

Five Years of Multi-Disciplinary Academic and Industrial Research: Lessons Learned¹

Gerrit Muller and Maurice Heemels

Eindhoven University of Technology, Embedded Systems Institute
P.O. Box 513, 5600 MB Eindhoven
gerrit.muller@esi.nl maurice.heemels@esi.nl

Abstract

The Boderc project was the first broad scale multi-disciplinary project with closely cooperating industrial and academic participants performed by the Embedded Systems Institute. After five years we evaluated the project and collected the lessons learned from this project. We look at the underlying project philosophy, Industry-as-Laboratory, at process and organization issues, and at the project results.

The article describes the industrial context, the design of the system (a high-volume printer), and the challenges of creating system designs in the industrial context. We discuss the research propositions to tackle these challenges. The central theme of the research is the use of multi-disciplinary models to predict and analyze system performance and to explore design options.

Boderc Project Introduction

Boderc Project. Early on in the Boderc project, the goal was defined as shown in Figure 1. The goal of Boderc is to develop a model-based methodology that supports multi-disciplinary design (space exploration) by predicting system performance. The developed models, methods and techniques should in particular be applicable in the early design phases and

must satisfy industrial application constraints. They should be usable in the industrial context with its particular people, processes and economic constraints related to design time, effort and costs. Moreover, the economic constraints and the traditional processes of the manufacturer of the product restrict the design space a priori by posing constraints on the design. Most parts in a new design will not be revolutionary, existing solutions and technologies and way of working will be re-used. The methodology should be effective for this constrained design space.

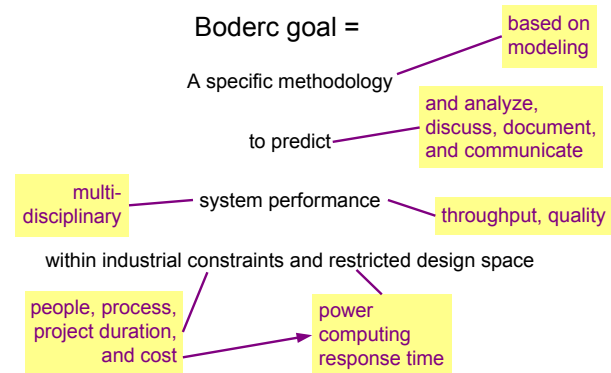


Figure 1. Boderc Goal

During the Boderc project the awareness emerged that it is not only about *predicting* system performance. The methodology and models force to make design choices quantitative and explicit which enables the analysis of various design options, communication between engineers from different disciplines and to commence the

¹ This work has been carried out as part of the Boderc project under the responsibility of the Embedded Systems Institute. This project is partially supported by the Netherlands Ministry of Economic Affairs under the Senter TS program.

design with all disciplines involved in the beginning of a project. Also modeling of (parts of) the system increases the understanding and insight in the design. All these factors lead to shorter design iterations and more confidence in the consequences of design choices. In the end, better products are delivered faster.

Industry-as-Laboratory. The Boderc project uses the *industry-as-laboratory* approach, as proposed by Colin Potts (Potts 93) and visualized in Figure 2.

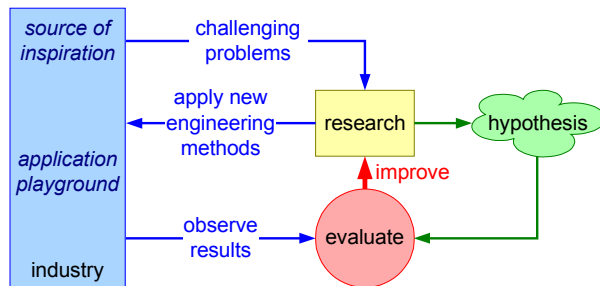


Figure 2. Industry-as-laboratory

The industry-as-laboratory approach exploits the actual industrial setting as a test environment, which warrants that the research question is based on real industrial problems. The Boderc research team, consisting of a mix of academic and industrial people, investigates a new product engineering methodology. A research hypothesis is formulated on the new methodology. The methodology is applied in the industrial setting and the results of these experiments are observed and used to evaluate the hypothesis. Coupled to the multi-disciplinary design problems for high-tech systems discussed in the beginning of this paper, the research hypothesis of the Boderc project was chosen as:

The product creation lead time will be reduced significantly by the use of multi-disciplinary models during the early product development phases.

The term *Carrying Industrial Partner (CIP)* is used for the company that provides the problem and the industrial setting. The CIP of Boderc is *Océ Technologies B.V.*,

which creates high-volume document printing systems.

The industrial context. One of the product families that is designed by Océ is a range of high-volume printers and printers, see Figure 3.



Figure 3. The Domain: Printers by Océ

The application context is best characterized by document printing systems that are highly productive, reliable, and user-friendly. These systems can print on several sizes of media, different weights, automatically on both sides and include stapling, booklet production, or other types of finishing. In order to be perceived as reliable devices, such printers must be very robust with respect to variations in media. As the printing speed is rather high (typically 1-2 images per second), timing requirements are tight and advanced mechatronics are indispensable. This indicates that variations in timing parameters that relate to paper and image transport must be controlled up to a high degree. This becomes the more apparent if one realizes that the positioning of images on paper has tolerances well below 1 mm.

When considering the embedded control of these systems, one should think of controlling multiple sheets that travel the paper path simultaneously and synchronizing this sheet flow with the imaging process. In Figure 4, an overview of a printer is presented. When the printer is in normal operation, a sheet is separated from the trays in the paper input module (PIM), after which it is sent to the paper path that transports the sheets accurately in the direction of the print engine, where the image is fused on a sheet of paper. After that, the sheet is turned for duplex

printing, or transported by the paper path to the finisher.

Multi-disciplinary methods. The Boderc research falls typically within the category of multi-disciplinary design methods as opposed to the more conventional mono-disciplinary research areas like mechanical, electrical or software engineering. The latter research fields are relatively mature. Some bi-disciplinary approaches exist, for instance hybrid systems theory (van der Schaft99) that combine continuous dynamical models (using e.g. differential equations) typically describing the physical part of a high-tech machine and discrete models (e.g. finite state machines or automata) to describe the software behavior. The hybrid field is relatively immature and many issues are at present unsolved (at least at the large-scale needed for industrial usefulness). However, the industrial need for analysis / synthesis methods for high-tech machines in which this ‘hybrid interaction’ plays an important role, will stimulate the research in this domain over the years to come.

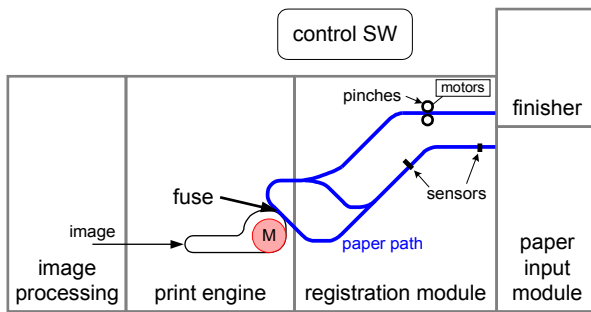


Figure 4. Overview of a printer

The translation of system requirements to detailed mono-disciplinary design decisions spans many orders of magnitude. The few statements of performance, cost and size in the system requirements specification ultimately result in millions of details in the technical product description: million(s) of lines of code, connections, and parts. Figure 5 shows this dynamic range as a pyramid with the system at the top and the millions of technical details at the bottom.

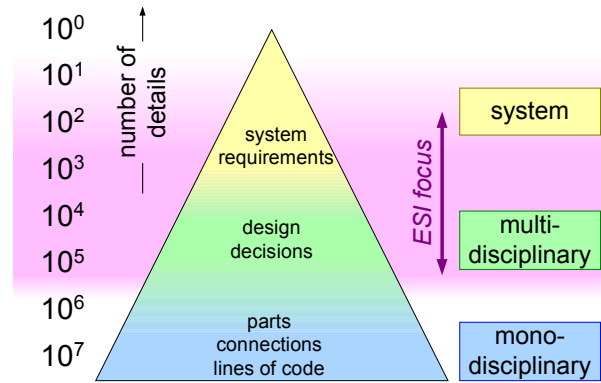


Figure 5. Exponential pyramid

The methodologies to be established by ESI, including the Boderc results, address the multi-disciplinary area and aim at coupling the academic research to industrial practice. It is the area of translating hundreds of system level requirements into tens of thousands of design choices.

Project results

In this section we summarize the results and position them in the design pyramid, see Figure 6. The numbers refer to the chapters of the Boderc symposium book (Heemels06). It indicates that the results are reasonably well distributed over the different abstraction levels. Most PhD-theses are connected to the existing scientific body of knowledge, a level of detail that goes into more detail than necessary in an industrial context due to the current academic standards. However, the continuous pull towards multi-disciplinary knowledge has resulted in several theses that range from detailed scientific up to a certain level of multi-disciplinary design.

The system-level reasoning used in the Boderc project was bundled in the Boderc method, that consists of a high-level framework, where more specific plug-ins are used to make it concrete and practical.

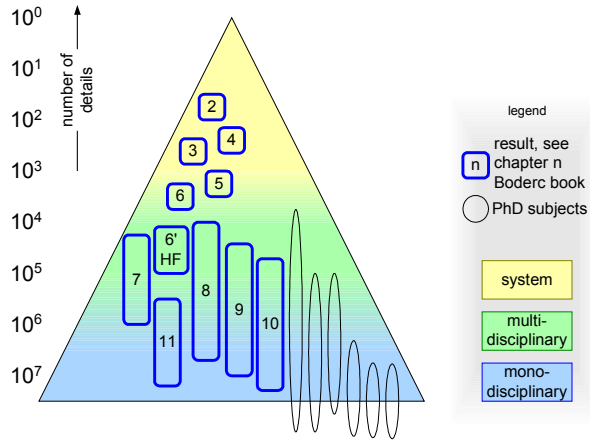


Figure 6. Boderc results

Submethods. Boderc explored a few submethods as system level plug-ins: the key drivers technique, threads-of-reasoning, and budget-based design (Chapters 3, 4, and 5 of the Boderc book, respectively). The key driver models for past as well as for future projects were highly appreciated by the industrial partner. The key driver method couples the main customer objectives to the technical requirements for the system and provides overview in the relationships between them. The submethod of threads-of-reasoning was used internally in the project, to relate industrial needs to (potential) research questions and modeling efforts. The value of these threads is the positioning of work and the relation between a local exploration and the more global context. Budget-based design was used mainly for power considerations for a printer. However, we derived general guidelines (a method) on how to setup budgets and use them in a supportive manner for design purposes.

System-level models. The industrial appreciation of research results is a source of inspiration for further research, as can be seen by the results on kinematic modeling (Happy Flow) as discussed in Chapter 6. To learn from this successful industrial model, we identified the success factors of this particular model in Chapter 6. This should form a stepping stone to arrive at clear guidelines on how to set up effective models in an industrial

context. In other domains with similar kinematic problems, like in mailing systems, there is already a strong interest in the particular model. The success of (the type of models as) the Happy Flow model created a demand for developing a similar type of models for thermo- and power-modeling (Chapter 7), which is an ongoing activity within Océ. Other ‘system-level’ models are also considered in Chapter 8 and 9. Chapter 8 focussed on how to evaluate the overall control architecture in terms of response times, CPU load, etc. Chapter 9 described models that are related to printing quality. New printer technologies were assessed via ‘virtual’ printer models with respect to their printing quality.

Detailed modeling. The study of stepper motors in Chapter 10 has a somewhat less system-level flavor as the before mentioned plug-ins. Océ had important reasons to replace the DC motors by stepper motors. For this purpose, Chapter 10 investigates the possibilities and impossibilities of stepper motors and aims at building a profound understanding of stepper motors that lead to practical design rules.

Also the work of the PhD students was stretched more to the multi-disciplinary design domain (see Figure 6) than in conventional research at universities.

- Chapter 12 provided an overview of techniques for state-of-the-art performance analysis for embedded real-time system architectures. Based on these experiences, an indication was given which method is used best under which circumstances to successfully support the decision making process for the architecture.
- Chapter 13 presented a model-driven design approach for real-time systems. This approach enables the analysis of real-time systems and allows automatic software code generation from the model that preserves the properties analyzed in the model.

- Chapter 14 takes a control engineering view on the controlled system and reduces the real-time software behavior to a model consisting of a varying time-delay. This chapter proposed analysis methods and techniques for the synthesis of controllers that are robust against these time-varying delays (i.e. jitter and latencies caused by computation and communication).
- For the control design of the drives of the paper transport system, Chapter 15 proposed a hierarchical control paradigm based on supervisory control is proposed. A systematic analysis and design procedure based on low-level controllers for the motors in combination with high-level sheet control was proposed.
- Chapter 16 described the design and application of event-driven control, which allows for a varying sample time in controllers. Event-driven control can have major benefits with respect to resource utilization like processor and communication load, while still maintaining a good control performance. See also (Sandee06).
- In Chapter 17, a systematic design trajectory was proposed for the combination of real-time controllers and physical / mechanical processes. A design path was indicated in which stepwise the original (simulation) models of both plant and controller are replaced by their real implementations.

Chapter 11 discussed ways to simulate real-time embedded software together with its environment, being of a physical / mechanical nature. One approach, Software-in the-Loop, is now used at Océ as a way to early test the functionality of paper path control software. This leads to faster feedback and design cycles and therefore better products.

The above indicates that several activities were carried out that connect more detailed knowledge (mono-disciplinary models) with multi-disciplinary design choices (system

level models). This is indispensable for the design process as outlined in the overall Boderc method. One specific example ('textbook example') was already discussed in Chapter 2. In the above mentioned work, successful multi-disciplinary results were achieved, based on a more detailed understanding. Primary value of these activities is to enable the multi-disciplinary reasoning, without the need to cope continuously with all details.

Industry-as-Laboratory research approach

The intention of the industry-as-laboratory approach is twofold:

- to better connect academic research to industrial needs and to focus on results with industrial feasibility.
- to unfreeze industrial participants from constraints imposed by their current context and current knowledge and to be perceptive for unconventional techniques.

The reward for this investment is that academic researchers obtained triggers for new, industrially relevant research directions and that industrial engineers were stimulated to try out multi-disciplinary models and design methods in actual development projects. Examples of the improved connection are event-driven control design in Chapter 16 and the evaluation of embedded systems architecture in Chapter 12 to mention just two. Examples of the unfreezing include the tool coupling ideas as described in Chapter 11, the use of key drivers in Chapter 3, budget-based design in Chapter 5 and many others. Especially, the fact that researchers at Océ could work in printer development projects without having to contribute directly was very beneficial. These researchers obtained the exploratory space to try new techniques and methods on actual industrial problems without the tight time-to-market

constraints that the developers themselves are faced with.

The research created useful industrial *models* on one hand, and did benefit the advancement of multi-disciplinary *methods* on the other. As an example, the success of the Happy Flow model (Chapter 6) had a direct effect on reducing the effort and time needed to design the paper transport system and the print job scheduling. At the methodological level, Happy Flow was used to identify properties that effective industrial design models should satisfy. From these properties, guidelines can be derived on how to build successful industrial models. Also the making of actual budget models in the project (Chapter 5) was successful in itself, but resulted also in a more methodological view upon the use of budgets.

Along the lines of the industry-as-laboratory research approach, we will evaluate the original research hypothesis against the findings in the Boderc project. The research hypothesis of Boderc was formulated as

The product creation lead time will be reduced significantly by the use of multi-disciplinary models during the early product development phases.

The question arises whether or not the research hypothesis is true and if it is true, what possible evidence is brought by the results of the Boderc project, as described in this book. A very strong ‘true’ can be given, even if we only focus on the one of the Boderc modeling activities: the Happy Flow model. Initial experience with the Happy Flow model showed significant savings in product creation lead time. On top of this reduction in the product creating lead time by Happy Flow, we believe that the use of the other Boderc models, like the virtual printer models, the heat flow modeling, the investigation in stepper motors, the evaluation of embedded system architectures, to mention a few, reduce the product creation time even further. Computations were not made to assert the

economical value of these and other Boderc modeling activities. However, considering the broad use of the models within Océ, we conjecture that they must have a positive effect on the reduction of the product creating time, as otherwise developers and engineers would not have embraced them.

As the developed models predict the performance and consequences of specific design choices more accurately than previous state-of-the-practice models, uncertainty and risks are reduced for later stages. This means that less conservative designs become feasible resulting in better products. For instance, the Happy Flow enabled a better prediction of the paper transport systems and as such smaller printers could be built.

Lessons learned in process and organization

Of course, the development of model-based design methodologies for high-tech systems cannot be solved by one project like Boderc. Boderc made one proposal for a design methodology based on the experience obtained. Although a first step has been made, additional projects are needed to do *research* on methods. These additional projects must apply the researched methodology in different settings, and re-evaluate the hypothesis. The industry-as-laboratory approach has a long term character:

- Each industrial application family requires significant time and effort to understand the necessary domain-specific knowledge.
- Multiple industrial applications are required to support methodological conclusions.

For the benefit of future large-scale industrial research projects, we will collect our lessons learned in the Boderc project. This is especially important as Boderc is innovative in the process model that it uses for performing research.

Tension between mono-disciplinary academia and multi-disciplinary industry.

The tension in this type of project is between the need for depth for mono-disciplinary academic partners and the need for short term industrially applicable and multi-disciplinary results of the industrial partners. The tension is most severe for students pursuing their PhD degree, as they are typically defending it within mono-disciplinary faculties. As a consequence, this tension is visible in the positioning of the subjects of the PhD-students, as shown in Figure 6. The required scientific depth pulls the students downward into the mono-disciplinary field. However, as can be observed there are some PhD results that stretches over several orders in the design pyramid. This is a clear benefit of a project like Boderc: the eye towards industrial applicability and system-level design is more profound in the Boderc (sub)projects than in the traditional research at universities. However, towards the end of the project the PhD students retracted more and more towards their own individual work on the PhD thesis, which is understandable on one hand, but caused some disintegration of the project team on the other.

To value system-level research more at PhD level, an opportunity lies in creating the possibility of receiving a PhD degree in 'multi-disciplinary or system engineering schools' that go beyond the traditional engineering faculties as often encountered at universities.

Duration of the project. If we consider the development of the project members from mono-disciplinary towards multi-disciplinary, then we see that we needed at least two years for this growth. When we started we expected that this growth would take only one year. This means that we needed more time for the total project than the 4 years as originally planned. After two learning years at least two years of exploration and application are needed, followed again by at least one year of

consolidation. A total project duration of 5 to 6 years would solve this problem, at least if we target for the original level of multidisciplinary methods. However, this clashes with the need for short-term usable results as is often desirable from an industrial point of view.

Multi-disciplinary curriculum. Another solution to reduce the long learning phase could be the educational part of the PhD students. The first year was typically filled with mono-disciplinary classes within their own domain as this is customary for the PhD students in general. For future projects we recommend to create a multi-disciplinary curriculum for the PhD-students working in ESI projects. This would give the project members basic knowledge of other design disciplines. As a consequence, we expect that they (better) oversee consequences of design choices for other disciplines. Building a common multi-disciplinary device in the first year would also be a good means to learn cross-disciplinary thinking. The purpose of such a curriculum is twofold. First, a faster learning curve in the multi-disciplinary industrial setting and secondly, scientific results that fit higher in the design pyramid of Figure 5. Ideally we would like PhD students with a T-shaped thesis: sufficient depth in the mono-discipline, the vertical part of the T, connected to the multi-disciplinary problem, the horizontal ledger of the T.

Clear initial problem statement. Another remedy for the long learning phase is to have a clear problem statement at the beginning of a project. In the beginning of the Boderc project we started with mainly a collection of industrial problems that were faced during the final integration of a high-tech system, where the (sub)designs of the disciplines meet. We still had to extract the problem statement approach from these symptoms. We anticipated that the integration problems were caused by design decisions in the early design phases of which the consequences were not

considered thoroughly across disciplines. From that we inferred the problem statement. Particularly in large-scale research projects, we recommend to prepare a sharp problem statement and approach before the project has even started.

We observed that about 18 months after the project start the PhD subjects were fixed. However, 2 years after the project staff several highly attractive PhD subjects surfaced, e.g. modeling of system impact of the use of stepper motors and thermo modeling. These topics would have been closer to system level and multi-disciplinary reasoning than the more deeper and mono-disciplinary topics that were chosen now. A clear problem statement at the beginning of the project might have helped to identify the PhD subjects earlier and it might have helped to select the right mix of disciplines. Projects started at ESI since 2005 perform a more extensive problem analysis before starting the project. It is too early to evaluate the effectiveness of this measure.

Project team composition. Another point of discussion is the team composition. For instance, one could question if PhD students are in the right phase of their personal development curve to do this type of research. They have the requirement to write a PhD thesis. From an academic point of view, PhD students are desirable as it forms one of the foundations of academic groups. But the requirement of developing sufficient novel contributions in a mono-disciplinary area forms an obstacle in obtaining system-level design techniques and models. A better balance could be obtained by involving more postdocs in the research as a remedy. The PhD students in the project were either *dynamics* and *control* oriented, or *software* and *digital electronics* oriented.

The members of the carrying industrial partner (CIP, the industrial partner that indicated the research problems) were typically young researchers, which had not yet developed their system engineering or system

architecting skills extensively at the beginning of the project. Also they still had to explore the application field of printers and printers. This resulted in the fact that domain specific knowledge was not readily available. More experienced engineers, instead of young researchers, is a solution although it is harder to unfreeze them from their project duties. Although it seems a high investment for industry to make their key engineers available for research projects, we believe that in the long run this would be very beneficial for all parties involved including themselves. Typically the first industrial Boderc workshop in which the CIP developers with more system overview were present, resulted in sharper discussion that arrived easier at the essence of the industrial design problem. A benefit for industry is that the young researchers were confronted with academic thinking, system level reasoning and industrial practice. These assets make them very valuable for the industry.

The non-CIP industrial people had more industrial and system-level experience. They turned out to be catalysts in the process of the project (especially in the beginning). The non-CIP industrial people are typically allocated to the project for two days per week. They found it hard to contribute in their part-time allocation. Part of the available time is needed for communication and recapturing what the other project members have been doing during their absence. The time left is not sufficient to actually build models. The project could benefit more from the existing industrial knowhow if these industrial participants would also be full-time available. In hindsight we might have created a more balanced team in terms of experience by replacing one or two PhD students by post docs, but also getting more senior CIP people in the project (at more days per week).

Summary lessons learned. In summary, the lessons learned with respect to process and organization:

- Even more attention is needed for the composition of the project team, in the balance experience-inexperienced, in the balance industrial-academic and in the balance mono-disciplinary and multi-disciplinary.
- The industrial problem is rather broad and also the original project goal was not really crystallized at the start of the project. This hampered the fast start of the project. When the goal has to be discussed in the beginning of the project, it is better to let the PhD students start later.
- Also the research topics of PhD students should be clear at the start of the project and most importantly, should match the overall research goal. In the first year of Boderc the PhD topics were selected, while we believe that in year two we were better prepared to make the selection. As PhD students form a major part of the work force and should take care of the momentum in the project, it is very important that they contribute directly to the overall project goal.
- The project team was too *dynamics* and *control engineering* oriented due to an unclear initial problem formulation. It is desirable to have more disciplines in the project team to arrive at a better balance.
- Part-time people can only be effective in a coaching role. The real research work (exploration, application, and consolidation) requires full-time people.
- Communication across disciplinary boundaries is really very difficult, as experienced throughout the project.
- It is very beneficial to provide plans for the classes that the PhD students attend in the first year. These classes should fit the overall multi-disciplinary project problem. In particular, some basic classes with respect to the specific application domain and classes outside the student's own discipline are considered valuable.
- The mix of project members in disciplines and background in the first year was a good preparation for the first industrial Boderc workshop. A critical success factor of this workshop was the presence of CIP engineers with system level overview. The participants were able to iterate between system requirements and mono-disciplinary design choices.

Concluding remarks

Of course, a pioneering project like Boderc is based on the 'spirit of an entrepreneur.' With good faith we started the project. Since there was little to no experience with research projects of this size and type, the Boderc project had a high learning character. As such, it is inevitable that there is room for improvement in the process and organizational part. Most importantly, we have to learn from this experience for the future.

The next generation of projects of the Embedded Systems Institute already benefits from the lessons learned in Boderc, which underlines the innovative nature of the Boderc project. Although the Boderc project may have suffered a bit from its pioneering position, we can be very satisfied with the outcomes, as described in the previous sections of this article.

If one takes a 'business-oriented view' on the Boderc project, one can say that it generated return on investment for the involved companies and academic groups that are clearly above expectation. For the academic groups this was typically realized via graduations of master and PhD students, published papers and the Boderc impact on their curricula. ESI was able to attract 3 research fellows for its staff via the Boderc project and Boderc helped to put ESI on the world map as a leading centre in the area of embedded system engineering. Océ saved a lot of design effort and time in current and future


projects due to the development of many valuable models, techniques and methods.

Using a more ‘soft view,’ the Boderc project created a lot of *awareness* both within academia and industry with mutual understanding and respect for individual positions, capabilities and strengths. The difficulties in multi-disciplinary and system-level design became more explicit and as such created a first step in addressing them. Academics were confronted with industrial needs, while industry learned to untangle itself occasionally from the time pressure present in product development projects. This led to the development of models, techniques and methods that were truly relevant in industrial practice. It also initiated the cooperation of all disciplines, already from the earliest phases of design projects. Of course, the derived methodology, understanding and models should be refined further and validated in future projects. But all aspects taken into account, the Boderc project made an excellent first step in developing model-based methodologies for designing high-tech systems.

References

- Maurice Heemels and Gerrit Muller (editors). “Boderc: Model-based design of high-tech systems; A collaborative research project for multi-disciplinary design analysis of high-tech systems.” Symposium book, Embedded Systems Institute, 2006.
- Colin Potts. “Software-engineering research revisited”. *IEEE Software*, Vol. 10, No. 5:19–28, September/October 1993.
- Heico Sandee. “Event-Driven Control in Theory and Practice”. PhD thesis TU/e, 2006.
- A.J. van der Schaft and J.M. Schumacher. An introduction to hybrid dynamical systems. *Lecture Notes in Control and Information Sciences*, Vol. 251, 1999.

Biography



Maurice Heemels (born 1972) received the M.Sc. degree (with honours) in Mathematics and the Ph.D. degree with the highest distinction from the department of Electrical Engineering (EE) of the TU/e (Technische Universiteit Eindhoven) in 1995 and 1999, respectively. He received the ASML-award for the best Ph.D. thesis of the TU/e in the period 1999/2000.

From 2000 until 2004 he worked as assistant professor in the control systems group (EE) of the TU/e. In 2001 he worked three months as a visiting professor at the ETH in Zurich, Switzerland.

Since 2004 he is working at ESI (Embedded Systems Institute) as research fellow. In 2004 he spent three months at the R&D lab of the printer manufacturer Océ in Venlo (The Netherlands).

Since June 2006 he is affiliated as associated professor at the Control Systems Technology group of the Mechanical Engineering department of the TU/e. His research interests include modeling, analysis and control of hybrid and non-smooth dynamical systems.



Gerrit Muller received his Master’s degree in Physics from the University of Amsterdam in 1979. He worked from 1980 until 1997 at Philips Medical Systems as system architect. From 1997 to 1999 he was manager System Engineering at ASML. From 1999 - 2002 he worked at Philips Research. Since

2003 he is working as senior research fellow at ESI (Embedded Systems Institute). In June 2004 he received his doctorate. The main focus of his work at ESI is on System Architecture methods and on education of future System Architects. Special areas of interest are:

- Ways to cope with the exponential growth of size and complexity of systems. Examples of methods to address the growing complexity are product lines and composable architectures.
- The human aspects of systems architecting (which in itself is a crucial factor in coping with the above mentioned growth).

More information (System Architecture articles, course material, curriculum vitae) can be found at: Gaudí systems architecting <http://www.gaudisite.nl/>