Science 2011 - 36th International Symposium, MFCS 2011, Warsaw, Poland, August 22-26, 2011. Proceedings, volume 6907 of LNCS, pages 60–71. Springer-Verlag, 2011.

[DLLMW11] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, and Zheng Wang. Time for statistical model checking of real-time systems. In Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011., pages 349–355, 2011.

[DLLMPVW11] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, Danny B. Poulsen, Jonas V. Vliet, and Zheng Wang. Statistical model checking for networks of priced timed automata. 2011. In Proceedings of FORMATS 2011.

[DLLPW11] Benoit Delahaye, Kim G. Larsen, Axel Legay, Mikkel Larsen Pedersen, and Andrzej Wasowski. Apac: a tool for reasoning about abstract probabilistic automata. 2011. To appear in Proceedings of QEST 2011.

[BJLLS11] Sebastian S. Bauer, Line Juhl, Kim G. Larsen, Axel Legay, and Jiri Srba. Extending modal transition systems with structured labels. Mathematical Structures in Computer Science, 2011.

LLTW11] Kim G. Larsen, Axel Legay, Louis-Marie Traonouez, and Andrzej Wasowski. Robust specification of real time components. In Uli Fahrenberg and Stavros Tripakis, editors, 9th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS 2011), volume 6919 of Lecture Notes in Computer Science, pages 129–144, Aalborg, Denmark, September 2011.

[BLLNW11] Sebastian Bauer, Kim G. Larsen, Axel Legay, Ulrik Nyman, and Andrzej Wąsowski. A modal specification theory for components with data. In 8th International Symposium on Formal Aspects of Component Software, Oslo, Norway, September 14-16, 2011, 2011. (Best Paper Award)

CISS, INRIA and RWTH

[KSDLLPW11] Joost-Pieter Katoen, Falak Sher, Benoit Delahaye, Kim G. Larsen, Axel Legay, Mikkel Larsen Pedersen, and Andrzej Wasowski. Abstract probabilistic automata. In Proceedings of the 12th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI), 2011.

[KSDLLPW11b] Joost-Pieter Katoen, Falak Sher, Benoit Delahaye, Kim G. Larsen, Axel Legay, Mikkel Larsen Pedersen, and Andrzej Wasowski. New results on abstract probabilistic automata. In Proceedings of the 11th International Conference on Application of Concurrency to System Design (ACSD), 2011.

-- The above are new references, not present in the Y3 deliverable --

3.5 Keynotes, Workshops, Tutorials

Keynote: The disappearing computer
Twan Basten – Devlab Café, Development Laboratories, Eindhoven, the Netherlands, 29 April 2011

Keynote: Performance prediction and optimization for Wafer Scanners
Jeroen Voeten - Dutch Model Checking Day 2011, Delft, the Netherlands, 17 June 2011

Keynote: Using a Commercial Model Checker at Philips Healthcare
Jozef Hooman - System Validation seminar, University of Twente, the Netherlands, 23 May 2011

Keynote: Compositional Model Checking using Verum's ASD:Suite at Philips Healthcare
Jozef Hooman - MBSD seminar, Radboud University, Nijmegen, the Netherlands, 1 July 2011

Keynote: Experiences with a Compositional Model Checker in the Healthcare Domain
Jozef Hooman - International Symposium on Foundations of Health Information Engineering and Systems (FHIES 2011), Johannesburg, South Africa, 30 August 2011

Keynote: AUTOSAR Timing Extension and a Case Study for Schedulability Analysis
Sara Tucci - ArtistDesign Workshop on Real-Time System Models for Schedulability analysis University of Cantabria 7-8 February 2011

Keynote: Applying Model Driven Engineering to RTES: Technologies, Standards and Experiences
Sara Tucci - ES-week Workshop on Time Analysis and Model-Based Design, from Functional Models to Distributed Deployments, Taipei, 2011

Keynote: The Digraph Real-Time Task Model, Wang Yi, invited talk, Workshop on Rigorous Embedded Design 2011, April 10th, 2011, Salzburg, Austria (within EuroSys 2011).

Keynote Lecture: Thomas A. Henzinger
Computational Science versus Computer Science, Ninth Basel Computational Biology Conference (BC2), Basel, Switzerland, June 2011.

Invited talk, Kim G Larsen: The 9th International Workshop on Java Technologies for Real-time and Embedded Systems - JTRES 2011, York 26-28 October 2011. Timing and Performance Analysis of Embedded Software Systems Using Model Checking.

Invited talk, Kim G Larsen: De 17e Nederlandse Testdag, 29 November 201. University of Twente, Enschede, The Netherland.

Invited talk, Kim G Larsen: ARTIST Summer School, Aix-les-Bains, France, September 4-9, 2011. /wwwtcs.inf.tu-dresden.de/wata2012/

Invited taklk, Kim G. Larsen: ARTIST Summer School in China, IOS/ISCAS, Beijing, August 8-12, 2011.
www.artist-embedded.org/artist/Overview,2239.html

Invited talk, Kim G Larsen; PDMC, 10th International Workshop on Parallel and Distributed Methods in verifiCation, July 14, 2011, Cliff Lodge, Snowbird, Utah.
www.pdmc.cz/PDMC11

Invited Panelist, Kim G. Larsen: Microsoft Software Summit , Paris, France, April 14, 2011, research.microsoft.com/en-us/events/ss2011

Invited talk, Kim G Larsen: RED, Rigorous Embedded Systems, Salzburg, Austria, April 10, 2011. www.artist-embedded.org/artist/Programm,2288.html/

Invited talk, Kim G Larsen: iWIGP, International Workshop on Interaction, Games and Protocols, Saarbrücken, Germany, March 27, 2011.
www.etaps.org/programme/76-programmeiwigp

Invited talk, Kim G Larsen: ROCKS, Rigorous Dependability Analysis using Model Checking Techniques for Stochastic Systems, Workshop, March 26, Saarbrücken, 2011. www.etaps.org/programme/66-programmerocks/

Invited talk, Kim G Larsen: World Conference, Development Tools Sessions, Nürnberg, March 3, 2011.
Invited Lecture: Christoph Kirsch,
Virtualizing Time, Space, and Power for Cyber-Physical Cloud Computing, ARTIST Workshop on Rigorous Embedded Design, Salzburg, Austria, April 2011.

Invited Lecture: Thomas A. Henzinger,
From Boolean to Quantitative Synthesis, Eleventh Annual Conference on Embedded Software (EMSOFT), Taipei, Taiwan, October 2011.

Invited Lecture: Thomas A. Henzinger

Ten Years of Interface Automata, ACM SIGSOFT Impact Paper Award Lecture, 19th Annual Symposium on Foundations of Software Engineering (FSE), Szeged, Hungary, September 2001.

Invited Lecture: Thomas A. Henzinger
Quantitative Reactive Models, Workshop on Synthesis, Verification, and Analysis of Rich Models (SVARM), Saarbrucken, Germany, April 2011.

Invited Lecture: Thomas A. Henzinger
Formal Methods for Composing Systems, Design Automation and Test in Europe (DATE), Grenoble, France, March 2011.

Panelist: Christoph Kirsch, Vehicular Wireless Networks: What should the future hold? International Symposium on Wireless Vehicular Communications (WiVeC), San Francisco, California, September 2011.

Conference: The 6th IEEE International Symposium on Industrial Embedded Systems (SIES 2011), Mälardalen University, Västerås, Sweden.
June 15-17, 2011.
TRENTO has co-chaired this conference, which is concerned with all aspects related to modelling and developing embedded systems, with particular emphasis on their application in a variety of industrial environments. The considered application range from SoCs, which are making inroads in to the area of industrial automation, to automotive and safety-critical systems.
In particular, at this year conference, TRENTO and IST-Austria have organized a special session dedicated to various aspect of robust design with a keynote speech by Jean-François Raskin on the Synthesis of Robust Controller and Games With Imperfect Information, and a set of three invited papers on specification, control and design methodologies.

INRIA Rennes Gipsy Workshop on Games, Logic and Security in Nov. 2011 (http://www.irisa.fr/prive/pinchina/GIPSy/gipsy11.html).

CISS, Aalborg, 9th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS 2011, Phønix Hotel, Aalborg, Denmark, 21-23 September 2011 http://formats2011.cs.aau.dk/

Timing aspects of systems from a variety of computer science domains have been treated independently by different communities. Researchers interested in semantics, verification and performance analysis study models such as timed automata and timed Petri nets, the digital design community focusses on propagation and switching delays, while designers of embedded controllers have to take account of the time taken by controllers to compute their responses after sampling the environment.

Organizers: Alexandre David, Kim G Larsen, Claus Thrane, Rikke W. Uhrenholt

LSV, CISS ao.: 3rd Workshop on Games for Design, Verification and Synthesis. Colocated with CONCUR'11, Aachen (Germany), 10 September 2011 http://www.lsv.ens-cachan.fr/Events/gasics10/

The aim of this workshop was to bring together researchers working on game-related subjects, and to discuss on various aspects of game theory in the fields where it is applied. The workshop was composed of two invited talks, together with contributed talks on the following (non-exhaustive) list of relevant topics:

- Adapted notions of games for synthesis of complex interactive computational systems
- Games played on complex and infinite graphs
- Games with quantitative objectives
- Game ith incomplete information and over dynamic structures
- Heuristics for efficient game solving.

Organizers: Kim G. Larsen, Nicolas Markey, Jean-François Raskin, Wolfgang Thomas.

Workshop: Design framework concept and tool
Hristina Moneva, Teade Punter, Roelof Hamberg – ESI workshop for industry with participation from companies Océ, ASML, Philips Healthcare, and Vanderlande, Eindhoven, the Netherlands, November 11, 2011

Workshop: A Design Framework for Model-based Development of Complex Systems
Hristina Moneva, Roelof Hamberg, Teade Punter – AVICPS (Analytic Virtual Integration of Cyber-Physical Systems Workshop), Vienna, Austria, November 29, 2011

Workshop: Synchron Workshop 2010 and 2011
INRIA organized through its Aoste team the 17th edition of Synchron in Frejus. The seminar is a rather informal event, of one-week duration, meant to gather international experts together with junior researchers and PhD/postdoc students in a studious while festive atmosphere. Days are given to formal presentations, and evenings may be spent in further
talks and informal demos. In 2010 the Synchron seminar attracted over 50 participants, and acknowledged the active support of Artist-Design. The 2011 edition of Synchron has been hold in Fontainebleau in December 2011
http://www.artist-embedded.org/artist/Synchron-2010,2206.html

Tutorial Day: Formal Methods in Computer-Aided Design (FMCAD 2011)
Verimag has organised the Tutorial day of this conferences on the theory and applications of formal methods in hardware and system verification. FMCAD provides a leading forum to researchers in academia and industry for presenting and discussing ground breaking methods, technologies, theoretical results, and tools for reasoning formally about computing systems. It covers formal aspects of computer-aided system design including verification, specification, synthesis, and testing.
http://www.cs.utexas.edu/users/ragerdl/fmcad11

-- The above is new material, not present in the Y3 deliverable --

4. Internal Reviewers for this Deliverable

- Bruno Bouyssounouse (Verimag)
- Susan Graf (Verimag)