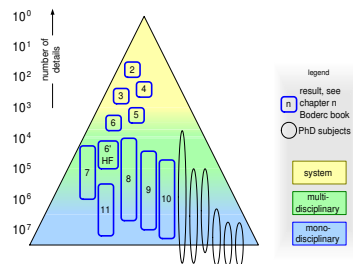


Five Years of Multi-Disciplinary Academic and Industrial Research; Lessons Learned

-



Maurice Heemels and Gerrit Muller

Embedded Systems Institute

Den Dolech 2 (Laplace Building 0.10) P.O. Box 513, 5600 MB Eindhoven The Netherlands

maurice.heemels@esi.nl gerrit.muller@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.

Distribution

This article or presentation is written as part of the Gaudí project. The Gaudí project philosophy is to improve by obtaining frequent feedback. Frequent feedback is pursued by an open creation process. This document is published as intermediate or nearly mature version to get feedback. Further distribution is allowed as long as the document remains complete and unchanged.

All Gaudí documents are available at:
<http://www.gaudisite.nl/>

version: 1.0

status: preliminary draft

February 10, 2011

1 Introduction

The design of high-tech mechatronic systems like wafer steppers, electron microscopes, copiers, etc. is a complex process. Multiple “classical” engineering disciplines need to make the overall design in close co-operation. Typically, electrical, mechanical, software and other engineering disciplines together determine the functioning of the final product. Especially in the early design phases, the design of the product is vulnerable for erroneous design choices as many others will be based on it subsequently. These erroneous choices tend to show up in later phases during the integration or even the manufacturing itself. Late in a project, the “repairs” are more difficult and can lead to a much longer development period than planned and/or a less optimal product.

The main reasons for non-optimal design choices are illustrated in Figure 1 and are summarized below.

- A common language and background between multiple engineering disciplines is lacking that enables the reasoning about system properties. As a consequence, the effects of choices, made by one discipline, cannot be overseen for another discipline; wrong assumptions are made on what other disciplines guarantee on their subdesign; confusions and misunderstanding are present about definitions of specific terms and priorities differ over disciplines.
- Many design choices are made in an implicit way, based on on experience, intuition and “gut-feeling.” That way, it is hard to communicate the reasons and to discuss the design or particular choices in it. Decisions are sometimes not well-founded by quantitative arguments, but can be forced by seniority and “shouting loudest.”
- Especially dynamic, time depending aspects of a system are complex to oversee. There are not many tools and methods available to support the time varying aspects in a design, in contrast to many static or steady-state aspects.
- Out-of-phase project evolution is another important factor. A typical example of the latter is that the mechanical design often precedes the electronic design, which on its turn precedes the software design.

The effects of the above design complications are amplified by the size and complexity of high-tech machines (typically millions lines of codes, tens of thousands mechanical components like pinches, springs, belts, motors, bolts, etc.) The more complex the machine and the more people involved in the design, the stronger the effects. Of course, the four mentioned reasons are not the only ones that complicate the design. Other factors like organizational or political, or geographically scattering of the design team contribute too. Those latter issues are important

Many multi-disciplinary problems in product development

- Mechanical engineering precedes
- Electronics engineering precedes
- Software engineering

Most of the problems show up late in engineering and in the integration phase

- For instance mechatronics assumes 1 ms response
- Software promises 10 ms response

Lack of systematic approaches to detect / solve these problems in early phases

- Lots of tuning, trial and error
- Unpredictable project timing and costs

Figure 1: Typical Industrial Problems in Mechatronics Systems

as well, but are of a different dimension and more related to business management. We believe that the aforementioned reasons can be relieved by the use of models that capture the system behavior and a reasoning method that indicates how and when to use them. That is why the Bode-RC project was initiated by the Embedded Systems Institute, Océ Technologies, Imtech, Chess and 6 academic groups of the universities of Eindhoven, Nijmegen and Twente.

2 The Boderc Project

Early on in the Boderc project, the goal was defined as shown in annotated form in Figure 2.

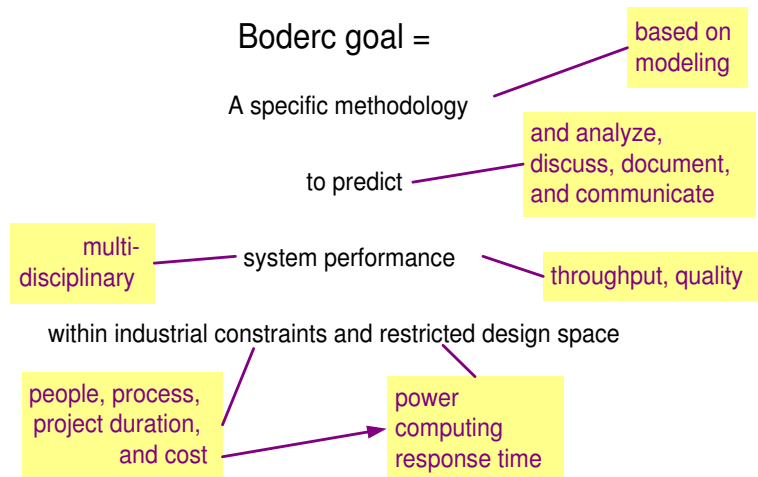


Figure 2: Boderc Research Project Goal

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, which constraints the design space. The methodology should be effective for this constrained design space.

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 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.

2.1 Research Method: Industry-as-Laboratory

The Boderc project uses the *industry-as-laboratory* approach, as proposed by Colin Potts [3] and visualized in Figure 3.

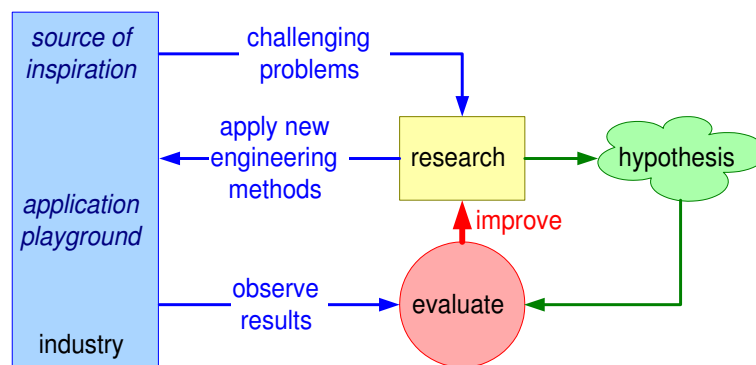


Figure 3: Industry as Laboratory: Research of Engineering Methods

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 chapter, 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, which creates high-volume document printing systems.

2.2 The industrial context

One of the product families that is designed by Océ technologies is a range of high-volume printers and copiers, see Figure 4.



Figure 4: The Domain: Printers and Copiers 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 copiers must be very robust with respect to variations in media. As the printing speed is rather high (typically above 1 image 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 5 overviews of a copier are presented. When the copier 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.

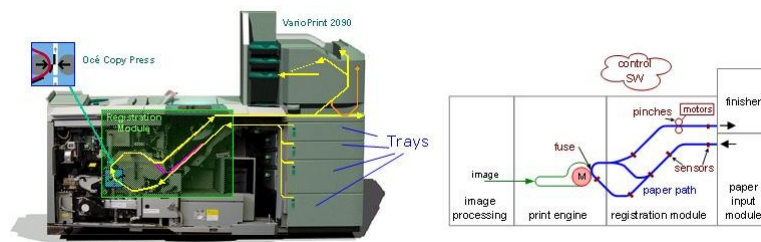


Figure 5: Illustration and schematic picture of a copier

The creation of a new copier is taken as the carrier for the Boderc project.

3 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, although some doubts exist about the maturity of software engineering [2]. Some bi-disciplinary approaches exist, for instance hybrid systems theory [5] 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.

Researchers in the mono-disciplinary areas are used to well-defined problems that can be researched in depth with solutions most often based on mathematical rigor. A lot of uncertainty pops up when we move to multi-disciplinary problem solving. The problem itself is only partially defined, while at the solution side different formalisms have to interoperate, such as discrete (software) and continuous (mechanical) models. Figure 6 shows a categorization of the design methods with as vertical axis the degree of multi-disciplinary interaction. The form of the method is an indication how well the method is defined and how much uncertainty is left.

In the industrial context the *system* level is often relatively well-defined in a systems requirement specification. Such a specification describes the functionality of the system and quantifies the main performance characteristics. The translation of these requirements into mono-disciplinary design choices, however, is still full of uncertainty. A lot of uncertainty is caused by the many (dependent and interfering) design dimensions that have to be managed at the same time. In Figure 6 the methods at this level are called *multi-objective design methods*.

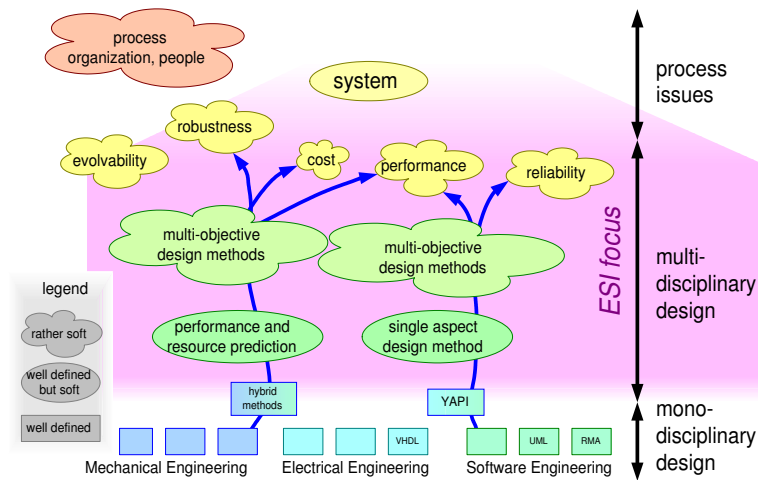


Figure 6: From Mono-Disciplinary to System Design

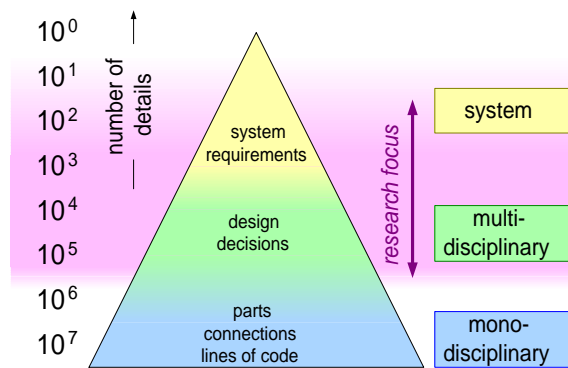


Figure 7: Exponential Pyramid

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 7 shows this dynamic range as a pyramid with the system at the top and the millions of technical details at the bottom.

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. In Figure 6 this is the range from *single aspect* to *multi-objective* design methods. In the pyramid, Figure 7, it is the area of translating hundreds of system level requirements into tens of thousands of design choices.

3.1 The Embedded Systems Institute

Boderc is the first project in a long line of ESI projects within the field of multi-disciplinary creation methods. This is a young research field, which is called *embedded systems engineering* by ESI. The existing scientific disciplines have little experience in this field, most experience can be found in industry.

The mission of ESI is *to advance industrial innovation and academic excellence in embedded systems engineering (ESE) with its vision to create and apply together with its partners world-class ESE methods*. The developed methodologies must support all aspects of the creation: specification, design, integration, test and validation.

4 Lessons Learned

The explorative discussions for the Boderc Project took place in 2002 and the project formally started also in 2002. The practical full blown take-off of the project was in early 2003, due to the time needed to fill all project vacancies (mainly the PhD positions). This means that we look back at five years in total on the first ESI project from preparation until final symposium.

We will look at several different aspects, to wit

- the research approach: industry-as-laboratory
- process, organization and people issues
- the research results

4.1 Industry-as-Laboratory Research Approach

The intent of the industry-as-laboratory approach, as also discussed in Section 1 is to address the gap between industry and academics. We have had to work hard at two sides of this gap:

- to connect academic researchers to industrial needs and to focus on results with industrial feasibility.
- to unfreeze industrial participants from contemporary constraints and to be perceptive for unconventional techniques.

The reward for this investment is that academic researchers get triggers for new industrially relevant research directions (e.g. event-driven control design and supervisory control to mention just two), and that industrial engineers are stimulated to try out multi-disciplinary design methods in actual development projects (e.g. kinematic modeling, the use of key drivers, budget-based design, et cetera) The combination

of industrially inspired academic research and the try-out of multi-disciplinary methods did benefit the advancement of multi-disciplinary methods.

4.2 Project Organization and Process

This section describes the process and evolution of the Boderc project. The master plan for the Boderc project is shown in Figure 8. The first year was used to explore and to learn from existing systems. In The second and third years the focus shifted to the creation of new products. The last year was required for the consolidation.

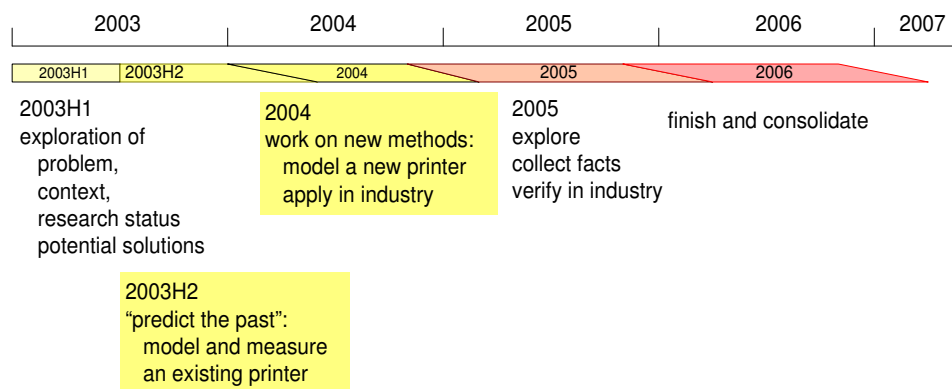


Figure 8: Master plan Boderc Project

In Figure 9 we superimposed several Boderc project activities on the design pyramid presented in Section 1, annotated with the project timing.

4.2.1 2003: Exploration and predicting the past

The activities start at a very detailed level. In a first workshop the researchers of different disciplines showed the other researchers their modeling tools applied on a small part of the paper path. The workshop after half a year of multi-disciplinary models shifted the focus slightly upwards. The main lesson learned for the project members was that communicating across disciplinary boundaries is really difficult. The actual industrial problem of barriers between the design disciplines had been repeated on a small scale. This experience was of great value for the remainder of the project as the project members were confronted with one of the main difficulties in multi-disciplinary design in industrial practice. The final models at the end of the first year of the project were more or less at the level of connected mono-disciplinary models.

During this first year of the project the team worked on an existing product, which was in the lab of the mechanical engineering department of the Technische Universiteit Eindhoven. This phase was called "predicting the past" and had the

advantage that the models could be validated against the real system behavior. A benefit of modeling an existing machine is that the target is rather stable; the design and its main parameters do not change suddenly. In the early design phases of a new product this is, of course, not the case. So, typically this was a learning phase

At the end of this first exploratory phase, the PhD students fixed their research topics:

- Supervisory control
- Event-driven control[4]
- Networked control systems
- Performance prediction of digital control architectures
- Hardware/software modeling
- Hardware/software in the loop.

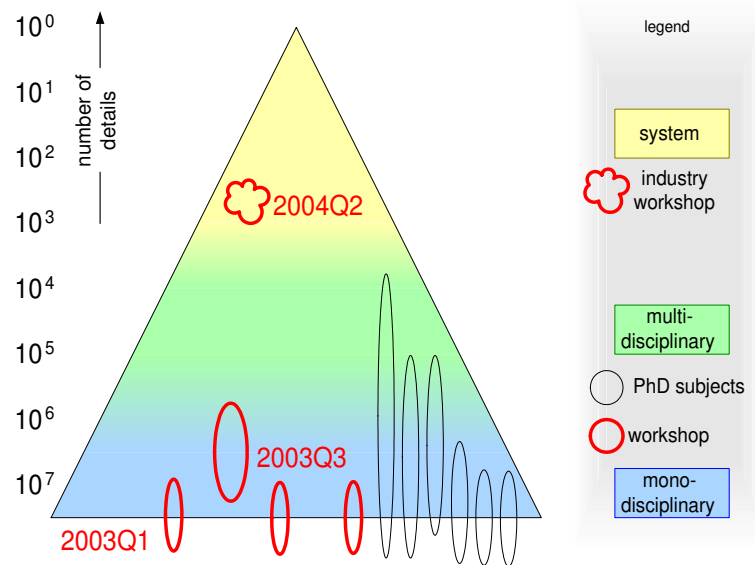


Figure 9: Positioning of workshops and PhD-work in the level of abstraction pyramid

4.2.2 2004: Predicting the future

The second project year introduces more real-world uncertainty by cooperating with the developers of a new product at Océ. This new product is more or less

defined at the top level (cost, price, size, performance, and introduction date), but all of these targets may change due to feasibility or market issues. The kick-off of this project phase, in the industry workshop, was on purpose focused at the system level, with heavy participation of the carrying industrial partner. Eight system issues, such as cost, performance, size, weight, and power, were discussed at the level of hundreds of design parameters. This sandwich approach supported the Boderc project members to create true multi-disciplinary models in the second half of 2004. Three main topics were identified at this workshop, which would be used in the next phase of Boderc to “predict the future.” The selected topics were power and space, throughput and response and the control architecture as shown in Figure 10. This workshop was a good demonstration of the skills obtained by the Boderc members in the first year as sharp and relevant discussions were held with the product developers and revealing important design choices and tradeoffs.

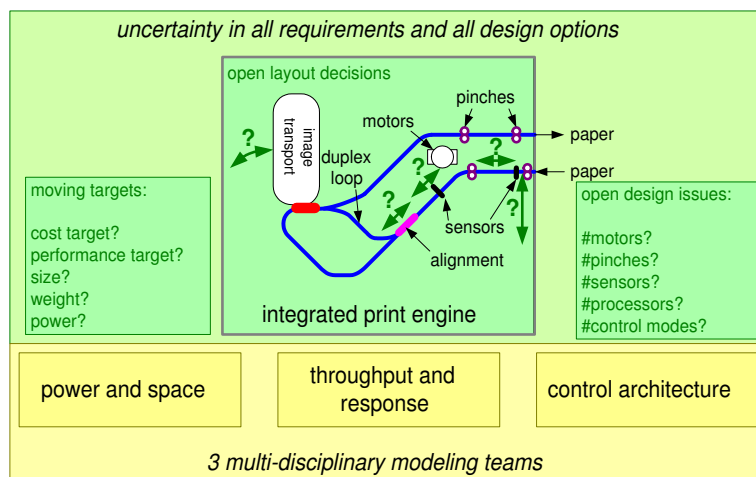


Figure 10: Modeling the Future

4.2.3 2005: Predicting the future and individual freedom

The modeling teams were transformed into aspect teams in which the way of cooperation was less rigid. *Modeling* teams of the previous phase worked weakly on fixed days, with around 5 people together at ESI during longer periods (several months), while *aspect* teams had different formats like sprints (one week working fully on one aspect) or regular meetings (e.g. four half day meetings spread over two months only to synchronize the individual activities). Also the group sizes differed and some aspect teams consisted of just one individual. The previous modeling team *power and space* transformed into a team working on the usage of stepper motors (instead of DC motors) in the copier paper path. As mentioned before, the *control architecture* team remained active as well, partly because two

PhD students dedicated their thesis on this topic, and partly, because this topic remained directly relevant to the development of the new copier at Océ. Comparison of various model-based methods for performance prediction were compared via an industrial case study of a car-navigation system. The kinematic model supporting the scheduling of sheets and images and the design of the mechanical layout grew mature in this phase. Various other smaller topics (including the individual PhD work) were pursued in this period.

This phase can be characterized by individual freedom. The PhD students felt the time-pressure to be productive with respect to their PhD thesis, that had to be finished within a year. At this point the team work was mainly abandoned in favor of the success of the PhD work. The work on the digital control architecture was however continued within a team. Also this team followed a course together entitled “Execution Architectures,” which gave an incentive to their team work. They learned the state-of-practice in this domain and their follow up work was related to bridging the gap between the academic research and industrial practice.

In this stage also a workshop was organized such that the team members, being submerged in their individual work, were stimulated to take a system and industrial application perspective on their current work. This helped many academic people to see how their work can be applied and validated in industrial practice. As a result, event-driven control was applied in the control of the image part of the copier. Cost price reasons urged the use of low-resolution encoders for which conventional control algorithms fail. Event-driven control fits this challenging problem. Also the work on supervisory control, which was applied experimentally on a small paper path at first, was scaled up to a real paper path of a complete copier.

In this phase also work on power and heat flows was started up as it was crucial in the development for the new copier. Models were (and are still being built) with characteristics of the kinematic model, as this was one of the most industrially successful models created within Boderc. An additional quality attribute related to printing accuracy was taken up as well. Via virtual copier models new copier concepts are tested out with respect to printing quality.

Also the Boderc methodology crystallized in this phase. The methodology gives place to activities like model(ing formalism)s, methods, techniques and tools is. This framework supports the making of sound design tradeoffs and the investment of time and effort only on essential choices and not on wasting design effort on non-issues. This ‘reasoning backbone’ uses as plug-ins the more specific in-depth models, techniques and tools for dedicated questions in the design.

4.2.4 2006: Consolidate

The divergence of the project in the previous phases, that was necessary to guarantee the success of the PhD students, changed again in this phase. The reason was the

delivery of the final overall results of the Boderc project in one common book and final symposium. The work was oriented mainly on consolidation. Research on, for instance, the virtual printer models and the power and heat flow in the copier continued. Also the success factors of the kinematic model were identified in this phase, which might lead to guidelines how to build successful industrial models.

4.3 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 good 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 is a long term investment:

- Each industrial application requires significant time and effort to understand the industrial setting (acquiring the domain specific knowledge).
- Many industrial applications are required to support methodological conclusions.

The tension in this type of project is between the need for depth for mono-disciplinary project partners and the need for short term industrially applicable and multi-disciplinary results of the industrial partners. The tension is most severe as students pursuing their PhD degree are typically defending it within mono-disciplinary faculties. As a consequence, this tension is most visible in the positioning of the subjects of the PhD-students, as shown in Figure 9. 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. An opportunity lies here 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 seen at universities. The Schaeffer school of Engineering (Steven’s, Hoboken, NJ, USA) is a good example in this. 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 disintegration of the project team.

If we look at the development of the project members, especially the PhD students, 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 need 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 multi-disciplinary methods. However, another solution could be found in the educational part of the PhD students. The first year was typically filled with mono-disciplinary knowledge in 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 in ESI projects. Project members that have basic knowledge of other design disciplines are expected to oversee consequences of design choices in a broader perspective. Building a common multi-disciplinary device in the first year would also be a good means to learn the cross-disciplinary thinking. The purpose of such a curriculum is twofold: a faster learning curve in the industrial setting and a scientific result that fits somewhat higher in the level of abstraction pyramid. 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.

The CIP members were typically young engineers, which have not yet developed their system engineering or system architecting skills extensively. Also they had to explore the field of printers and copiers implying that domain specific knowledge was not available as one would expect with more experienced people. Typically the industry workshop 2004 Q2 in which the CIP developers with more system overview were present, showed that with this knowledge present discussions are more to the point and arriving easier at the essences of the design problem. However, the education of the young CIP-members and their growth in both academic and system level thinking on one hand and the industrial experience gained makes them very valuable for the industry. Also the non-CIP people had more industrial and system-level experience, which turned out to be a catalyst in the process. The non-CIP industrial people are typically allocated to the project for two days per week. They find it difficult 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. The time left is not sufficient to actually build models. The project could benefit more from the existing industrial know-how 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.

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 and in the balance industrial-academic.
- The industrial problem is rather broad and also the original project goal was not really crystallized yet. This hampered the fast start of the project. If still the goal has to be discussed in the beginning of the project, it is better to let

the PhD students start later. Also in the Boderc case, the Project team was too *control* oriented due to an unclear main objective. It is desirable to have more disciplines in the project team.

- 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 in the entire project.
- The mix of project members in disciplines and background in the first year was a good preparation for the industry workshop in Q2 2004. 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 disciplinary design choices.

4.4 Project Results

The results of the Boderc Project have been bundled in a symposium Book **cite to be added**. The research results of the project are positioned again in the design pyramid, see Figure 11. It shows that the ordering of the Boderc symposium book is top-down. It also shows that the results are reasonably well distributed over the different abstraction levels.

The contents of PhD-theses is also shown in this figure. Most PhD-theses are connected to the existing scientific body of knowledge, a level of detail that goes beyond the bottom of the pyramid. As a consequence, the PhD theses are well founded, with mathematical rigor. However, the continuous pull towards multi-disciplinary knowledge has resulted in several theses that range from detailed scientific up to multi-disciplinary design.

The system-level reasoning used in the Boderc project was bundled in the Boderc methodology, that consists of a high-level framework, where more specific plug-ins are used to make it practical. Boderc explored a few system level plug-ins: key-drivers technique, threads-of-reasoning, and budget-based design. The key driver models for the past as well as for the future were highly appreciated by the industrial partner. The industry-as-laboratory approach facilitates this cross fertilization; this in contrast with conventional research projects that are never confronted with customers or markets. Threads-of-reasoning were 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. Also some PhD students used this approach to explore and clarify their individual research goals and the tradeoffs they investigate. Also the particular results of PhD students can be seen as plug-ins in the high-level method.

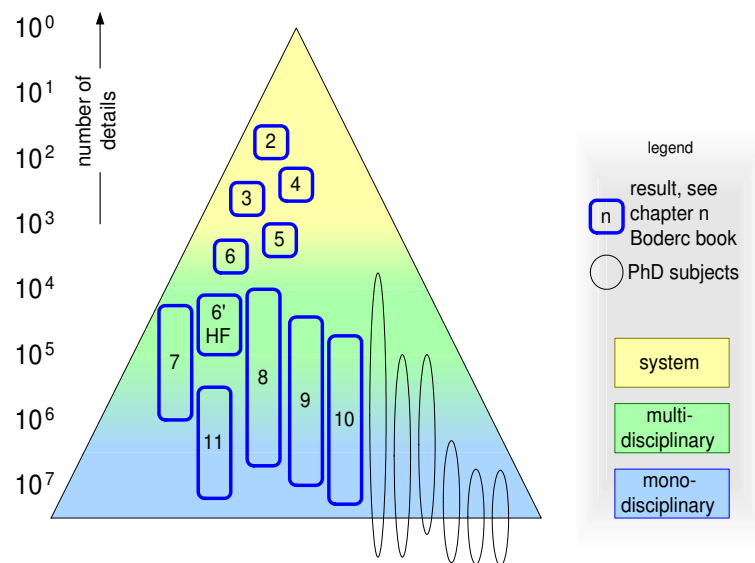


Figure 11: The project results positioned in the level of abstraction pyramid

The industrial appreciation of research results is a source of inspiration for further research, as can be seen by the results on kinematic modeling. The kinematic model was a big industrial success and from a financial point of view this model in itself saved more money than invested in the whole Boderc project. The reason is that it saves at least one complete design cycle for developing a new printer. To learn from this successful industrial model, we identified the success factors of this particular model. This should be 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. Also from the kinematic model a desire was risen to develop similar type of models for power and thermo-modeling.

Several activities have been carried out that connect more detailed knowledge (mono-disciplinary models) with multi-disciplinary design choices (system level models), as is indispensable for the design as outlined in the overall Boderc design methodology.

- micro-benchmarking and performance modeling
- power and heat modeling: both budgets and dynamical models are used.
- stepper motor characterization and analysis
- virtual printer models for print quality
- trade-offs in event-driven control

In all cases 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. Extracting the generic steps from these experiences resulted in the high-level method.

References

- [1] Gerrit Muller. The system architecture homepage. <http://www.gaudisite.nl/index.html>, 1999.
- [2] David L. Parnas. Software engineering: An unconsummated marriage. *Communications of the ACM*, page 128, September 1997. This article can also be found in "Software Fundamentals, Collected Papers by David Parnas", Addison-Wesley.
- [3] Colin Potts. Software-engineering research revisited. *IEEE Software*, Vol. 10, No. 5:19–28, September/October 1993.
- [4] J.H. Sandee. *Event-driven control in theory and practice - trade-offs in software and control performance*. Eindhoven University of Technology, Netherlands, 2006. Ph.D. thesis.
- [5] A.J. van der Schaft and J.M. Schumacher. An introduction to hybrid dynamical systems. *Lecture Notes in Control and Information Sciences*, page Vol. 251, 1999.

History

Version: 1.0, date: October 4, 2006 changed by: Gerrit Muller

- completely rewritten version, based on Chapters 2 and 18 of the Boderc Symposium book written mostly by Maurice Heemels.

- changed status to preliminary draft

Version: 0.2, date: July 19, 2006 changed by: Gerrit Muller

- changed 5 into Five in title
- added T-shaped thesis
- changed thesis into theses when plural

Version: 0.1, date: July 18, 2006 changed by: Gerrit Muller

- Added broader set of disciplines required
- added inter-disciplinary communication difficult, but achieved

Version: 0, date: July 8, 2006 changed by: Gerrit Muller

- Created, no changelog yet