



Project no.: **ICT-FP7-STREP-214755**

Project full title: **Quantitative System Properties in Model-Driven Design**

Project Acronym: **QUASIMODO**

Deliverable no.: **D5.7**

Title of Deliverable: **Case Studies: Validation**

Contractual Date of Delivery to the CEC:	Month 24
Actual Date of Delivery to the CEC:	Month 24 (1 Feb 2010)
Organisation name of lead contractor for this deliverable:	Saarland University
Author(s):	Henrik Bohnenkamp, Holger Hermanns, Kim G. Larsen, Jean-François Raskin Jan Tretmans, Frits Vaandrager, Jiansheng Xing
Participants(s):	P01 AAU, P02 ESI, P03 CNRS, P05 SU
Work package contributing to the deliverable:	WP 5
Nature:	R+P
Version:	0.1
Total number of pages:	16
Start date of project:	1 Jan. 2008 Duration: 36 month

Project co-funded by the European Commission within the Seventh Framework Programme (2007-2013)		
Dissemination Level		
PU	Public	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

Abstract:

This deliverable summarises the Quasimodo activities on the modelling verification, analysis, and testing of the Quasimodo case studies in year 1 and 2.

Keyword list: Templates.

Contents

1	Introduction	3
2	HYDAC: Accumulator Charge Controller	5
3	CHESS: Wireless Sensing	7
3.1	Clock Synchronization Validation	7
3.2	Model-Based Testing of a Wireless Sensor Network Node	8
3.3	Quantitative Evaluation of the gMAC protocol	9
4	TERMA: Model-Based Schedulability Analysis	10
5	OCE: Adaptive Scheduling of Data Paths	12
6	ASML: Design space exploration	13
7	Conclusion	14
	Bibliography	14

1 Introduction

This deliverable reports on progress made in the modelling, verification and analysis of the QUASIMODO case studies. The description of the case studies is part of Deliverable D5.2 [11].

HYDAC. This case study is based on a product which has been developed by HYDAC, but is not yet available on the market. The problems and tasks described here for this concrete product are also easy transferable to other products, so the HYDAC has a great interest in the knowledge transfer provided by this EU-project. The product is an accumulator-charge controller (ACC) which optimizes the energy and the wear of the used components, especially the pump.

For this case, we managed to enforce safety properties in an efficient way, such that the controller consumes nearly the least possible amount of energy. This is obtained in a systematic way using a chain of automatic tools for the synthesis, verification and simulation. As we will explain below, this case study shows that our tools have reached a level of maturity that allows us to tackle interesting and relevant industrial control problems.

CHESS. The first case provided by CHESS concerns the design of a selfbalancing scooter, the Chessway. This is a meta-stable system with interesting control challenges. It exposes design problems in all engineering disciplines involved: mechanics, electronics and software development. As such, it is more than a case study. Actually, the initial problem description is rather vague, and thus more a general a design problem than a quantitative verification problem. Therefore we considered this problem not as a verification task, but as a task to be considered with respect to modelling process improvements, on which we report in Deliverable D1.3 [18]. The case is therefore not covered in this deliverable.

The second CHESS case study is about wireless sensor networks protocols, particularly the G-MAC protocol developed by CHESS. Sensors communicate wirelessly with neighboring sensors thus forming a sensor network. Sensors rely on these neighbours to forward their messages such that these messages eventually reach their destination where they can be processed.

The progress made on this case study is outstanding. Not only did we find problems in the original clock-synchronization protocol as designed by CHESS, we also managed to identify a revised algorithm that satisfies the relevant properties for the scenarios considered. Furthermore, this case is central for our test generation activities, and we studied trade-offs between energy consumption and collision probabilities resulting from different internal timing parameters of the protocol used by CHESS.

TERMA. This case study considers the ACC ASW software, a system for satellite attitude and orbit control used within the Herschel and Planck satellite systems.

The architecture of the Herschel/Planck ACC ASW reflects the aim of maximizing the commonality between the Herschel and the Planck software. The architecture is primarily Herschel-based, since it generally has more sensors and actuators than Planck does. The architecture has been described using UML.

During year 2, focus has been on schedulability analysis. For this Terma has so far performed classical worst case response time analysis. Though applicable in most case, this (overly) conservative method declares one of the configurations non-schedulable. A new (extended) methodology in which both task, *resources* and scheduling principles are modelled as timed (stop-watch) automata allows for a more precise analysis declaring all configurations schedulable. Also, the method is efficient (results obtained within 2 minutes per configuration), provides schedule visualization as a Gantt chart and is able to replicate possible deadlock violations in a simulator. In year 3, the usability take-up of the method within Terma will be further evaluated.

OCE. After year 1, we added a fifth case study, the Oce case, to our selection of cases. It is concerned with the data path of a printer/copier encompassing the complete path of the image data (the bit stream) from source (for example the network) to target (the imaging unit). This data path has to be adaptive because its properties heavily influence the image quality of the end product as well as system behavior aspects that have an eminent effect on usability. Due to its complexity, it provides an excellent challenge for the new analysis and synthesis techniques that are being developed within Quasimodo. Below we report good progress on applying our integrated methodology to this case.

Another case study which we report on has been tackled in the context of a research cooperation with the company ASML.

2 HYDAC: Accumulator Charge Controller

Participants

- Jean-François Raskin and Pierre-Alain Reynier, Centre Fédère en Vérification, Belgium
- J. J. Jessen and Kim G. Larsen, Aalborg University, Denmark
- Pierre-Alain Reynier, National Center for Scientific Research, CNRS, France
- Franck Cassez, NICTA Australia

Contribution The HYDAC case study is about controller design for a system composed of: (1) a machine which consumes oil, (2) a reservoir containing oil, (3) an accumulator containing oil and a fixed amount of gas in order to put the oil under pressure, and (4) a pump. When the system is operating, the machine consumes oil under pressure out of the accumulator. The level of the oil, and so the pressure within the accumulator (the amount of gas being constant), can be controlled using the pump to introduce additional oil in the accumulator (increasing the gas pressure). The control objective is twofold: first the level of oil into the accumulator (and so the gas pressure) can be controlled using the pump and must be maintained into a safe interval; second the controller should try to minimize the level of oil such that the accumulated energy in the system is kept minimal.

In our attack on this problem, we showed how to apply recent tools for the automatic synthesis of robust and near-optimal controllers to this real industrial case study. We employed three different classes of models and their supporting existing tools, UPPAAL-TiGA for synthesis, PHAVER [8] for verification, and SIMULINK [14] for simulation, in a complementary way.

We first developed an approach for the synthesis of a correct controller for a timed system. This controller is based on UPPAAL-TiGA applied on a very abstract untimed game model for synthesis and on SIMULINK for simulation.

To solve the HYDAC control problem, we use three complementary tools for three different purposes: UPPAAL-TiGA for synthesis, PHAVER for verification, and SIMULINK for simulation. For the synthesis phase, we show how to construct a (game) model of the case study which has the following properties: (1) it is simple enough to be solved automatically using algorithmic methods implemented in UPPAAL-TiGA; and (2) it ensures that the synthesized controllers can be easily implemented. To meet those two requirements, we consider an idealized version of the environment in which the controller is embedded, but we put additional constraints into the winning objective of the controller that ensure robustness of winning strategies. As the winning strategies are obtained in a simplified model of the system, we show how to embed automatically the synthesized strategies into a more detailed model of the environment, and how to automatically prove their correctness using the tool PHAVER for analyzing hybrid systems. While the verification model allows us to establish correctness of the controller that is obtained automatically using UPPAAL-TiGA, it does not allow us to learn its expected performance in an environment where noise is not completely antagonist but follows some probabilistic rules.

For this kind of analysis, we consider a third model of the environment and we analyze the performance of our synthesized controller using SIMULINK. To show the advantages of our approach, we compare the performances of the controller we have automatically synthesized with two other control strategies. The first control strategy is a simple two-point control strategy where the pump is turned on when the volume of oil reaches a floor value and turned off when the volume of oil reaches a ceiling value. The second control strategy is a strategy designed by the engineers at HYDAC with the help of SIMULINK.

This work is reported in [5].

Perspective The design of controllers for embedded systems is a difficult engineering task. In the case considered, we managed to enforce safety properties in an efficient way, such that they consume the nearly least possible amount of energy. We devised a systematic way to develop models and to use a chain of automatic tools for the synthesis, verification and simulation of a provably correct and near optimal controller for a real industrial equipment. We believe that this case study shows that our tools have reached a level of maturity that allows us to tackle interesting and relevant industrial control problems.

3 CHESS: Wireless Sensing

3.1 Clock Synchronization Validation

Participants

- F. Heidarian, Schmaltz, J.M. Schuts, F. W. Vaandrager, and F. Zhu, ESI: Radboud University Nijmegen, the Netherlands.

Contribution The Wireless Sensor Network (WSN) case study provided by our partner CHESS uses an epidemic (gossip) communication model. In order to meet strict energy constraints, a Time Division Multiple Access (TDMA) protocol is used. This limits the period in which nodes are active and allows them to switch to an energy saving mode for the remainder of the time. One of the greatest challenges in the design is to find suitable mechanisms for clock synchronization.

This is a main focus of our work on this case study. We developed a model of the actual clock synchronization algorithm, as it is being used in the gMAC protocol. We model this as a network of timed automata and checked instances using the Uppaal model checker. In doing so, we established that in certain cases a static, fully synchronized network may eventually become unsynchronized if the Median algorithm is used, even in a setting with infinitesimal clock drifts.

We therefore investigated a slight variations of the gMAC clock synchronization algorithm that does not have the correctness problems of the Median algorithm. Again we modelled the protocol as a network of timed automata, and verified various instances using the Uppaal model checker. Next, we developed a full parametric analysis of the protocol for the special case of cliques (networks with full connectivity), that is, we give constraints on the parameters that are both necessary and sufficient for correctness. These results have been checked using the proof assistant Isabelle. Finally, we arrived at a negative result for the special case of line topologies: for any instantiation of the parameters, the protocol will eventually fail if the network grows. This result suggests a variation of the fundamental result of Fan and Lynch on gradient clock synchronization for a setting with logical clocks whose value may also decrease.

This work is reported in [10, 13]

Perspective Analysis of clock synchronization algorithms for wireless sensor networks is an extremely challenging area for quantitative formal methods. One challenge is to come up with the right abstractions that will allow us to verify larger instances of our model. Another challenge is to make more detailed (probabilistic) models of radio communication and to apply probabilistic model checkers and specification tools such as PRISM, mcpta, or UPPAAL pro.

Our model checking results are promising since – despite the small number of nodes that can be analyzed, model checking provides valuable insight in the behavior of protocols for wireless sensor networks, insight that is complementary to what can be learned through the application of simulation and testing.

3.2 Model-Based Testing of a Wireless Sensor Network Node

Participants

- Jan Tretmans, ESI: Radboud University Nijmegen, the Netherlands

Contribution Another activity within the Wireless Sensor Network (WSN) case of Chess is model-based testing. The goal of this testing activity is to perform a protocol conformance test of the gMAC protocol layer, i.e., checking whether the gMAC implementation of a node in isolation behaves in compliance with the protocol rules in the gMAC protocol specification. Such a conformance test is important for checking the correctness of a single node, in particular, if in later stages of the Chess WSN project nodes in a single network might be supplied by different manufacturers based on the same gMAC specification.

The conformance test will be performed using model-based testing, i.e., a model of the required behaviour according to the gMAC specification is constructed and tests are automatically generated from this model and then executed on the gMAC protocol stack of a WSN node.

The goals of this activity for Quasimodo are:

- to check the quality of the Chess gMAC protocol implementation with respect to Chess' own gMAC protocol specification;
- to provide a way of assessing compliance of future third party protocol implementations with respect to the Chess gMAC protocol specification;
- to assess the real-time and data-intensive model-based testing approaches developed in Quasimodo in the context of the Chess WSN case study;
- to compare the different model-based testing tools Uppaal-Tron and TorX(akis) developed in Quasimodo;
- to investigate the difference between models used for verification (see Section 3.1) and for model-based testing;
- to investigate testability in the context of WSN: how to design and develop so that future testing and adapter development is facilitated.

We have chosen a phased approach for testing the gMAC protocol entity. In the consecutive phases, the real-time software and hardware will be tested in increasingly more realistic environments:

phase 1: testing of the gMAC software only in a simulated host environment with simulated time;

phase 2: testing of the software on the target hardware with simulated time;


phase 3: testing of the software on the target hardware in real time.

Perspective The current status is that an initial model has been developed, based on the models used for verification, but adapted to the specific use for model-based testing. First experiments have been performed in phase 1: based on the initial, abstract model tests were automatically generated with the model-based test tool TorXakis, and these tests were on-the-fly (on-line) executed on the gMAC protocol software in a host environment.

3.3 Quantitative Evaluation of the gMAC protocol

Participants: Haidi Yue, Joost-Pieter Katoen, Henrik Bohnenkamp (RWTH Aachen)

Contribution As already discussed in Deliverable 5.5 [3] we have developed a quantitative model of the gMAC protocol using the language MoDeST for which tool support is being developed within Quasimodo. The focus of our model lies on the Time Division Multiple Access (TDMA) scheduling is employed to allow multiple sensor nodes to share the same transmission medium. Collisions happen if two or more nodes send messages to another node at the same time. Due to the energy limitation, the number of active slots in TDMA should be chosen as small as possible. However, the smaller this number, the greater the probability of collisions becomes. We investigated the relation between the number of active slots, collisions, and energy consumption. The results have now been extended in two directions: Modelling of mode mobility, and analysis of energy efficiency. In particular, we investigated the efficiency of the gMAC protocol in terms of speed of information dissemination vs. energy consumption. The main result here is that there is a trade-off between the number of active slots available per frame, and the energy needed to infect every node in the network with a bit of information. This can be seen in the figure below, which shows results for a 15x15 grid network, and where the *circle-lines* show the energy consumption (right y-axis) versus the time (counted in frames), and the black, curved lines (left y-axis) show the ratio of infected nodes (i.e., nodes that have received a message) versus the number of frames. Obviously, for the chosen parameters, 7 active slots per frame are most energy efficient.



Similar results are obtained when considering the possibility that a node randomly decides whether to send in its send slot or not, and with node mobility.

This work is published in [20].

Perspective Work on this case study is ongoing. The results we have so far are based on some simplifying assumptions, especially with respect to transmission range, where a unit-disc radio model is employed. We currently investigate how more realistic radio models influence the measures of interest. We chose the physical radio model described in [9], where the signal to interference and noise ratio (SINR) is used to determine whether a node receives a message or not. Using this radio model, we consider not only static grid structures as before, but networks with random node-placement. We plan to compare the results not only with those obtained with the unit-disk model, but also with a classical slotted Aloha protocol model.

4 TERMA: Model-Based Schedulability Analysis

Participants

- Jacob Illum, Kim G. Larsen, Marius Mikucionis; Aalborg University, Denmark.
- Steen Palm; TERMA.

Contribution The Herschel/Planck mission consists of two satellites. Herschel and Planck have different scientific objectives and thus the sensor and actuator configurations differ, but both satellites share the same computational architecture. The common architecture consists of a single processor, real-time operating system (RTEMS), basic software layer (BSW) and application software (ASW).

Terma A/S has performed classical worst case response time analysis by analyzing [15] and [16] resulting in [17]. The analysis is based on classical scheduling framework [4]. The goal of this work was to show that ASW tasks and BSW services are schedulable on one processor and no deadline is violated. The framework uses preemptive fixed priority scheduler and a mixture of priority ceiling and priority inheritance protocols for resource sharing and extended deadlines (beyond period). One of the results of [17] is that the system is *not* schedulable on Herschel in event processing configuration, however it is argued that such situation has never been observed in testing and hence the result is too pessimistic.

The goal of the Quasimodo contribution [12] is to apply a model-based approach allowing schedulability analysis to be carried out as model checking. With this approach a more precise analysis is possible (and hence more optimistic but still realistic results) by making more assumptions explicit in the model with better control over task and resource modeling than the classical scheduling framework can offer.

Extending both the Times tool [2] (supporting only highest locker protocol [7]) and the UPPAAL scheduling framework [6] (support for multi-processor scheduling but no support for shared non-CPU resources) the contribution [12] develops a methodology to solve the scheduling

problem using the UPPAAL model checker with extensive application of the new feature of stop-watches (in natural modelling of preemption, more accurate estimates of worst case blocking-times and worst-case response times). Our model-based analysis is much more optimistic than the classical (conservative) scheduling method: all tasks are schedulable (at least for the first 12 cycles) and only one task is subject to blocking (in contrast to the classical method where all tasks except one is subject to blocking).

Perspective We have shown how UPPAAL model-checker can be applied for schedulability analysis of a system with single CPU, fixed priorities preemptive scheduler, mixture of periodic tasks and tasks with dependencies, and mixed resource sharing protocols. A number of issues still remain to be dealt with, including:

- Is the model fair with respect to the actual system? We will evaluate this by also applying the methodology to new missions under development.
- So far, schedulability has “only” been established for the first 12 cycles. Work on extending this to an infinite number of cycles is ongoing by identification of the hyper-period of the system, and by running UPPAAL on more powerful (cluster) platforms.

5 OCE: Adaptive Scheduling of Data Paths

Participants

- I. AlAttili, F. Houben, G. Igna, S. Michels, F. Zhu, and F. W. Vaandrager, ESI: Radboud University Nijmegen, the Netherlands

Contribution The data path of a printer/copier encompasses the complete path of the image data (the bit stream) from source (for example the network) to target (the imaging unit). In order to reach Oce's objective of genuine system adaptability, also the data path has to be adaptive because its properties heavily influence the image quality of the end product as well as system behavior aspects that have an eminent effect on usability. At run-time changes in the environment (or in the observed image quality, using a feedback mechanism) may for instance require the use of different algorithms in the data path, deadlines for completion of computations may change, new jobs may suddenly arrive, and resource availability may change. To realize this type of behavior in a predictable way is a major challenge. Currently it is already impossible to quickly evaluate cost, energy, performance aspects of various data path implementation solutions at design-time. This does not only hold for changing, adaptive functionality (steered by load, content, print quality), but even for a given fixed functionality. Partner ESI/RU is involved in a big project (named Octopus) with Oce in which Uppaal is used to make detailed models of the datapath of printer/copiers and to analyze their behavior. Due to their complexity, these models provide an excellent challenge for the new analysis and synthesis techniques that are being developed within Quasimodo.

We applied UPPAAL-TiGA to automatically compute adaptive scheduling strategies for a simplified version of the Oce case study. As far as we know, this is the first application of timed automata technology to an industrial scheduling problem with uncertainty in job arrivals. We focussed on a timed automata model reflecting uncertainty, which is due to the fact that the arrival time of new printer jobs is typically unknown. Arrival times are the most significant source of uncertainty in this application domain. For this setting, we managed to apply UPPAAL-TiGA to automatically compute adaptive optimal scheduling strategies. G. Igna constructed a detailed Uppaal model of the datapath of a new machine that is currently under development at Oce, and used Uppaal for design space exploration. A joint paper with J. Illum from AAU on the use of guided search techniques for deriving schedules for this model is currently under preparation.

This work is reported in [1].

Perspective Even though the work of G. Igna has demonstrated that it is possible to build accurate Uppaal models of the datapath of realistic printer designs, it is hard to manage the complexity of these models: the details of the Uppaal models are hard to understand for the engineers at Oce, and often a single small change in the design requires changes at many places in the Uppaal models. Therefore, we are currently working on a software tool that translates high-level models of such embedded systems, defined in terms of parameterized partial orders, to Uppaal, and that supports design space exploration based on analysis of the Uppaal models. Since many companies within the embedded systems area face similar scheduling problems as

Oce, we expect that several of these companies will be interested in such a tool once we have demonstrated that it is able to effectively support the design process at Oce. A scenario is that an external tool vendor takes over maintenance of the tool after the Octopus project has ended.

6 ASML: Design space exploration

Participants

- Rom Langerak, Jaco van de Pol, and Jiansheng Xing, ESI: Universiteit Twente, the Netherlands
- Jan Tretmans, ESI: Radboud University Nijmegen, the Netherlands

Contribution In a research cooperation with ASML on the design space exploration for motion control applications, we have considered an additional case study, which is not an official QUASIMODO case study. We studied a multi-processor platforms where processors are interconnected by Rapid Input-Output (RIO) packet switches. Motion control applications are characterized by feedback/feed forward controllers and periodic execution. The main challenge is how to map a specific application on the platform such that periodic timing constraints (all packets arrive their destinations before the periods) are met.

We have constructed a system model for analyzing packet latencies, according to a given mapping and then end-to-end latencies are analyzed; second, all possible mappings are explored for finding the optimum one that satisfy the timing constraints. We first used the modelling language POOSL, a system-level description language that supports such method in constructing and/or refining models. Based on the developed models, POOSL supports functional and performance analysis with simulation based techniques and thus significantly reduces the risk of expensive design-implementation iterations.

From the POOSL model, we obtained approximate results for worst-case latencies and average-case latencies. However, as motion control applications are safety-critical and time-critical, worst-case latencies are strict timing constraints. Formal verification of worst-case latencies as well as functional logics is demanded. Motivated by this requirement, we transformed the model to a network of timed automata models. We were able to study worst-case latencies and functional logics can be verified with UPPAAL model. We then investigated a method to record time (clock value) is proposed such that approximate worst-case latencies can be analyzed by simulation in UPPAAL. We are currently exploring if this approach is also applicable to best-case latencies and average-case latencies, and we are comparing with the results obtained via POOSL.

This work is reported in [19].

Perspective The deficiency of our approach is that UPPAAL cannot handle large models and is simply not made for simulation-style analysis. We are nevertheless going to focus on larger case and see if more abstractions can be made for the UPPAAL model. We anticipate that by focusing

on a specific performance metric such as worst-case latencies and by omitting non-related details larger models for realistic applications can be handled.

7 Conclusion

The case studies considered in Quasimodo are deliberately chosen to be very diverse. This enables us to study strengths and weaknesses of our approach, and to identify attack points for further technological improvements.

The CHES gMAC protocol is a very fruitful and inspiring research case. Not only did the tools we develop help us to identify problems in the original clock synchronisation setup, and suggest ways to overcome this. We also are applying model based testing of implementation against its specification in an exemplary way.

For the TERMA case study, we managed to devise a methodology that solves intricate scheduling problems in the presence of preemption, accurate estimates of worst case blocking-times and worst-case response times, arriving at schedules of better quality than what classical approaches deliver.

The successes on the HYDAC case study shows that our tools have reached a level of maturity that allows us to tackle interesting and relevant industrial control problems. This is achieved by a skillful combination of analysis and simulation tool, and as such constitutes a blueprint for further successful industrial applications.

The Oce data path case helped us identify the need for enhanced tool support to handle the model creation process. Therefore, we are working on a software tool that translates high-level models of Oce-like embedded systems, defined in terms of parameterized partial orders, to Uppaal, and that supports design space exploration based on analysis of the Uppaal models. The need for support in design space exploration is also apparent in the ASML case study we reported on.

Bibliography

- [1] I. AlAttili, F. Houben, G. Igna, S. Michels, F. Zhu, and F. W. Vaandrager. Adaptive scheduling of data paths using Uppaal Tiga. In S. Andova et.al, editor, *Proceedings First Workshop on Quantitative Formal Methods: Theory and Applications (QFM'09)*, volume 13 of *Electronic Proceedings in Theoretical Computer Science*, pages 1–12, 2009.
- [2] Tobias Amnell, Elena Fersman, Leonid Mokrushin, Paul Pettersson, and Wang Yi. Times - a tool for modelling and implementation of embedded systems. In *TACAS '02: Proceedings of the 8th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 460–464, London, UK, 2002. Springer-Verlag.
- [3] Henrik Bohnenkamp, Joost-Pieter Katoen, Kai Mittermüller, Holger Hermanns, Julien Schmaltz, Faranak Heydarian, Frank Cassez, Haidi Yue. Quasimodo Deliverable D5.5. *Case studies: Models*. Public Document. January, 2009.

- [4] Alan Burns. Preemptive priority based scheduling: An appropriate engineering approach. In *Principles of Real-Time Systems*, pages 225–248. Prentice Hall, 1994.
- [5] Franck Cassez, J. J. Jessen, Kim G. Larsen, Jean-François Raskin, and Pierre-Alain Reynier. Robust and optimal controllers - an industrial case study. In *Proceedings of HSCC'09*, 2009.
- [6] Alexandre David, Jacob I. Rasmussen, Kim G. Larsen, and Arne Skou. *Model-based Framework for Schedulability Analysis Using UPPAAL 4.1*. Taylor and Francis, 2009. To appear in CRC Press Book on "Model-Based Design of Heterogeneous Embedded Systems".
- [7] Elena Fersman. *A generic approach to schedulability analysis of real-time systems*. Acta Universitatis Upsaliensis, 2003.
- [8] G. Frehse. Phaver: Algorithmic verification of hybrid systems past hytech. *Int. Journal on Software Tools for Technology Transfer (STTT)*, 1(1–2), 2008.
- [9] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Trans. Inf. Theory*, 46(2), March 2000.
- [10] F. Heidarian, J. Schmaltz, and F. W. Vaandrager. Analysis of a clock synchronization protocol for wireless sensor networks. In A. Cavalcanti and D. Dams, editors, *Proceedings 16th International Symposium of Formal Methods (FM2009), Eindhoven, the Netherlands, November 2-6, 2009*, volume 5850 of *Lecture Notes in Computer Science*, pages 516–531. Springer, 2009.
- [11] Holger Hermanns, Kai Sven Mittermüller, Teun van Kuppeveld, Jan Storbak Pedersen, Poul Hougaard. Quasimodo Deliverable D5.2. *Preliminary descriptions of case studies*. Public Document. July, 2008.
- [12] Kim G. Larsen, Marius Mikučionis, Brian Nielsen, Steen Palm, Jacob I. Rasmussen. Quasimodo Deliverable D5.7b. *Model-Based Approach for Schedulability Analysis*. Confidential Document. January, 2010.
- [13] M. Schuts, F. Zhu, F. Heidarian, and F. W. Vaandrager. Modelling clock synchronization in the Chess gMAC WSN protocol. In S. Andova et al., editor, *Proceedings Workshop on Quantitative Formal Methods: Theory and Applications (QFM'09)*, volume 13 of *Electronic Proceedings in Theoretical Computer Science*, pages 41–54, 2009.
- [14] Simulink, 2008. <http://www.mathworks.com/products/simulink/>.
- [15] Terma A/S. Herschel-planck acms acc asw requirements specification. Technical report, Terma A/S, Issue 4/0.
- [16] Terma A/S. Software timing and sizing budgets. Technical report, Terma A/S, Issue 9.

- [17] Steen Palm. Herschel/plank acc asw: Sizing, timing and schedulability analysis. Technical report, Terma A/S, 2006.
- [18] Frits Vaandrager. Quasimodo Deliverable D 1.3. *Modelling Process Improvement*. Public Document. January, 2010.
- [19] J. Xing, R. Langerak, J. van de Pol, J. Tretmans, J.P.M. Voeten and B.D. Theelen. Performance analysis of POOSL Model using UPPAAL. In preparation 2010.
- [20] Haidi Yue, Henrik Bohnenkamp, and Joost-Pieter Katoen. Analyzing energy consumption in a gossiping mac protocol. In *Proc. MMB '01*, LNCS. Springer, 2010. To appear.