
Computer Science Aalborg University

Research Evaluation
1996–2000

Department of Computer Science
Fredrik Bajers Vej 7E • DK 9220 Aalborg Ø • Denmark

Computer Science, Aalborg University—Research Evaluation 1996–2000

Edited by Luca Aceto

Cover design by Alberto Bongi

Cover reproduction by Handelstrykkeriet, Aalborg

Printed by Handelstrykkeriet, Aalborg

Copyright ©2001 by Department of Computer Science, Aalborg University. All rights reserved.

R-01--5002

ISSN 1601--0590

Preface

It is the policy of the Faculty of Engineering and Science at Aalborg University that its research units be evaluated every five years. Following that policy, the present report documents the third research evaluation of the Computer Science Department and covers the period 1996–2000.

According to the Faculty’s guidelines, the objectives of the evaluation are threefold. First, the evaluation should assess whether there is a satisfactory agreement between allocated internal and external research resources and the research accomplished. Second, the evaluation should determine whether the Department’s actual research efforts are consistent with its goals and plans for the period, as they are expressed in the long-term research plans of the Faculty of Engineering and Science. Third, the evaluation is intended to aid the Department in planning its future research efforts and its research organization.

Realizing that a research evaluation itself is a learning experience that provides an opportunity to rise above the daily routine and reflect upon research, we agreed upon an additional objective. Within the Department, the evaluation—the process as well as the final report—should constructively aid the staff members in evaluating and improving their effectiveness as researchers, group leaders, and administrators of research at the departmental level. In particular, the evaluation presents an opportunity to assess the effectiveness of departmental policy in promoting high quality research.

The research evaluation process that has led to the production of this document began in April 2000, when the members of the Department discussed the structure of the process. Among the early, important tasks, a four-member evaluation committee was appointed. Professor Finn V. Jensen accepted to be the local member of the committee (being local, he would be without a vote in the committee), and we began the search for external members. The Department was delighted when Søren Damgaard (IBM Denmark), Professor Matthias Jarke (RWTH Aachen, Germany), and Professor Moshe Vardi (Rice University, U.S.A.) agreed to be on the committee. We felt that this committee of recognized and experienced senior computer scientists from academia and industry would be able to cover all the rather diverse research areas in the Department and provide an insightful and wide-ranging evaluation. (Short curricula vitae of the members of the evaluation committee may be found in Chapter 2 of this document.)

Next, the general structure of the evaluation process and of the report were agreed upon. The structure of the evaluation report was mostly based on that of its predecessor. The main novelty in the general structure of the evaluation process was the internal evaluation of drafts of the departmental research units' contributions to the evaluation report. This process of internal evaluation went on as follows. The report of each research unit was critically read by one member of each of the other three research units. The senior members of the research unit under evaluation then discussed the unit's draft with the internal reviewers. This internal evaluation process was useful, as it invariably raised questions about the research approach, results and organization of the research units that were not addressed in the original drafts. Based on the outcome of these internal evaluations, the research units then proceeded to produce the final versions of their contributions to the preliminary evaluation booklet. I was given the responsibility of steering the process, coordinating with Finn V. Jensen when appropriate, and of editing the research units' contributions. By early December 2000, a 214-page preliminary version of the research description (on which Part II of this report is based) was completed and sent to the evaluation committee members.

In January 2001, a two-day research evaluation seminar was held at a conference facility. All research staff, including Ph.D. students, long-term visitors, and secretarial staff from the Department were invited. On the first day, the Department's research units gave oral presentations of their research, allowing the committee and other attendees to ask questions on both the presentations and the written contributions submitted by the units. During the afternoon of the first day, and part of the morning of the second, the evaluation committee worked in isolation (apart from a discussion with the leaders of the research units in which the overall organization of the Department was clarified to the committee members), and the attendees took part in various activities related to the general theme of research evaluation. Afterwards, the committee presented its evaluation to the Department and answered questions.

After the seminar, the evaluation committee finalized its written evaluation, which forms Chapter 3 of this document. In parallel, the research units produced the final versions of their contributions, generally expanding somewhat the substance of their preliminary descriptions, based on the input from the seminar. It is those slightly updated descriptions that, possibly subjected to some minor editing, appear as the chapters in Part II of this report.

Part I of the report is devoted to the actual evaluation of the research carried out at the Department. It provides a brief description of the formal context of the evaluation, short curricula vitae of the members of the committee, and then presents the committee's evaluation. As mentioned above, Part II has a chapter for the Department in general and for each of the research units. It covers the background and organizational context of the Department and the overall organization of the Department's research. Moreover, it presents the research plan for 1996–2000 and the resources available for implementing that

plan. Finally, it evaluates the global administration and organization of research and presents a plan for the period 2001–2005. The chapters for the research units survey their research activities in the period 1996–2000, and describe their organization and staff. Descriptions of research collaborations, Ph.D. projects, and service and research-related activities then follow. Finally, self-evaluations and plans for the next period are given. Part III provides a brief summary in Danish of the report.

Designing the research evaluation and producing this report was intrinsically a consensual and collective effort involving to varying degrees all the researchers in the Department. I thank them all for their assistance and support. However, Ivan Aaen, Erik Ernst, Henrik E. Jensen, Uffe Kjærulf and Peter Axel Nielsen deserve special thanks for their work as editors of the contributions by each of the research units and of the Department as a whole. Christian S. Jensen, who edited the previous evaluation report, made my editorial work much easier by providing essential information on the structure of the previous report. Hans Hüttel translated the executive summary of the committee evaluation into Danish. Mikkel Christiansen converted the figures in Chapter 4 into encapsulated PostScript. The cover page of the report was designed by architect and designer Alberto Bongi.

Any infelicity in this report is solely the editor's responsibility.

Aalborg, Denmark
June, 2001

LUCA ACETO

Contents

I	The Evaluation and its Framework	1
1	Evaluation Guidelines	2
2	The Evaluation Committee	4
3	Evaluation Report	6
3.1	Evaluation Process	7
3.2	Evaluation of the Units	7
3.2.1	Databases and Programming Technologies	7
3.2.2	Decision Support Systems	8
3.2.3	Distributed Systems and Semantics	10
3.2.4	Information Systems	11
3.3	Evaluation of the Departmental Strategy	12
3.3.1	Management of Growth	12
3.3.2	Management	14
3.3.3	Hiring	15
3.3.4	Cooperation with Industry	16
3.4	Suggestions for Future Review Processes	16
II	The Unit's Research Report	18
4	The Department of Computer Science	19
4.1	Introduction	19
4.2	Background, Conditions, and Resources	19
4.2.1	Computer Science at Aalborg University	20
4.2.2	A Brief History of the Department	20
4.2.3	Research Funding	21
4.3	The Department's Organisation	23
4.3.1	Research Units	24
4.3.2	Research-Based Education	24
4.3.3	Ph.D. Studies	25
4.3.4	Collaboration with IT Industry: Nouhauz	25

4.4	Own Evaluation of the Department	26
4.4.1	The Department's Research Plan for 1996-2000	26
4.5	The Department's Plan for 2001-2005	27
4.5.1	Strategy	27
4.5.2	Policy Making and Deployment	28
5	Database and Programming Technologies	29
5.1	Profile of the Unit	29
5.2	Activities and Results	30
5.2.1	Research Results	30
5.2.2	Framework Projects	33
5.2.3	Specific Projects	38
5.3	Organization and Staff	42
5.4	Collaborations	48
5.4.1	Collaborations with Academia	48
5.4.2	Collaborations with State Agencies	51
5.4.3	Industry Collaborations	51
5.5	Ph.D. Projects	52
5.5.1	Project Descriptions	53
5.6	Services	62
5.6.1	Program Committees	62
5.6.2	Conference Related Services	65
5.6.3	Editorial Services	67
5.6.4	Reviewing	67
5.6.5	Other Academic Services	68
5.7	The Unit's Own Evaluation	72
5.7.1	Database Systems	72
5.7.2	Programming Systems	74
5.8	Research Plan 2001-2005	75
5.8.1	Research Profile	75
5.8.2	Mobile Internet Services	76
5.8.3	Non-Umbrella Projects	78
5.8.4	Additional Considerations	80
5.9	Bibliography	82
6	Decision Support Systems	101
6.1	Profile of the Unit	101
6.2	Activities and Results	102
6.2.1	Research Results	102
6.2.2	Projects	106
6.2.3	Patents and Products	107
6.2.4	External Funding	108
6.3	Organization and Staff	108

6.3.1	Staff	108
6.3.2	Research organization	110
6.4	Collaborations	111
6.5	Ph.D. Projects	114
6.5.1	Overview	114
6.5.2	Project Descriptions	114
6.6	Services	118
6.7	The Unit's Own Evaluation	121
6.8	Research Plan 2001–2005	123
6.9	Bibliography	125
7	Distributed Systems and Semantics	133
7.1	Profile of the Unit	133
7.2	Activities and Results	134
7.2.1	Research Results	134
7.2.2	Projects	148
7.3	Organization and Staff	153
7.3.1	Staff	153
7.3.2	Research Organization	159
7.4	Collaborations	160
7.5	Ph.D. Projects	164
7.5.1	Overview	164
7.5.2	Project Descriptions	164
7.6	Services	171
7.6.1	Editorial Services	171
7.6.2	Conference and Workshop Organization	172
7.6.3	Programme Committees	172
7.6.4	Invited Presentations	174
7.6.5	Evaluation Committees	175
7.7	The Unit's Own Evaluation	177
7.8	Research Plan 2001–2005	179
7.8.1	Research Plans for the Existing Activities	179
7.8.2	A Strategy for the Unit	185
7.9	Bibliography	187
8	Information Systems	201
8.1	Profile of the Unit	201
8.2	Activities and Results	202
8.3	Organization and Staff	207
8.3.1	Staff	207
8.4	Collaborations	212
8.5	Ph.D. Projects	213
8.5.1	Overview	213

8.5.2	Project Descriptions	213
8.6	Services	222
8.6.1	Editorial Boards	222
8.6.2	Keynote Speaker	223
8.6.3	Scandinavian Research Community	223
8.6.4	IFIP Working Group 8.2	223
8.6.5	Program Committees	223
8.6.6	Organising Committees	224
8.6.7	Reviewing	225
8.7	Own Evaluation	226
8.8	Research Plan 2001–2005	227
8.9	Bibliography	229
III	Summary in Danish	240
9	Resumé	241
9.1	Formål	241
9.2	Evalueringsprocessen	242
9.3	Denne Rapport	243
10	Evalueringsudvalgets Konklusion	244

Part I

The Evaluation and its Framework

Chapter 1

Evaluation Guidelines

The guidelines from the Faculty of Engineering and Science that constitute the formal context for the research evaluation are paraphrased below.

1. The evaluation is conducted for each research unit separately and covers the research unit's work since the previous evaluation.
2. The evaluation of each research unit is conducted every five years. The specific time is set by the unit's institute.
3. The evaluation is conducted by an evaluation panel of normally three persons, with at most one from the research unit and not less than one from outside the university. Normally, at least one member must be employed by an international institution. A panel member from the unit has no vote on the panel. The panel is appointed by the Faculty following the recommendation of the research unit.
4. The evaluation includes the following:
 - (a) Published research results; (b) other research results, such as results yet to be published and results that are unpublished in a traditional sense, but which are disseminated to industry, public institutions and organizations, etc. through curricula or otherwise; (c) research activities in progress; (d) other research activities, including international research collaboration, participation in scientific congresses, conferences, and symposia, editorship, refereeing, and evaluation of theses and applicants for academic positions; and (e) research collaboration.
5. The evaluation has these objectives.

First it aims at assessing whether there is a satisfactory agreement between allocated internal and external research resources and the research accomplished. Second, it must determine whether the research unit's actual research efforts are consistent with the unit's goals and plans for the

period as they are expressed in the long-term research plans of the Faculty of Engineering and Science. Third, the evaluation is intended to aid the research unit in planning its future research efforts and its organization of research.

6. The evaluation panel prepares a preliminary evaluation that is made available to and is discussed with the unit's staff. On this basis, the final report is prepared and submitted to the Faculty with comments from the research unit and the Institute.

The evaluation report may be written in English or Danish. If it is written in English, a summary in Danish must be included. Conversely, if it is written in Danish, a summary in English must be included.

7. The research unit organizes and conducts its research evaluation while adhering to the above general guidelines, so that the evaluation fits best with the wishes and activities of the unit.

Realizing that a research evaluation is a learning experience that provides a welcome opportunity to rise above the daily routine and reflect upon research, the Department of Computer Science agreed upon an additional objective: Within the Department, the evaluation—the process as well as the final report—should constructively aid the staff members in evaluating and improving their effectiveness as researchers, group leaders, and administrators of research at the departmental level. In particular, the evaluation presents an opportunity to evaluate the effectiveness of departmental policy in promoting high quality research.

Chapter 2

The Evaluation Committee

Søren Damgaard

Søren Damgaard is EMEA SMB Operations Manager, IBM Global Services. He studied Physics and Mathematics at Aarhus University, where he obtained his doctorate in Physics in 1981. In 1987 he took a Master in Business Economics at the Copenhagen Business School. Søren Damgaard has worked both in academia and industry. He was associate professor at Aarhus University in 1979–1980, and research fellow at CERN, Geneva, in 1980–1982. Since 1982 he has been mostly holding managerial positions at IBM Denmark. Søren Damgaard has held several appointed roles, mostly as chairman and member of Danish boards for Natural Science Education. Most notably, since 1997 he has been Chairman of the Advisory Board for Natural Science Education (Ministry of Education, Denmark), and since 2001 he is Chairman of the Research Council, Danish Industry Group ITEK.

Matthias Jarke

Matthias Jarke is professor of Information Systems at Aachen University of Technology (RWTH Aachen) and executive director of the Institute of Applied IT at the GMD National Labs in Birlinghoven, Germany. Jarke studied business administration and computer science and took his doctorate in business informatics at the University of Hamburg. Prior to joining Aachen in 1991, he held faculty positions at New York University and the University of Passau, Germany. In his research, he investigates information systems support for cooperative processes in engineering, business, and culture. He has served as coordinator of three Esprit projects and is involved in two national centers of excellence; the German government appointed him to the strategy group defining the German IT research strategy for the years 2002–2006. Jarke served as program or conference chair of several international conferences and is Editor-in-Chief of Europe's oldest

database journal, Information Systems. In 1999, he was elected Vice President of the German computer society, GI.

Moshe Y. Vardi

Moshe Y. Vardi is the Karen Ostrum George Professor of Computational Engineering and Chair of Computer Science at Rice University. Prior to joining Rice in 1993, he was at the IBM Almaden Research Center, where he managed the Mathematics and Related Computer Science Department. His research interests include database systems, computational-complexity theory, multi-agent systems, and design specification and verification. Vardi received his Ph.D. from the Hebrew University of Jerusalem in 1981. He is the author and co-author of over 150 technical papers, as well as a book titled “Reasoning about Knowledge”. Vardi is the recipient of 3 IBM Outstanding Innovation Awards and a co-winner of the 2000 Gödel Prize. He is an editor of several international journals and is a Fellow of the Association of Computing Machinery.

Chapter 3

Evaluation Report

Søren Damgaard Matthias Jarke Moshe Vardi
IBM Denmark RWTH Aachen Rice University

Executive Summary

The department of computer science at Aalborg University is now well on its way towards becoming one of the top computer science departments in Europe, and is already one of the strongest in Scandinavia. This can be seen, among other things, from the fact that the leaders of those three units that are active in core computer science, all appear among the 0.5% most-cited computer science authors in the ResearchIndex created by the NECI Scientific Literature Digital Library project in the US.

In the reporting period, the department has reorganized into four units, each led by an internationally well-known researcher. The international standing of each of the four units, and of the department as a whole, has improved significantly over the past five years, in terms of publications, industrial cooperation, and overall impact. Much of this success can be attributed to active decisions taken by the department during the last five-year period.

However, the enormous demand for computer science creates a strategic need for managing the foreseeable growth to the long-term advantage of the department with a host of associated novel challenges. Among others, these challenges include the definition of growth areas, a stronger willingness to accept management responsibility, and strategies for hiring and career planning. An excellent preparation of, and very clear and open discussions at the review have helped the reviewers frame a set of specific recommendations in this regard.

3.1 Evaluation Process

The main thrust of the evaluation was through a department-wide meeting in which the four research units provided overviews of their research activities and plans. This was preceded by a detailed self-evaluation report, which was also focused on the four research units. The reviewers found the process open, frank, and friendly. It was clear that the department itself benefited from the opportunity for everyone to get a general overview of the variety of research activities in the department. The reviewers recognize that, while this kind of self-evaluation has become relatively common in North America, computer science in Aalborg is still playing an advanced and in some parts pioneering role within continental Europe. Many of the comments provided would not have been possible within most other European universities, thus the large number of discussion items and recommendations below should be seen as a positive rather than a negative sign.

The two-day meeting was organized by first presenting the work and strategy of each individual unit, then discussing overall departmental strategy. The remainder of this report follows the same sequence, ending with some recommendations about further possible improvements of the review process itself.

3.2 Evaluation of the Units

In the sequel, a brief evaluation of each of the four individual units—databases and programming technologies, decision support systems, distributed systems and semantics, and information systems—is provided. In some cases, the unit-oriented discussion also covered broader issues which were particularly visible in one department, but may also apply to the department as a whole. In these cases, comments are in part repeated when discussing the departmental level, but still retain their relevance for the units where they were discussed.

3.2.1 Databases and Programming Technologies

This unit was created by merging the previously separate temporal database and programming language groups. Both groups have a rather different working culture, the database group traditionally strongly focussed on one topic, the programming language group consisting of largely individual researchers with relatively loosely related topics. There were two separate presentations of the existing work by each subgroup together with an overall vision presented by the unit leader.

Since the last review five years ago, the database subgroup has broadened its scope of research into new domains such as data warehousing and data mining with the goal of optimally exploiting its world-leading core competency in temporal databases. The group has accomplished a clear definition of top publication

outlets for its research results, and has obtained an impressive track record in these outlets. The international linkages in Europe and worldwide are being retained by extensive research exchange activities, joint projects and papers. The previously somewhat haphazard cooperation with industry has been strategically strengthened, e.g. by selecting strong strategic partners in industry as well as pursuing own start-up strategies.

Summarizing, it can be said that the database subgroup has now joined the small league of top database groups in Europe in terms of research quality, and of visibility both in the scientific and industrial communities. For quite a while, this impressive success story has rested essentially on the shoulders of one person, with all associated advantages and risks. It is now time to actively manage the strong growth of the group by distributing leadership, and it is heartening to see that first successful steps in this direction have been taken in the last couple of years. An additional remark of the reviewers is that the size of the group is now large enough to consider a portfolio strategy in which besides the existing short- and medium-term oriented research goals also some higher-risk long-term topics could be included.

The programming technologies subgroup has undergone significant change in the reporting period, with one key member leaving the department. Considering these difficult circumstances, the results achieved by the group can be considered quite respectable, including the publication record. However, these results largely concern individual or small-team work in different sub-fields of programming language technology and software engineering. The members of this subgroup individually, and the department as a whole, should seriously consider the different options for alternative futures: (a) continuing the attempt of creating a coherent whole by attaching this work closely to the database programming aspects of the database group, (b) complementing the Information Systems unit into the classical software engineering community (see remarks there), or (c) leaving the choice to each individual senior researcher which unit to join.

For the future of the unit as a whole, an exciting vision of being a key player in the growing field of mobile internet services research was presented. This vision seems to optimally address several regional strengths, and indeed could be supported on a grand scale by a strategic research proposal with other local players, including linkages into other schools of the university. Additionally, presentations by the other units show a lot of work being closely related to this topic as well; this synergy should not be lost by the separation into units (see also section on departmental strategy).

3.2.2 Decision Support Systems

The Decision Support System (DSS) unit has not undergone any significant changes during the evaluation period and has not merged with any of the units from the previous period. The reviewers find the research performed by the De-

cision Support Systems group of top value both commercially and technically in the area where the group has chosen to work.

Bayesian networks is a somewhat narrow area. However, the group has continued to demonstrate very successfully during the evaluation period that a theoretical discipline can find practical value as demonstrated in the SACSO project together with an external strategic partner (HP). The DSS group has benefited positively from cooperation with HP in applying the research on a real life example of using Bayesian networks in troubleshooting printer problems. Furthermore it is very positive that three patents have been filed in the period. This is an important way of demonstrating impact of research. The three patents were filed by HP and HP also provided the necessary legal support in the filing process.

If patents more broadly become an area of importance it is recommended that this is taken up as a university role instead of a case by case decision with an external partner.

The people involved from HP in the SACSO project are physically located in the same building as the researchers from the Decision Support Systems group. It was pointed out that this had led to a pleasant working atmosphere where there was no ‘them’ and ‘us’ syndrome.

The core DSS group is relatively small consisting of 8–9 people of which 3 are permanent positions including one full professor. The group is led very much based on a strong teamwork idea where the overall research theme keeps the group together. During the verbal presentation strong emphasis was put on explaining that the backbone of the organization of the group is the Wednesday workshop and the Friday seminars as well as physical conditions supporting agenda-less meetings. It was explained that the premises for success hinged on two things: 1) that all researchers have a common platform of knowledge from which the issues grow and 2) all researchers have a direct interest in all issues.

The advantages of this way of organizing the units research was explained as giving a stimulating working atmosphere as well as a quick integration of newcomers leading to a very productive environment. The disadvantages mentioned was that this led to a narrow scope, inbreeding and little room for expansion. The risk is that the group in the longer term will lose its potential if it does not attract researchers who can challenge the research profile. Furthermore the loose structure puts an upper limit to the size of the group in that the structure will probably break down once the group increases by 2–4 more people.

In the self-evaluation report it is mentioned that the collaboration with other research units of the department has been limited to a couple of presentations between the DSS group and the Distributed Systems and Semantics group. It is pointed out then that “There seems to be basis for further cross fertilizations, as the techniques and methods under study in the two units seem to have common traits”. The reviewers will encourage exploiting this in the next evaluation period.

Collaboration with groups within Denmark seems modest presumably because of small overlap between the narrow DSS research profile and research profiles

of the other Danish universities. The international collaboration is stronger as evidenced by a number of leading roles in the top conferences of the field.

During the evaluation period the unit has supervised five PhD students. Three have graduated in the period and two are expected to graduate within the next 2 years. This was considered lower than the DSS group would like. It was pointed out that it was difficult to attract new PhD students to work with the DSS group. On one hand because the field is considered too narrow and on the other hand because external partners could offer job opportunities that the students considered more attractive.

The research plan presented for the period 2001 to 2005 focuses on advanced data analysis (data-mining), decision-theoretic and game-theoretic agents (e-service, auctioning) and mentions the possibility for import of a new area. The reviewers find the research plan rather sketchy and would like to see more tangible and elaborate ideas presented. The group is encouraged to more clearly define goals and milestones for the next period. As discussed in Section 3.4, a five-year plan should focus on research priorities, the plan for pursuing them, success criteria and milestones.

3.2.3 Distributed Systems and Semantics

The Distributed Systems and Semantics (DiSS) group was created in the present evaluation period as a merge between the Computer Systems and Semantics group to better exploit the underlying interrelationships at the research topic level.

The Semantics group has been the oldest and within classical computer science the internationally best-known team in the department. The reviewers are impressed that this standing was even improved during the past five years, not only as previously by high-quality publications but now also by successful first steps into technology transfer. It is striking to see the very high number of activities that have taken place.

The results obtained in verification and validation as demonstrated for instance by the successful tool UPPAAL are excellent and very promising. In this specific area it is of course tempting to put theory into practice and also try to share the tool with other interested users. The group decided to invest resources in building a graphical user interface to the tool to attract new users and lower the barrier of entry for use. This strategy without any doubt proved successful in terms of actually attracting many new users. This is a good way of testing the strength and usefulness of the theory as well as the tool itself. The fundamental question is always how much coding should the group actually engage in from a research point of view and where is the border between sufficient and too much? During the discussion it became clear that the group did not use its own research resources but decided to hire an external programmer for the work. This is a nice example of a good and rational decision.

The group has been involved in an impressive number of activities with external local companies. For example, the group helped the validation and development of real-time protocols in B&O by locating a bug in the software that the company by itself had been unable to for a number of years. The frustration expressed in the presentation was that despite this successful result the group seemed to have failed building strong relationship with B&O at management level. This is a classical example. Cooperation with companies needs to involve management on both sides to assure that mutual benefits are clearly defined. Probably the management in B&O was not aware that a real problem had successfully been solved by the DiSS group. Thus, the DiSS group did not get the full recognition it deserved. By assuring stronger problem recognition and thus buy-in at management level, this can be avoided. But it assumes that the research group is prepared to step up to a challenge that requires leadership in dealing with external relations.

The review shows that the group has gone through a long and successful period of management by enthusiasm. The merger with another group into the distributed systems and semantics unit does not just pose additional technical opportunities and challenges, but also means that the unit leader should take charge at the management level to help the unit develop a strategy beyond many nice individual results and to strengthen not just the external cooperations (which are evidenced by numerous joint papers with international coauthors) but also the internal cooperation within the unit. Given the actuality of the topic and the track record of the team, the reviewers are quite optimistic that this additional step can be taken successfully.

3.2.4 Information Systems

The Information Systems unit has a long-standing tradition as a leader in the so-called Scandinavian approach to software engineering which is characterized by a strong emphasis on work impacts of information technology and participative design. Many of these ideas have by now influenced the (Business) Information Systems community both in research and practice. In the last few years the unit has therefore faced the trade-off between continued strategic leadership in the Nordic region and the challenges of global visibility and competitiveness. To remark on this trade-off, the reviewers would certainly not want to suggest to give up the richness of regional diversity in favor of following the “standard” of US-defined research questions and approaches; rather, the IS unit should bring their own approach actively into the international debate.

Report and presentation demonstrated that the group has understood this challenge, partly mentioned already in the previous review, very clearly and has taken significant and successful action. Major advances can be noted in the record of international publications which can now be considered very good, as well as in obtaining suitable external research funding and strengthening the

infrastructure for intensive cooperation with industry. The reviewers appreciated, in particular, the additional focus on linking research with innovative approaches to student education. They also note the strong external placement record of Ph.D.'s from this group—in the long term one of the most important measures of impact. Summarizing, the unit—even though it is one of the oldest and most established in the department—has shown its liveliness by re-juvenating itself in many dimensions and by further increasing its international visibility and future potential.

The unit defines its research approach as looking at Software Engineering from an Information Systems perspective. However, looking at the publication record, the practical implication has been that most publications from the group—even where they investigate typical SE topics such as software process improvement—have appeared in outlets (good ones!) of the Information Systems community. While this approach clearly makes paper acceptance easier due to greater similarity of cultures, pursuing this avenue alone appears strategically too defensive. The reviewers believe that publications of this very valuable cross-disciplinary approach would have a much higher impact when confronting the Software Engineering community within Computer Science as well. One of the possible means to make progress towards this goal could be to team up with people inside or outside the department who are already active in this community, in order to get the debate started in a language accessible to the SE community.

Under the umbrella of the Information Systems unit, a strong multi-person research activity in the highly promising direction of interactive media has gained momentum, without being yet visible in the official structure of the department. The recent hiring of a prominent senior researcher at the professor level should present an excellent opportunity for this activity to define, in the plan for the next five years, its own research agenda and increase the overall image of the department by making this activity structurally visible.

3.3 Evaluation of the Departmental Strategy

During the evaluation of the four units some broader issues emerged, which seem to involve not only the units but also the department as a whole. Most importantly, the foreseeable growth—caused by increasing student demand—needs to be actively managed to continue the way towards becoming one of the top computer science departments in Europe.

3.3.1 Management of Growth

Computer science today is a popular discipline among undergraduate students. The world-wide shortage in information-technology workers essentially guarantees good jobs for holders of computer-science degrees. This popularity drives the

recent increase in the enrollment in the computer-science program. In the short term, such an increase causes hardship due to increased teaching load. In the long run, however, the enrollment increase drives growth, as academic staff positions are driven by student enrollment. The department, however, needs to give more thought to what should its growth strategy be.

The department is currently organized into four research units. This organization brought stability, and enables the senior staff to lead the units in a natural way. However, in a strong growth situation, the current organization may impose an overly rigid structure. For example, the prevailing sentiment in the department seems to be that growth ought to be funneled into the existing four research units, since teaching load is currently managed through the four units. The reviewers believe that that this poses a risk of freezing the department in its current research areas. It was noted, for example, that the current spectrum of research activities in the department is reorganized and broadened but not significantly different than what was already in place five years ago. Promising research areas the department wishes to expand into, such as interactive media or mobile internet services, were only mentioned at the unit level and should be made more visible at the departmental level. One example is mobile internet services. The reviewers believe that this is indeed a promising area and would like to encourage broadening the scope of this direction beyond the Databases and Programming Technology Unit. It is also important that this seems to be an area with a potential not only for cross-unit collaboration but also for collaboration with other departments in the university.

There is no question that the current organization served the department well by providing leadership to the research areas in a natural way and by putting in place a framework for smooth departmental operations. Faced, however, with significant growth over the next five years, the department needs to take a hard look at its operating model. Computer science is a dynamic discipline. The department needs not only to grow but also to evolve. Does the current organizational structure promote innovative evolution? What would take the department to the next level? What must be given up to facilitate a move to the next level?

One obvious place to start relaxing the current rigid structure is to separate the tight coupling between research units and teaching allocation. The computer science curriculum is roughly divided into introductory courses (such as introductory programming, algorithm and data structures, machine organization and discrete math) and advanced courses. While it is natural to couple advanced courses with research areas, the introductory courses could be viewed as a departmental responsibility, as, in principle, every academic staff member should be able to teach such courses. A looser coupling between the research units and the teaching allocation would facilitate more flexibility in hiring.

3.3.2 Management

The reviewers sensed some reluctance by the senior staff members to engage in “management”, as this seems to contradict the democratic character of the department. The reviewers, however, believe that it is the responsibility of the senior staff member to provide leadership. Much of this leadership can be provided via role models, but for the department to move to the next level, the senior staff members should shoulder more explicit leadership role. This is of urgent importance at this point, since the department needs to formulate its strategic directions before embarking on an ambitious growth program.

For the department to enhance its standing, its focus should be in maximizing its impact. Impact, however, can be maximized in many ways. The undergraduate program provide IT workers to the Danish economy. The graduate program provide advanced IT workers as well as academic staff members to Danish universities. The research program contributes to the intellectual infrastructure of the Danish economy and enhances the regional and international reputation of the department, enabling it to attract high-quality graduate students and staff members. The senior staff members should lead the department in discussions on impact enhancement. How much focus should there be on technology transfer? What is the role of patents in enhancing technology transfer? What is the role of the Nohauz industry club in enhancing industrial relations? What is the “return on investment” for software development in the research program (e.g., vehicle for research, vehicle for technology transfer, etc.)? What publication fora maximize the visibility of the department research program? What is the proper balance between medium-term and long-term research? What should be the balance between Nordic visibility and international visibility? What is the importance of cross-unit and cross-departmental collaborations? What is an optimal level of external funding? What is a desired level of Ph.D. production? How should impact be measured? (Citation count was mentioned as one, admittedly coarse, measure of research impact.) There are no definitive answers to these questions, but frank discussions would lead the department and the research units to their own answers.

In addition to intellectual leadership, the role of management is the focusing of attention and resources. The loose management style of the department means that key issues, such as graduate recruiting and staff recruiting, are not getting the required attention. The department needs to identify its key tasks and ensure that each such task is assigned to a staff member or a committee of staff members. The fundamental principle of accountability dictates that for every key task there should be a responsible person. The reviewers believe that the department will see immediate results from such a focusing of attention and resources.

The importance of physical co-location in developing and managing a joint strategy should not be under-estimated. The reviewers note that the co-location of all units in the E-Building has improved cooperation in the department, but

the growth makes it urgent to make plans for physical expansion in order to avoid a re-scattering of the department.

3.3.3 Hiring

There is no doubt that an ambitious growth program would very quickly encounter the difficulty posed by the scarcity of qualified personnel, as other Danish computer science departments undoubtedly are embarking on similar growth programs. It is important, however, that the department does not relax its currently high standards in hiring. The reviewers believe that the department can succeed in growing only if it recruits globally. Denmark has a reputation as one of the more open countries in Europe, where English is universally spoken. The department should take advantage of this reputation and scout for international talent.

One danger that the department should be careful of is that of inbreeding. It is very tempting, when facing a stiff competition for human resources, to try to hold on to “our own”. Hiring one’s own doctoral students, however, results in stagnation and lack of intellectual diversity. The department should view placement of its graduates in other universities as a measure of success, not as a failure to retain them. One would hope that other Danish computer science departments would take the same high-minded perspective, resulting in stronger computer science departments across Denmark. Where the department’s own graduates are occasionally hired for permanent academic positions, that should happen only after they have spent a few years elsewhere, to ensure that they have broadened their intellectual perspective. (It was noted that Israel, a small country with an internationally strong computer science program, informally requires graduating doctoral students to spend at least two years on an international assignment before assuming academic staff positions.)

The reviewers wish to note for the record that the ability of the department to succeed in attracting high-quality personnel is conditioned on the university’s making academic positions in computer science attractive for prospective applicants. Computer science professionals are facing an unprecedented international demand for their talents. Danish computer science departments will unquestionably fail to grow their computer science departments and maintain their high quality without ensuring that these departments are provided with adequate space and that the gap between academic and industrial salaries does not expand to make academic positions rather unattractive. The reviewers believe that the current situation regarding these issues is not quite satisfactory and is in danger of getting much worse in the coming years.

3.3.4 Cooperation with Industry

One fundamental question is external relationship with companies and to what extent the units should be encouraged to build strong relationships. The underlying assumption in most of the self-evaluation has been that the more external cooperation the better. The reviewers believe it is important to have a well defined policy and strategy not only at group level but also at departmental level. Elements in a strategy could be to define the goals in terms of

- Funding levels,
- No of head counts from external partners,
- Patents filed,
- No of Ph.D. students on joint projects,
- Agreed policy on hiring of students,
- Publishing policy, and
- Other research impact.

The case with UPPAAL and the cooperation with HP are successful examples of applied research. The department and the groups should develop a visible strategy for basic research and applied research including cooperation with external partners for the next evaluation period. This strategy should be translated into an operational plan for the period with specific goals including specification of the metrics by which the achievement will be measured. For instance the difficulty some of the groups seem to have in attracting/keeping new Ph.D. students versus the attractiveness for graduated students to work with the external collaboration partners was mentioned as an issue.

3.4 Suggestions for Future Review Processes

As noted above, the reviewers found the process open, frank, and friendly. At the same time, the reviewers found three areas in which the evaluation process itself could be improved.

Firstly, in several units the evaluation criteria were rather vague. This means not only that the reviewers were not provided with clear criteria by which to measure the success of these research units. It also means that the research units themselves did not seem to have clear criteria to evaluate their own success. Furthermore, there seems to be a reluctance for either the department or the research units to formulate such criteria. The sentiment was expressed that it is up to individuals to define their success criteria. If, however, individuals can

define their own success criteria, then success becomes tautologous. The reviewers believe that the department and the research units should define clear evaluation criteria (see discussion on impact maximization in Section 3.3.2). The next review panel can then judge both the soundness of criteria and the department's success in meeting them.

The second weakness in the process is the lack of uniformity in the research plans of the four units. These ranged from being either too specific or too general and vague. The reviewers believe that five-year research plans ought to be strategic. They should state the top three to four research priorities for the coming five years, the reason for their importance, and the units' plans for pursuing these research priorities. The plans should also include the units' success criteria and, if appropriate, milestones for measuring progress.

Third, the narrow focus of the evaluation on the four research units alone made it hard to assess the quality of the department's research activities as a whole, as this quality has dimensions that go beyond individual research units. Such dimensions include the very organization of the department into the four research units, strategy for further growth, industrial relationships, external funding, intellectual property, doctoral program, and the like. To be clear, there was no reluctance to address these issues, and when they did come up they were discussed in a frank and open manner, enabling the reviewers to offer the comments in Section 3.3. Still, these issues were initially not central to the evaluation process. As some of the above comments by the reviewers indicate, it is their belief that such issues have a major impact on the overall quality of research in the department.

Part II

The Unit's Research Report

Chapter 4

The Department of Computer Science

In this chapter we provide a general introduction to the Department of Computer Science with a particular focus on research. We provide some background and a brief history, an overview of the organisation, the research topics covered by the department's research units, and the resources. Furthermore, we give a general description and evaluation of the department's research plan for the period 1996-2000. Finally, we sketch the overall research plan for the next five year period.

4.1 Introduction

The department's research has software, use and performance of software, as well as information and data as its subject. In particular, there is research in use of software in organisations, software engineering, management of software engineering, programming and languages, data management, data analysis and data mining, techniques for decision support, networks and protocols, techniques and models for distributed and parallel software, and tests.

The research approach is fundamentally constructive and embraces the analytical mathematical research, the experimental research with algorithms, systems, techniques and methodologies, as well as the analytical empirical research.

Computer science is connected to mathematics, engineering, and to the human and social sciences.

4.2 Background, Conditions, and Resources

We first describe the local organisational context of the department, then its history and finally the resources.

4.2.1 Computer Science at Aalborg University

Aalborg University was established in 1974 and is the newest university in Denmark. It consists of three faculties:

- The Faculty of the Humanities;
- The Faculty of Social Sciences;
- The Faculty of Engineering and Science.

The Department of Computer Science is part of the Faculty of Engineering and Science, which presently is divided into 13 departments. The head of department and the board manage the department. The board consists of the head of department, 2 scientific staff (also part of the daily management), and 1 administrative/technical staff. The head of department and the board are elected among their peers for a three-year period.

The Faculty of Engineering and Science decides the budget for the department, based on the budget it gets from the Ministry of Research and the Ministry of Education. A budget model taking produced teaching as a primary parameter determines the department's internal funding including number of scientific staff and funding for support staff, laboratories, offices and research. The department influences the internal funding by deciding to teach more or less than expected in the budget model. The department is responsible for managing and conducting research and for supplying teaching to the study boards. The department is also responsible for administrating the internal funding. The Danish research councils, the European Union, Center for IT Research, and IT companies provide external funding for research projects.

Study boards govern education and degree programmes. The department's scientific staff performs the teaching requested by the study boards. The department provides teaching to the Natural Science Study Board, the Electrical Engineering Study Board, and the Basic Education Study Board. A study board consists of equally many scientific staff and students elected for a three-year period. A member of the scientific staff heads the study board.

4.2.2 A Brief History of the Department

The history of the department goes back to 1976, when a minor degree (i.e., in practice a bachelor degree) in computer science was established in combination with a graduate degree (called cand. scient.) in mathematics. In the early 1980s, major degrees in computer science (cand. scient.) and computer engineering (cand. polyt.) were established with associated undergraduate studies in mathematics and electronics engineering, respectively. From this early beginning and until 1990, the number of awarded graduate degrees grew from less than 5 per year to approximately 40, and the number of scientific staff (full, associate,

and assistant professors) grew from 4 to 8. The first full professor was appointed in 1987.

From the mid-1980s to 1993, the Ministry of Research and Education reserved special budgets for computer science and engineering at all Danish universities. This meant that computer science was able to initiate a relatively high volume of Ph.D. studies (3-4 per year), and that the Faculty granted positions as assistant professor and associate professor to follow up on this. At the end of the present period (2000), the scientific staff consists of 38 full, associate, and assistant professors.

Computer science was previously together with mathematics in the Department of Mathematics and Computer Science that again was part of the Department of Electronic Systems. The department was first separated from mathematics in 1995, and then in 1999 it became an independent department at the same organisational level as Mathematics and Electronic Systems. As a consequence the department received its internal funding directly from the Faculty starting from 1999.

4.2.3 Research Funding

The internal research funding depends directly on the amount of teaching produced by the department. On top of this comes a significant amount of external funding for research projects in the department.

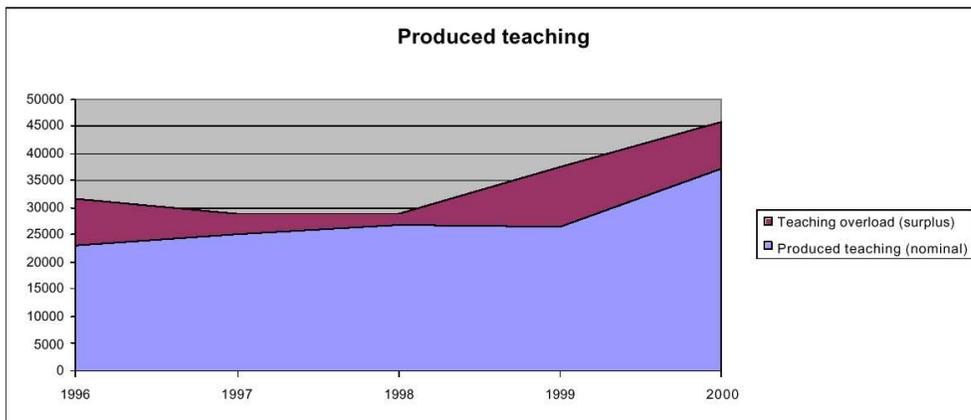


Figure 4.1: Development in produced teaching hours

The development in the produced hours of teaching is divided into the nominal teaching load put on each scientific staff member by the faculty and the teaching overload. The nominal produced teaching has increased as the number of scientific staff has increased. The teaching overload is on average 25%. The teaching overload rests on the shoulders of tenured scientific staff.

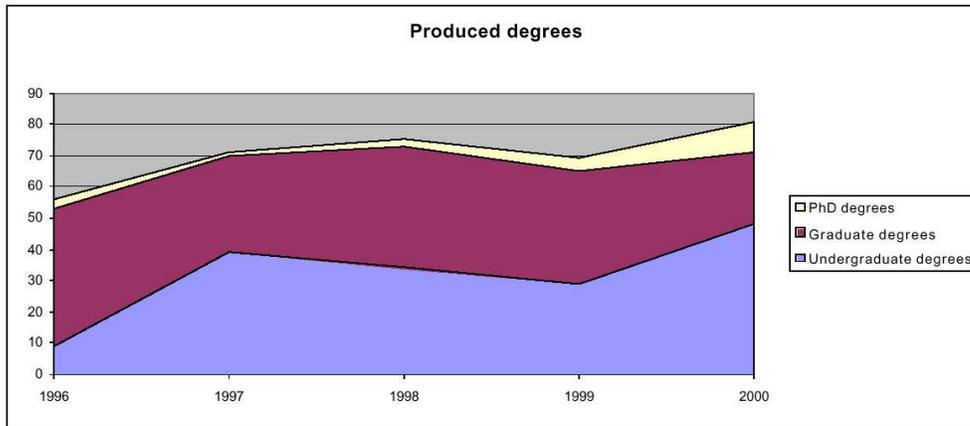


Figure 4.2: Development in produced degrees

The development in produced degrees is affected by two hidden structural changes. The way undergraduate degrees were counted changed in 1997, as all graduate students (all in a five-year programme) would get an undergraduate degree halfway through. The curriculum for the graduate degree changed with effect in 2000, where the students were given the option of studying computer science for an extra year, which many have chosen to do. The increase in Ph.D. degrees is significant.

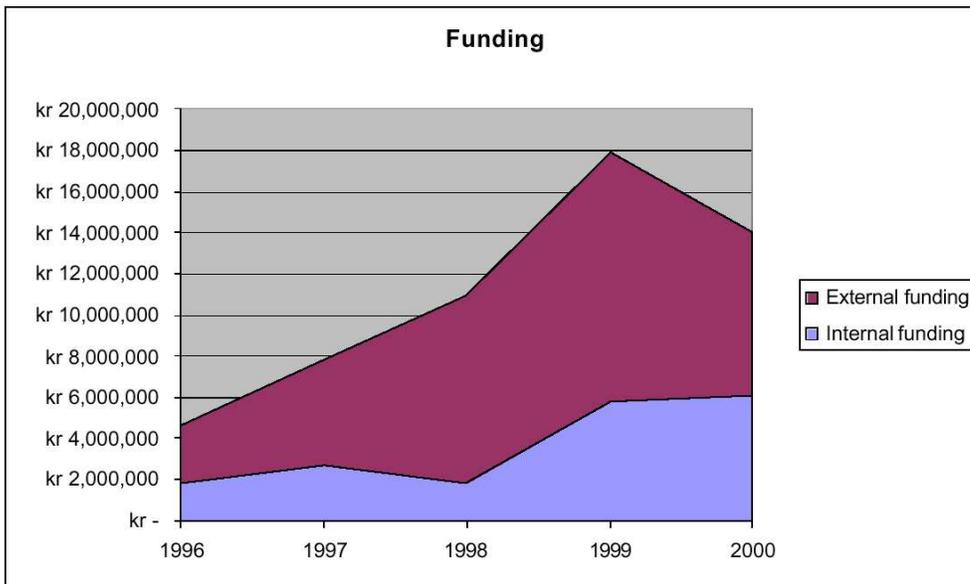


Figure 4.3: Development in the Department's funding

The development of funding shows two patterns. The internal funding from

the Faculty increased dramatically in 1999. During 1996-1998 the department was still part of the Department of Electronic Systems, that was then responsible for employment of the technical and administrative staff. After becoming a department on its own, the internal funding and the employment responsibility were transferred directly to the department. During the whole period there has been a significant increase in the external funding for research projects established within the department. External funding was on average at 7.4 mill. kr.

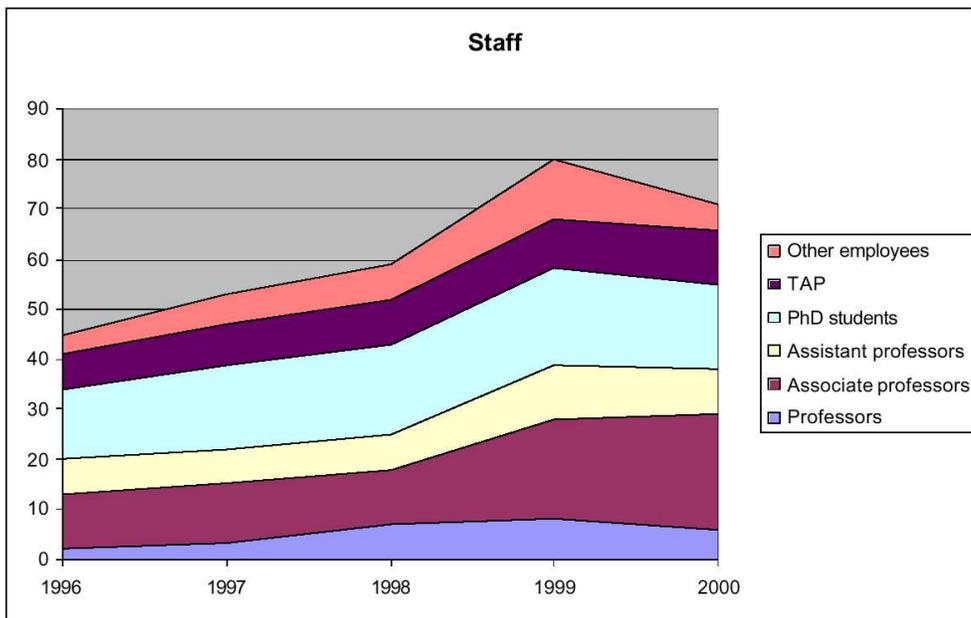


Figure 4.4: Development in staffing (full and part-time)

The development in staffing (counting all full-time and part-time staff) is shown in Figure 4.4. Professors and associate professors are tenured, while assistant professors are on three-year contracts. Ph.D. students are employed for three years, during which they are expected to complete their study. Technical and administrative personnel (TAP) are full-time. The category of other employees covers temporary research assistants employed in projects and hired help. Figure 4.4 shows a stable period during 1996-1998, followed by an increase in 1999 for associate professors with other categories remaining relatively stable.

4.3 The Department's Organisation

The department's organisation is described in terms of the research units, education, Ph.D. studies, and collaboration with the IT industry.

4.3.1 Research Units

Until 1998 the department had been organised in six research groups. In early 1999, the scientific staff divided itself into four research units (Database and Programming Technologies, Decision Support Systems, Distributed Systems and Semantics, and Information Systems). The main purposes were to make the department visible through its research profile and also to de-centralise the responsibility for research, planning, and teaching. Merging four existing research groups formed the Database and Programming Technologies Unit and the Distributed Systems and Semantics Unit. The research units form the basic social environment for the individual researcher.

The Database and Programming Technologies unit covers data management and techniques and tools for data access, in addition to the design, implementation and application of programming languages, their environments, and tools.

The Decision Support Systems unit covers techniques and tools for decision support systems. The work is primarily based on Bayesian nets.

The Distributed Systems and Semantics unit covers real-time and distributed systems, networks, formalisms for the description and analysis of computer systems, tools for verification and validation.

The Information Systems unit covers, in particular, the development and application of computerised systems in organisations.

4.3.2 Research-Based Education

Research and teaching are intimately connected. The teaching duties of the staff in computer science are mainly directed towards the following areas:

1. The basic education year at the Faculty of Engineering and Science that initiates all the students of the faculty to scientific studies, and in particular to the problem-oriented and project-organised study form used at Aalborg University.
2. Computer Science and Informatics at the undergraduate level (for students of mathematics, computer science, computer engineering, and physics).
3. Computer Science, Software Engineering, and Informatics at the graduate level.
4. Ph.D. supervision and courses within the Computer Science and Engineering Programme.

Aalborg University employs project-organised problem-oriented studies. Every term has a particular problem theme. The students divide themselves into project groups consisting of 5-8 members, and each group is assigned a full-time working room. Approximately half of the time, the students work in the groups on topics within the chosen problem theme. It is the group and its supervisor that agree on the particular topic. The other half of the time, the students take courses of a more traditional form. About half of the courses are directly relevant to the projects, and the other half form part of the general discipline. Each project group is assigned one or more scientific staff members as supervisors for the project. Scientific staff members serve both as project supervisors and as lecturers on courses.

The main part of produced degrees consists of five-year graduate studies in Computer Science, Software Engineering or Informatics. Towards the end of the fourth year, the students choose an area of specialisation for the final year within one research unit. The final project work is conducted in close collaboration with the research unit to ensure a close relationship between research and teaching.

In 1998 a new graduate degree in Informatics was established. The production of candidates with a master degree will therefore double from 1997 to 2003.

4.3.3 Ph.D. Studies

The Faculty's Ph.D. studies are organised by its Research School that offers three-year degree programmes. The research school governs the Ph.D. study in Computer Science and Engineering, while all activities and courses are carried out in the department.

The department offers every year 1-2 Ph.D. positions for full-time research over three years. Many additional Ph.D. positions are offered by externally funded research projects. A few Ph.D. positions are devoted to collaboration with an IT company. Few Ph.D. students are without a Ph.D. position.

Ph.D. students devote the full three-year period to a particular research project that is within the interests of one of the research units. This work is supplemented by Ph.D. courses and most likely also by a six-month visit at a university abroad.

4.3.4 Collaboration with IT Industry: Nouhauz

Collaboration with the IT industry has always played a role for many of the research units' activities. Many externally funded research projects are co-financed by IT companies.

During 2000 the department has set up an organisation for structuring this collaboration and for providing support for an expansion in this area. The organisation is called Nouhauz (i.e. "No House" and "Know Hows"), and its purposes are to structure further education for IT practitioners, courses, networks where IT

researchers and practitioners meet, students' projects conducted in companies, IT practitioners as teachers in our courses and projects, and research collaboration.

4.4 Own Evaluation of the Department

The self-evaluations by each of the research units are given in Sections 5.7, 6.7, 7.7 and 8.7. Here we give an evaluation of the conditions for research provided by the Faculty and the Department.

4.4.1 The Department's Research Plan for 1996-2000

The previous research evaluation contained a strategic research plan for the period 1996-2000. Its main goals were:

- The number of research units should not be changed. Rather the existing groups must be consolidated. New research topics should be initiated through external or internal collaboration. Especially, the possibilities for collaboration within the Department of Electronic Systems and with industrial partners should be exploited. The leadership of the department can be instrumental in bringing about such collaboration. However, each individual research group should decide on the final priorities among its potential research activities.
- Due to the growth of tenured positions, an additional full professorship should be allocated.
- There should be a controlled growth in the number of tenured positions.
- The problem of teaching overload must be solved. We note that hiring more permanent faculty conflicts with the previous goal.
- Assuming that the problem of teaching overload will be solved, the annual extra income (approximately DKK 500,000) from this activity will vanish. This implies the need for at least a doubling of the external funding in order to maintain the quality of the technical platform (workstations and network).

Some of these goals have been met, some have been changed, but none have been left untouched.

- The research units were changed in early 1999 because the scientific staff wanted to form larger units and thereby strengthen leadership and distribute responsibility for both research and education. Thus, all units needed to have sufficient size and staffing to take on this commitment. Further, the growth for each research unit was negotiated between the research units, which led to a five-year agreement.

- The establishment of Nouhauz has had and will have increasing impact on our collaboration with the IT industry, in particular with the regional IT industry.
- The expansion of the department's scientific staff has resulted in the allocation and employment of two new professors during 1996-2000, and there has been controlled growth in the number of tenured positions. At the same time, the department has succeeded in increasing the rate of growth.
- The problem of teaching overload has not been resolved. However, the department is now in a position where it no longer depends economically on the income from the teaching overload. Also, a mechanism for explicating the teaching overload has now been institutionalised and is used in the planning and allocation of teaching.
- The economical situation for the department has improved considerably. While external funding creates more research activities, it is not critical for the operation of the core business of the department.

4.5 The Department's Plan for 2001-2005

The department is an organisational structure containing the four research units. As such the department is responsible for developing and maintaining the common strategies for research and for providing the best possible conditions and resources for the research units.

4.5.1 Strategy

The department's strategy is three-fold:

- The department must strengthen its position nationally to attract more external research funding and to recruit the best scientific staff.
- The department must strengthen its position internationally to further more collaboration and further improve research quality.
- The department must improve its collaboration with IT industry to create symbiotic benefits.

The strategy and its deployment must be reflected in the research units' strategies and plans. Further, the deployment may require changes in organisational structure, management, leadership, and research profile.

4.5.2 Policy Making and Deployment

The strategy will require several changes in the organisation of research:

- It is the management's responsibility to re-create adequate processes for policy making and deployment involving the management, leaders, and scientific staff.
- A proper incentive structure for scientific staff must be created and inculcated.
- The management of size must be installed.
- The problem of teaching overload must be solved.

Chapter 5

Database and Programming Technologies

5.1 Profile of the Unit

Driven by the continued advances in hardware technologies, by the diffusion of the Internet, and by the increasing complexity of software systems, the areas of database management and programming languages and environments are faced with abundant research challenges. Today's software systems manage large amounts of traditional and non-traditional data, including temporal, spatial, spatio-temporal, dimensional, multimedia, and semi-structured data. The research objective is to develop technologies that meet programming and data management needs posed by software systems in general and by data-intensive applications, in particular.

The areas covered include general-purpose programming languages with an emphasis on languages in the object-oriented paradigm. Also covered are special-purpose languages, such as languages for management of different types of data. Also included are environments and tools that support the design, implementation, documentation, versioning, and configuration management in software development.

In the area of databases, the research relates to business intelligence, including data warehousing and on-line analytical processing, and to database integrity. In addition, it relates to temporal, spatial, spatio-temporal, and mobile databases, including conceptual modeling and database design, data models, query processing, indexing, and applications.

World-wide web related research covers semi-structured data management, Web application development, database search engines, and the use of functional languages for Web authoring and server-side programming.

The research approach has a technological focus and is primarily constructive in its outset, but also integrates experimental and analytical elements. Construc-

tive activities include the design of concepts and frameworks, as well as the design and implementation of algorithms, data structures, languages, and systems. Experimental activities cover the testing of constructed artifacts, including both prototype-based experiments and simulation-based performance studies. Analytical activities include complexity analysis and language evaluation. The emphasis is on the development of theoretically sound results that solve actual real-world problems in the medium term.

The unit covering database and programming technologies was established towards the end of the present evaluation period, by merging the research two groups, database systems and programming systems.

5.2 Activities and Results

This section first surveys the unit's research in four broad areas, thus covering most of the unit's research. Then follows descriptions of four broad, so-called framework projects that have received substantial external funding. Descriptions of six more focussed and smaller-scale projects conclude the section.

5.2.1 Research Results

The main research results from the database and programming technologies unit concern temporal and spatial databases, data warehousing, and programming languages and tools. This section briefly summarizes these results.

Data Warehousing

Within data warehousing, the efforts have concentrated on data models for multidimensional data, physical database design for multidimensional databases, the integration of multidimensional data with external data, and the accommodation of advanced application areas, covering clinical data warehouses and analysis of web-server logs.

Data modeling has been the topic of several studies. A conceptual data model, to be used for the design of warehouses, has been developed [109]. Other efforts have resulted in a logical model for complex multidimensional data [24, 88]. In this context, the important problem of handling imprecise data has been addressed [24, 91]. Another interesting area is the modeling of the operations to be performed on multidimensional data. One study [44] considers exclusively the MD-join, while other studies [24, 88] present a complete algebra for multidimensional objects.

In the multidimensional context, the creation of a number of materialized views to support efficient query processing is an important aspect of physical

database design. These views contain partial answers to queries, e.g., joins or aggregations of base data, so an important problem is the selection of the right views to materialize, given space or time constraints. Before the selection can proceed, it is necessary to construct the set of all possible views [43]. One study [45] has considered how to minimize the amount of data to be stored in order to obtain self-maintainability, i.e., updates to the base data can be processed in the data warehouse without access to the original base data. Previous research has only considered the problem of using pre-aggregated data for fast query response for a limited set of multidimensional structures, which far from handles all real-world situations. This has been extended significantly so that irregular hierarchies and fact-dimension relationships are accommodated [92, 166]. This aspect is being patented. The systematic management of temporal data in data warehouses has also been studied [57].

Language facilities, query processing techniques, and a prototype system have been provided that offer one solution to the problem of integrating multidimensional data with external data [70, 93].

Two complex application areas have been studied, namely clinical data warehouses and analysis of web logs (click-streams). The state-of-the-art of industrial products for clinical data warehouses was surveyed [85, 86], and central, special challenges posed to data warehousing technology by clinical data warehouses were identified [87]. The problem of analyzing sequences of clicks in web server logs has received initial attention [46].

Programming Languages and Tools

Concerning object-oriented languages, the support for the correct composition of modules and similar entities in systems with support for advanced separation of concerns has received attention [64, 66]. This work includes the provision of support for the specification of variability and encapsulation of modules in such systems [67]. Finally, new abstraction mechanisms in the object-oriented language **gbeta** have been designed and implemented, and the compilation system for **gbeta** has been developed.

Next, the functional programming paradigm has been shown to fit well with Web authoring using the XML and HTML markup languages [134, 152]. The HTML markup language has been mirrored in the functional language Scheme, and the abstraction mechanisms of Scheme have been used to obtain the advantages of XML in a fairly simple way. On this basis, a number of substantial applications have been developed that validate the thesis—that the programmatic approach of using functional programming for Web authoring—is a realistic alternative to the more conventional approaches, both with respect to the authoring of static Web documents and server side applications [83].

Research on program understanding has been conducted that takes its starting point in the research on literate programming and hyperstructure programming

environments. The idea of 'elucidative programming' has been coined as a variation of literate programming [82]. The main idea is to separate the documentation from the program, and to relate places in the documentation with abstractions in the program via hypertext links. Via two different, operational prototypes, the feasibility of supporting elucidative programming with practical tools has been demonstrated [81, 84, 22].

Some research concerns various aspects of configuration management. Within versioning models, focus has been on the modeling of data in software development projects, and change management has received special emphasis. In relation to the modeling of history in collaborations, a tool for versioning of hierarchical documents has been provided [158, 53], which uses version control concepts and mechanisms to enable the collaboration during the authoring of shared documents. Concerning general, object-based environments, the potential of an object-based approach to the modeling of software development environments was investigated as an alternative to traditional file-based approaches. Objects encapsulate files, enabling the association of attributes, relations, and actions with files. Finally, a taxonomy for change management has been explored with the objective of establishing a model for how to carry out the change management in software development organizations.

Spatial and Spatio-Temporal Databases

In spatial and spatio-temporal databases, much work has focused on conceptual and logical design and has resulted in comprehensive specifications of requirements, definitions, and notations of relevance in spatial and spatio-temporal applications [98].

Based in part on these requirements, spatio-temporal conceptual models that extend the Entity-Relationship model [31] and UML [26, 99] have been designed. With the objective of further facilitating the conceptual design activity, a set of modeling abstractions has been proposed [111] where each abstraction represents a semantically autonomous excerpt of a spatio-temporal conceptual database schema. In the same context, we proposed a model for the design of so-called phenomena in spatio-temporal databases [173].

Of relevance to logical design, two extensions to the relational model have been developed that accommodate spatial data [110, 77] versus spatio-temporal data [60]. Furthermore, a CASE tool [108] is being built that supports the proposed conceptual and logical models, including the translation of conceptual schemas to logical schemas. Finally, a framework to facilitate data mining in spatial environments has been proposed [30, 107].

Considering continuously (spatio-temporal) moving objects, the group has taken part in establishing a foundation for representing and querying such objects [14]. A framework for representing the trajectories of moving-point objects has been developed that emphasizes data captured via sampling using GPS and

considers the inherent uncertainty [94].

In query processing, focus has been on a variety of spatio-temporal indexing problems [135], including the indexing of the current and near-future positions of continuously moving objects [100], the indexing of the past positions of moving objects [96, 38], and the combined 2D or 3D spatial and now-relative bitemporal indexing of discretely moving objects [39].

Temporal Databases

The research results obtained so far by the database community in temporal data models and query languages clearly demonstrate that database applications that manage time-varying data may benefit substantially from built-in temporal support in the database system. The group took an active, often leading, role in progressing the state-of-the-art in the management of time-varying data at the international level. Among others, we made contributions in the areas of data modeling [17, 16, 63, 127], database design [18, 36, 37, 68], data models and query languages [10, 11, 12, 40, 101, 103, 153, 154], efficient query processing [27, 28, 29, 41, 55, 56, 62, 187], and systems architectures [58, 104, 105, 120].

A significant part of our work investigates the design and development of temporal database technology that is expressly transferable. Such technology allows for the continued operation of legacy code, for the harmonious coexistence of legacy and new application code, and for the reuse of programmer expertise and knowledge [11]. It also allows for maximal, effective, and efficient reuse of functionality already provided by current database management systems [10]. Together, these properties lead to a technology that is attractive to use and manageable to implement, as illustrated by several prototype systems that we have developed.

5.2.2 Framework Projects

Nykredit Center for Database Research, January 1998–December 2003

Two governmental reports from 1997 conclude that Danish computer science and engineering candidates are only moderately qualified in database technology and that the existing competence in database technology in higher education is divided among few researchers and is sparse and vulnerable in a national perspective. On the other hand, the ability to make effective use of the newest database technologies is essential to ensuring the competitiveness of companies and institutions, thus rendering it of strategic national importance to establish and maintain a solid and up-to-date competence in databases in the nation's higher-education programs in computer science and engineering.

Nykredit Center for Database Research has as its objective to contribute to increasing the competence in database technology in Denmark in the near and

longer terms. Specifically, the center is responsible for conducting database research at a high international level; the center is responsible for offering internationally oriented research training with the longer-term objective of contributing to the establishment of new database research groups at Danish universities; and to meet short-term demands, the center offers courses and collaboration in database technologies.

The center embraces and provides infrastructure to all other database research activities in the unit. Topics covered consequently include temporal and spatio-temporal databases, data warehousing, and world-wide-web data management. Four Ph.D. students have graduated so far, and more are underway, five of which are scheduled to graduate during 2001.

To maintain a high quality and relevance of the research activities, the project supports intensive collaborations among members of the unit and leading international researchers in each of these areas. A visiting faculty program makes it possible to invite guest researchers for short stays. So far, the center has hosted almost twenty visiting researchers. The project also supports its own researchers in visiting other research institutions. Researchers from the center also collaborate with a number of companies and institutions in Denmark, including EUMAN, Informix, Kommunedata, Kort- og Matrikelstyrelsen, MindPass, netnord, Novo Nordisk, Nykredit, Oracle, and TimeChain.

The center is governed by a steering committee with two representatives from each of the Nykredit corporation, Center for IT Research, and the researchers associated with the center. The Nykredit corporation contributes DKK 5,000,000 and is the main sponsor. Center for IT Research contributes ca. DKK 3,750,000. Kommunedata, the Danish Academy of Technical Sciences, and Kort-og Matrikelstyrelsen have contributed ca. DKK 2,800,000 (two Ph.D. stipends). Finally, university sources have contributed ca. DKK 2,500,000 (two Ph.D. stipends).

ChoroChronos, August 1996–July 2000

ChoroChronos aimed at enabling established European researchers in spatial and temporal databases to coordinate their work and integrate their findings in their respective areas of expertise. The project enabled these researchers to work together to unite their expertise in order to design, implement, and evaluate database technology for the handling of spatio-temporal information. Eleven members of the unit contributed to several technical aspects of the project. The contributions are documented in ca. 60 papers and also encompass substantial software. Key contributions are outlined next.

Within conceptual modeling, the unit's researchers have explored requirements and conducted case studies. Conceptual data models based on the Entity-Relationship model and UML have been developed that offer built-in support for accommodating a variety of spatial and temporal aspects of data. Perhaps most notably, support for complex spatial relationships and so-called spatio-temporal

phenomena has been provided, as has a prototype CASE tool.

In query languages, the relational model has been extended with the support for so-called spatial quanta and the core of a spatio-temporal SQL has been designed.

In relation to query processing and indexing, the unit has produced several results. Two indices have been designed for general bitemporal data that include two-dimensional, continuously growing regions (these occur due to the presence of the database variable *now* that denotes the current time). Both are based on the R-tree. Another R-tree based index has been developed that supports spatio-bitemporal data. It supports bitemporal data that is more general than that supported by the previous two indices, and it offers the ability to prioritize spatially or temporally selective queries. Only discrete spatial change is accommodated. Next, a framework for representing the trajectories of moving-point objects has been developed that emphasizes data captured using GPS and considers the inherent uncertainty. Finally, work has been conducted on the indexing of the trajectories of moving point objects, as well as on the indexing of the current and anticipated future positions of continuously moving objects.

Concerning architectural issues, the group has explored layered architectures and the associated techniques for DBMS implementation. Special attention is being given to the implementation of well-behaved temporal transactions while reusing a conventional transaction processor. New work has been initiated that generalizes the layered architecture, a central goal being to appropriately divide query processing between the layer and the underlying DBMS, based on the functional and performance capabilities of the layer and the DBMS.

The project was supported by the European Commission under the TMR Programme and involved also research groups from the National Technical University of Athens, Fern Universität Hagen, University of L'Aquila, UMIST (Manchester), Politecnico di Milano, INRIA, University of Thessaloniki (with the Agricultural University of Athens as a subcontractor), Technical University of Vienna, and ETH Zürich. Timos Sellis of the National Technical University of Athens was the project coordinator. Christian S. Jensen was the local scientist-in-charge. The total funding of the project was ca. EUR 1,600,000, of which the unit received ca. EUR 175,000.

EURESCOM P817, January 1998–December 1999

In today's dynamic and highly competitive environment, telecommunication companies (telcos) increasingly see data as a key asset. Telcos have to deal with large volumes of data and require high levels of performance, reliability, and security, which places exacting requirements on the data management systems they deploy. The introduction of a wide range of new telecommunication services, based Internet use, mobility, multimedia and intelligent networks, will further increase these demands.

The P817 project examines the reasons for the increasing role and importance of data management for telcos. The technical requirements for data management technology in telcos are considered, using a specific application area, call-detail records, to illustrate these requirements.

The deployment of data management technology in telcos is complex and often greatly underestimated. By their nature, data management systems are among the most complex software systems, and it is a major task to deploy and manage them effectively. Significant resources and investments must be dedicated to the design, installation, maintenance, and further development of robust and flexible data management solutions. A key factor is a new generation of component-based data management solutions that can be incorporated seamlessly into current and future telecommunication systems.

The P817 project includes participants from KPN Research in the Netherlands (project leader), Telia Research AB in Sweden, Deutsche Telekom Berkom GmbH in Germany, Tele Danmark A/S in Denmark, Telefónica I+D in Spain, BT Laboratories in the United Kingdom, and Portugal Telecom SA in Portugal. The project receives EUR 1,000,000 from EURESCOM, the European institute for collaborative research and strategic studies in all areas of telecommunications. Currently, 24 operators from 23 European countries participate in EURESCOM. M. Böhlen (scientist-in-charge) and Li Chen from the database group participated as a sub-contractor to TeleDanmark with a budget of EUR 96,000.

3D Visual Data Mining, August 1999–July 2003

Visual data mining provides methods for accessing, analyzing, and visualizing large amounts of data. The 3DVDM project combines expertise in databases, statistics, visualization, and perceptual psychology with facilities for immersive real-time interactive visualizations. The 3DVDM project develops new data analysis methods that exploit the human perceptual faculties as far as possible in the search for unknown structures and relationship in large data sets. Database technology selects and provides the desired data subsets from the large database to be subjected to appropriate statistical processing. Using expertise from perceptual psychology, suitable processing methods from computer vision and scientific visualization are applied to create data structures amenable to visualization and adequate for visual perception.

The 3DVDM project exploits the facilities of the VR Center Nord, a cutting-edge research and development center for virtual reality that encompasses a Cave, a Panorama, and an auditorium with a PowerWall. It is part of a center for research, development, and training in interactive multimedia that was established at Aalborg University in January 1998. The data analyst has various interface controls, which allow for real-time feedback to the system for controlling the current data selection, the statistical and visual processing, and the observer's position and orientation relative to the data visualized. The observer may explore

the (static) data by navigating around in it, may passively observe an evolving world, and may navigate in an evolving world.

The 3DVDM project is interdisciplinary, with participants from the Department of Computer Science, the Department of Mathematical Sciences, the Faculty of Humanities, and the Department of Electronic Systems at Aalborg University. The project is funded with DKK 7,000,000 by the Danish Research Council. Current participants include M. Böhlen (principal investigator, scientist-in-charge, database systems), L. Bukauskas (database systems), S. L. Lauritzen (scientist-in-charge, statistical analysis), J. Parner (statistical analysis), P. S. Eriksen (statistical analysis), P. Mylov (scientist-in-charge, perceptual psychology), E. Granum (scientist-in-charge, scientific visualization), and H. R. Nagel (scientific visualization).

TimeCenter, 1995–

TIMECENTER was conceived in 1995 as an international consortium for the support of temporal database applications on traditional and emerging DBMS technologies. TIMECENTER serves as the context for most of the group's research in temporal databases.

A wide range of database applications manage time varying information. These include financial applications such as portfolio management, accounting, and banking; record-keeping applications, including personnel, medical-record, and inventory; and travel applications such as airline, train, and hotel reservations, and schedule management.

Recent advances in temporal query languages and data models clearly demonstrate that database applications that manage time-varying data may benefit substantially from built-in temporal support in the database technology used. In spite of this, and although temporal databases has been an active area of research for fifteen years, temporal database technology has so far had little impact on practice. TIMECENTER aims to contribute substantially to changing this state of affairs.

While others are pursuing valuable organizational solutions to the general problem of effective technology transfer, TIMECENTER focuses on technical solutions and aims at developing temporal database technology that is expressly transferable. Such technology allows for the continued operation of legacy code, for the harmonious coexistence of legacy and new application code and for the reuse of programmer expertise and knowledge. It also allows for maximal, effective, and efficient reuse of the functionality already provided by current database management systems. Together, these properties lead to a technology that is attractive to use and manageable to implement.

TIMECENTER has led to the local spin-off company, TimeChain, one member of the project has created his own company, TimeConsult, and members of TIMECENTER served as consultants to the temporal data warehouse Silicon

Valley start-up, *if...*, that has since been acquired by Amazon.com. The TIMECENTER report series and web pages make more than 50 technical reports and substantial accompanying software available to the public.

TIMECENTER is codirected by Christian S. Jensen and Richard Snodgrass. The unit's research that falls within TIMECENTER is funded by the Danish National Technical Research Council, by DKK 3,500,000. The current participants external to the unit include researchers at University of Arizona, Washington State University, Microsoft, TimeConsult, University of Bologna, University of Berne, Chungbuk National University, Flinders University, Amazon.com, University of California, Riverside, and University of California, Los Angeles.

5.2.3 Specific Projects

Aspects of Configuration Management

Lars Bendix conducted research on various aspects of configuration management.

Versioning models for software development has its focus on researching the possibilities to model data related to software development projects. Special emphasis is put on how to manage changes. Initial work in this field centered around a generalization of the workspace concept [48, 49]. To further support this project, subworkshops were organized at NWPER96 [7] and NWPER98 [131], which pinpointed certain problem areas and established some requirements to the data model. This led to the birth of the unified extensional versioning model, which is an alternative to the traditional selection based versioning models [47]. Recently another subworkshop has been organized at NWPER2000 with the purpose of investigating how Open Source Software development differ from traditional methods and whether or not traditional models for change management can be used [130].

Tools for modeling history in collaboration started its life developing a tool for versioning of hierarchical documents [158, 53]. This tool uses version control concepts and mechanisms to enable the collaboration between more people when authoring a shared document. However, its underlying model lacked capabilities for handling the history of how the document evolves at a sufficiently detailed level. Therefore, an attempt was made to combine the capability of VTML for fine-grained change tracking with CoEd's capability to handle structures and collaboration [157, 54].

In *General Object-Based Environments*, the potential of an object-based approach was investigated in contrast to traditional file-based approaches to the modeling of software development environments. Objects encapsulate files, making it possible to add attributes, relations, and actions to these files. It turns out that the integration is primarily facilitated by the presence of attributes and relations. Automation, on the other hand, is obtained through the actions that objects can carry out. The action mechanism is augmented by triggers, which are

well-known from object-oriented programming languages. This way the possibility to react on specific events in the programming process is obtained [50, 51].

The aim of *A taxonomy for Change Management* is to establish a model for how to carry out the change management function in software development organizations. A preliminary structure for change management was designed [155], where the function is subdivided into six defined sub-functions. This model was used to analyse and measure the change management capability of 15 Danish and Italian software development companies. This work led to a refinement of the previous model [156]. Recently the model has been augmented with a process that can be followed to continuously improve the change management capability of project groups and to spread this capability to the whole organization [52]. The model and the process is currently being tested in collaboration with a number of Danish software development companies.

DOPU

The Documentation Of Program Understanding (DOPU) project was initiated in 1998. Its starting points were two-fold: (1) The experiences from the HyperPro project in which we developed programming environments based on hypertext structures. (As such, the DOPU project can be seen as a continuation of the Hypertext and Software Engineering project described in the 1991–1995 evaluation). (2) Observations about practical problems with the application of literate programming. The main goal of this project is to come up with tools and techniques that allow the programmer to retain the program understanding as written accounts in close connection to the program.

As a concrete outset we have developed elucidative programming as an alternative to literate programming (represented by the so-called WEB tools.) Elucidative programming makes it possible to relate units in a program explanation with units in a program. This better enables the documentation of transverse program themes than does literate programming. An elucidative program is presented in two parallel panes that allow for mutual navigation between the documentation and the program. The physical proximity of explanations and program fragments, as emphasized by literate programming tools, is hereby replaced by navigational proximity in the elucidative programming tools. The development of elucidative programs is supported by a documentation-enabled programming environment.

We have developed a LAML-based elucidative programming environment for Scheme, and in an on-going student project—now in its second year—we work on an elucidative programming environment for Java.

The main research themes in this projects are:

- The challenges of constructing and maintaining elucidative programs.
- Internet presentation of elucidative programs.

- Exploration of new areas which will benefit from a close coupling between program explanations and the program fragments.
- Evaluation of elucidative program development processes, both within industry and educational settings.
- Identification of documentation patterns.

The DOPU project is documented through a number of papers on elucidative programming [22, 81, 82, 84]. Thomas Vestdam (Ph.D. student) works in the DOPU project.

DYNAMO

Being active from 1995 to 1998, the DYNAMO project continued an activity described in the research evaluation for 1991–1995. This project aimed to use dynamic modeling techniques as an alternative to static modeling in an object-oriented design context. The main hypothesis was that designers think in terms of objects and object relations, and consequently that it may be beneficial to develop dynamic models before static models. As a derived hypothesis, it was stated that parts of the static models can be derived automatically from the dynamic models.

The DYNAMO project contributed with an abstract language for dynamic modeling. This language contains a few new OOD language constructs. In the DYNAMO project, a number of tools were developed, including an interaction diagram editor that introduces a more accurate graphical notation than used in other similar tools. As an alternative approach, we developed the idea of design by animation. This branch of the work was done in cooperation with the students Lars Iversen and Per Madsen. Finally, we researched how to automatically derive static models from dynamic models.

The DYNAMO project is documented in several publications [23, 71, 79, 80, 147, 148, 149, 150, 151, 165].

gbeta

The **gbeta** project was initiated by Erik Ernst in 1995 at the University of Århus, Denmark, and it has been continued at this department since late 1999. The main themes of the project is programming language design and implementation, in particular as a continuation of the Scandinavian research tradition that produced Simula (thereby initiating object-orientation as such) and BETA. The language **gbeta** is a total reconstruction of the language BETA on a more general foundation, thus providing stronger abstraction mechanisms in many different ways. The **gbeta** project has produced an informal definition of the language and a full (but inefficient) implementation. One ongoing activity is to produce a

better implementation—which encompasses many novel challenges compared to the implementation of other languages, since this language, e.g., supports both static typing and dynamic specialization of objects. Another ongoing activity is an investigation of the potential for expressing aspect-oriented programming (and other, related approaches to advanced separation of concerns) as a built-in facility of the language itself, not as an external macro or meta-programming activity. This may be achieved simply by learning more about how to use the language, and it may also lead to adjustments of the language design. Finally, the `gbeta` type system provides a notion of dependent types that makes it possible to handle mutually recursive families of classes in a type-safe manner; in other type systems there is either a loss of reuse opportunities (which causes code duplication and thus maintenance problems) or a loss of type safety. It is also an ongoing activity to express the semantics of core language formally, thereby making such things as the core of the type system more accessible for researchers working with other kinds of languages.

About 7 graduate students collaborate with Erik Ernst on `gbeta` in their final-year projects.

Hypermedia Research

The research on Hypermedia systems was carried out by Uffe Kock Wiil. In the beginning of the period, this line of research focused on the HyperDisco prototype [113, 114]. Later in the period, a new component-based open hypermedia system named Construct became the focus of this line of work. Construct is a testbed (development environment) to experiment with different new ideas and concepts such as structural computing and multiple open services.

The Flag model [115] was developed as part of the work on standards and references models for open hypermedia systems. This model covers interoperability between different hypermedia systems. Another open hypermedia system reference model (OHRA) was also proposed to the research community. A set of standards for navigational hypermedia (linking) was developed in the context of the open hypermedia system working group (OHSWG). As part of this work, a series of workshops was organized with the goal of promoting the international research collaboration in the area of open hypermedia systems.

LAML

Kurt Nørmark initiated the Lisp-Abstracted Markup Language (LAML) project in January 1998. The main idea is to unify static Web authoring and (server side) dynamic Web programming by use of functional programming techniques. Both static Web documents (single as well as multi-paged) and dynamic Web documents are authored as functional programs using LAML libraries. Using the LAML approach, it is possible to express a wide variety of Web documents

and services in a single, powerful language, as opposed to using a mixture of markup languages and programming languages. This uniformity enables the use of abstraction and automated solutions anywhere in documents, both of which are essential in order to make solutions that scale.

We have developed software that produce a mirror of HTML in the programming language Scheme [134, 152]. Based on this mirror we have produced a variety of libraries, document styles, and tools, many of which support educational activities. The IDAFUS distance education environment and the LENO lecture note system are among the most substantial applications using the LAML libraries [83].

From a research perspective, the following themes are central in this project:

- The combination of programming languages and markup languages in Web development.
- The use of the functional programming paradigm for static Web authoring as well as for server side programming (using CGI, for instance).

5.3 Organization and Staff

The table below summarizes the staff during the evaluation period. In some cases the counts among the spring and fall semesters, in which cases two numbers are given. Blank cells indicate zero counts.

Although not apparent from the table, the unit has undergone substantial changes in its staff during the evaluation period. Aspects of these changes are discussed in Section 5.7, where the unit's own evaluation of the period is given.

	1996	1997	1998	1999	2000
Full Professors					1
Research Professors			1	1	1
Associate Professors	3	3 2	1 2	2	2 5
Research Associate Professors	1	1			
Assistant Professors	4 3	3 3	3 2	2 4	4 3
Ph.D. Students	3 4	6 7	7 8	7 8	7 8
Teaching Staff			1	2 1	2
Visiting Professors		1	2	1 2	1
Visiting Associate Professors			1 2		
Visiting Ph.D. Students					1 3
Visiting Scientists		1	1	1	1

Lars Bendix

Academic Degrees

- 1996 Ph.D. in Computer Science, Aalborg University
1986 Cand.scient. (M.Sc.) in Computer Science, University of Århus

Positions

- 1998– Associate Professor, Aalborg University, Department of Computer Science
1997–98 Visiting Associate Professor, Computer Science Department, University of Bologna (7 months)
1995–98 Assistant Professor, Aalborg University, Department of Mathematics and Computer Science
1994–95 Teaching Assistant, Aalborg University, Department of Mathematics and Computer Science
1993 Visiting Professor, Siegen University, Department of Computer Science (10 months)
1991–92 Lecturer in Information Technology, European Business School, Parma, Italy
1989–91 Ph.D. scholarship, University of Århus, Department of Computer Science (stationed at University of Pisa, Computer Science Department)
1987–89 Consultant, BIT Consulting, Struer
1986–87 Visiting Researcher, University of Pisa, Computer Science Department (9 months)

Michael H. Böhlen

Academic Degrees

- 1994 Ph.D. (Dr.sc.tech.), ETH Zürich, *Managing Temporal Knowledge in Deductive Databases*, Prof. R. Marti
1990 M.Sc. (Dipl. Informatik-Ing. ETH), ETH Zürich, *Ein Typenkonzept für ein deduktives Datenbanksystem*, Prof. R. Marti.

Positions

- 1998– Associate Professor in the Database Systems Group, Aalborg University, Department of Computer Science
- 1998 Research Associate Professor in the Database Systems Group, Aalborg University, Department of Computer Science (2 months)
- 1995–98 Assistant Professor in the Database Systems Group, Aalborg University, Department of Computer Science
- 1994–95 Postdoctoral Researcher in the TempIS Temporal Database Group headed by Prof. R. Snodgrass, University of Arizona, Department of Computer Science, Tucson, AZ, USA
- 1990–94 Teaching Assistant in the Knowledge Based Systems Group headed by Prof. R. Marti, Department of Computer Science, ETH Zürich, Switzerland
- 1990–94 Lecturer at the HTL Engineering School in Grenchen, Switzerland, courses taught on *information theory, automata theory, micro processors, structured programming, Jackson structured programming, data structures and algorithms, object oriented programming, and operating systems.*

Erik Ernst

Academic Degrees

- 1999 Ph.D. in Computer Science, University of Århus
- 1996 Cand.scient. (M.Sc.) in Computer Science, University of Århus
- 1989 Diploma in classical guitar, Royal Academy of Music at Århus
- 1988 Music Teacher (cls. guitar), Royal Academy of Music at Århus

Positions

- 2000–02 Assistant Professor, Department of Computer Science, Aalborg University
- 1999 Research Assistant Professor, Department of Computer Science, Aalborg University (2 months)
- 1999 Student Intern 2, Sun Microsystems Laboratories, Inc., Mountain View, CA, USA (7 months)
- 1998 Visiting Scholar, Department of Computer Science, University of Washington, Seattle, USA (3 months)
- 1995–99 Ph.D. scholarship, Department of Computer Science, University of Århus
- 1991–95 Project Assoc. Programmer, DEVISE, Department of Computer Science, University of Århus
- 1988–89 System Developer, TemaData, Kerteminde, Denmark (10 months)

Heidi Gregersen

Academic Degrees

- 1999 Ph.D. in Computer Science, Aalborg University
- 1995 Cand.scient. (M.Sc.) in Computer Science, Aalborg University

Positions

- 1999–2003 Assistant Professor, Aalborg University, Department of Computer Science
- 1997-98 Visiting Scholar, College of Computing, Georgia Institute of Technology, USA (9 months)
- 1996–99 Ph.D. Scholarship, Aalborg University, Department of Computer Science
- 1995–96 Teaching Assistant, Aalborg University, Department of Mathematics and Computer Science

Christian S. Jensen

Academic Degrees

- 2000 Dr.Techn., Aalborg University
- 1991 Ph.D. in Computer Science, Aalborg University
- 1988 Cand.scient. (M.Sc.) in Computer Science, Aalborg University

Positions

- 2000– Professor, Aalborg University, Department of Computer Science
- 1999 Visiting Professor, University of Arizona, Department of Computer Science (7 months)
- 1998–2003 Nykredit Research Professor, Aalborg University, Department of Computer Science (5 years; currently on a one-year leave)
- 1994–2000 Associate Professor, Aalborg University, Department of Computer Science (on leave since January 1998)
- 1994–95 Visiting Associate Professor, University of Arizona, Department of Computer Science (7 months)
- 1990–94 Assistant Professor, Aalborg University, Department of Mathematics and Computer Science
- 1991–92 Visiting Scholar, University of Arizona, Department of Computer Science (14 months)
- 1988–90 Ph.D. Scholarship, Aalborg University, Department of Mathematics and Computer Science (stationed at University of Maryland, Department of Computer Science)

Bent Bruun Kristensen

Positions

- 1997–98 External Associate Professor (Ekstern Lektor) in Computer Science, Aalborg University, Institute for Electronic Systems, Department of Computer Science
- 1997 Reading Professor (Docent) in Computer Science, Aalborg University, Institute for Electronic Systems, Department of Computer Science

Kurt Nørmark

Academic Degrees

- 1987 Lic.scient. (Ph.D.) in Computer Science, Aarhus University
- 1983 Cand.scient. (M.Sc.) in Computer Science, Aarhus University

Positions

- 1991– Associate Professor in Computer Science, Aalborg University, Institute of Electronic Systems
- 1987–91 Assistant Professor in Computer Science, Aalborg University, Institute of Electronic Systems
- 1984–86 Visiting Scholar at the Computer Science Department, Stanford University
- 1984–87 Ph.D. Student, Aarhus University

Torben Bach Pedersen

Academic Degrees

- 2000 Ph.D. in Computer Science, Aalborg University
- 1994 Cand.scient. (M.Sc.) in Computer Science, Århus University

Positions

- 2000– Associate Professor, Aalborg University, Department of Computer Science
- 2000 Lecturer, Aalborg University, Department of Computer Science (3 months)
- 1999 Visiting Scholar, Lawrence Berkeley National Laboratory (7 months)
- 1997– Industrial Ph.D. Fellow, Kommunedata
- 2000
- 1994–97 Database Administrator/Database specialist, Kommunedata
- 1990–93 Teaching Assistant, Department of Computer Science, Århus University

Dieter Pfoser

Academic Degrees

- 2000 Ph.D. in Computer Science, Aalborg University
- 1996 Magister (M.Sc.) in Business Informatics, University of Linz, Austria

Positions

- 2000– Assistant Professor, Aalborg University, Department of Computer Science
- 1997– Ph.D. Student, Aalborg University, Department of Computer Science
- 2000
- 1996–97 Ph.D. Student, University of Maine (USA), Department of Spatial Information Science and Engineering

Kristian Torp

Academic Degrees

- 1998 Ph.D. in Computer Science, Aalborg University
- 1994 Cand.polyt. (M.Sc.) in Computer Science, Aalborg University

Positions

- 2000– Database Specialist, Logimatic
- 1998– Research Assistant Professor, Aalborg University, Department of Computer Science (on leave since July 2000)
- 1996–97 Visiting Scholar, Department of Computer Science, University of Arizona (10 months)
- 1995–98 Ph.D. Scholarship, Aalborg University, Department of Computer Science
- 1994–95 Software Developer, Århus Stiftsbogtrykkeri
- 1993–94 Visiting Student, College of Computing, Georgia Institute of Technology, USA (9 months)

Nectaria Tryfona

Academic Degrees

- 1995 Ph.D. in Computer Engineering and Informatics, University of Patras
- 1991 Dipl.eng. in Computer Engineering and Informatics, University of Patras

Positions

- 2000– Associate Professor, Aalborg University, Department of Computer Science
- 1997– Assistant Professor, Aalborg University, Department of Computer Science
- 2000
- 1995–97 Post-doctoral Research Associate, National Center of Geographic Information and Analysis, Department of Spatial Information Science and Engineering, University of Maine, USA
- 1991–95 Research Engineer, Computer Technology Institute, Patras
- 1991–95 Ph.D. Scholarship, Computer Technology Institute, Patras

Uffe Kock Wiil

Academic Degrees

- 1993 Ph.D. in Computer Science, Aalborg University
- 1990 Cand.polyt. (M.Sc.) in Computer Science, Aalborg University

Positions

- 1996 Visiting Assistant Professor, University of California, Irvine, Department of Information and Computer Science (7 months).
- 1993–97 Assistant Professor, Aalborg University, Department of Mathematics and Computer Science
- 1991–92 Visiting Scholar, Department of Computer Science, Texas A&M University (1 year).
- 1990–93 Ph.D. Student, Department of Mathematics and Computer Science, Aalborg University.

5.4 Collaborations

5.4.1 Collaborations with Academia

Richard T. Snodgrass, University of Arizona. R. T. Snodgrass was a key collaborator for several members of the unit. The following are examples of this collaboration. He hosted Christian S. Jensen during a seven-month sabbatical in 1999, at which time Giedrius Slivinskas also visited. Prior to this, he hosted and co-advised Kristian Torp during a similar period. He co-directs TimeCenter with Christian S. Jensen. He has completed six journal papers, two for each of *ACM TODS* and *IEEE TKDE*, in addition to a number of conference papers with members of the unit. The reader is referred to the publications reported towards the end of this chapter.

Leo Mark, Georgia Institute of Technology. Leo Mark serves on the steering committee for Nykredit Center for Database Research. In addition, Leo Mark has served as a co-advisor for Kristian Torp and Heidi Gregersen during their stays at Georgia Institute of Technology, and he has been with the unit for a total of 7 months as a visiting professor. Most recently, he has collaborated with Janne Skyt and Christian S. Jensen. The collaboration has led to several papers, two of which have been published and two of which are in submission.

Curtis E. Dyreson, Washington State University. Curtis E. Dyreson has collaborated with primarily Michael H. Böhlen, Christian S. Jensen, and Torben B. Pedersen. He has been a visiting associate professor with the unit for 4 months, and he has most recently hosted Ole G. Jensen for ca. 6 months. The collaboration has so far led to papers in *ACM TODS* and *Information Systems*, in addition to papers in the conferences *VLDB* and *SSDBM*.

Klaus R. Dittrich, University of Zürich. Klaus R. Dittrich was a visiting professor for 1 month during Summer 1999. He lectured on object oriented databases and discussed issues related to middleware and heterogeneous data management with members of the unit.

Martin Erwig, Ralf H. Güting, and Markus Schneider, Fernuniversität Hagen. Ralf H. Güting was a visiting professor for 1 month during Summer 1998. All three have conducted research in moving object management with Michael H. Böhlen and Christian S. Jensen. This has led to one paper in *ACM TODS*, and a project proposal is underway.

James Clifford and Tomas Isakowitz, New York University. These researchers collaborated with Rick Snodgrass, Curtis Dyreson, and Christian S. Jensen on providing a framework for defining the semantics of *now*-relative data. The results are documented in an *ACM TODS* paper.

Rosanne Price, Monash University. Rosanne Price was a visiting Ph.D. student for 6 months during 2000, and has also visited for two shorter periods. She collaborates with Nectaria Tryfona and Christian S. Jensen on issues related to the conceptual modeling of spatial and spatio-temporal data in the context of geographic information systems. A journal and a conference paper document the collaboration.

Timos Sellis, National Technical University of Athens. Timos Sellis was a visiting professor for 2 months during June 1998 and December 1999. He has taught a Ph.D. course on spatial data management, and has advised the unit on its research directions and organization.

Jens Otto Sørensen, Århus School of Business. Jens Otto Sørensen spent a six-month sabbatical with the unit during Fall 1998. In addition to taking part in the teaching, he studied aspects of commercial database products.

Ouri Wolfson, University of Illinois. Ouri Wolfson was a visitor during July 2000. A panel for the *ICDE 2001* conference was designed, and issues related to moving object management were discussed. Work is in progress on the indexing of objects movement that is constrained by road networks. This collaboration involves Christian S. Jensen, Dieter Pfoser, and Simonas Šaltenis.

Bernhard Seeger, Philipps-University Marburg. Bernhard Seeger was a visiting professor for 1 month during August and September 2000. He collaborated with Simonas Šaltenis and Christian S. Jensen, and work is in progress that concerns the indexing of moving objects. In addition, he co-chairs the program committee of SSD'2001 with Christian S. Jensen.

Scott T. Leutenegger and Mario A. Lopez, Denver University, Colorado, USA. During 1999, Simonas Šaltenis and Christian S. Jensen collaborated with Scott T. Leutenegger and Mario A. Lopez on the indexing of the current and predicted future positions of moving objects. Simonas Šaltenis spent the first half of 1999 at Denver University. The collaboration resulted in a paper in the *ACM SIGMOD* conference.

Algimantas Juozapavicius, Kazimieras Mickus, Arturas Mazeika, and Mindaugas Pelanis, Vilnius University. Michael H. Böhlen, Linas Bukauskas, Christian S. Jensen, and Simonas Šaltenis collaborate with these researchers on 3D visual data mining, indexing, and index visualization. Christian S. Jensen is on the advisory committee for Mindaugas Pelanis' Ph.D. study.

Meral Ozsoyogly and Gultekin Ozsoyoglu, Case Western Reserve University. These researchers hosted Janne Skyt during a six-month stay in Spring 1999.

Ulf Asklund, Boris Magnusson, and Jonas Persson, Lund University, Sweden. Lars Bendix collaborated with these researchers on version models for software development and on change Management for Open Source Software. Four papers document this collaboration.

Henrik Bærbak Christensen, Aarhus University, Denmark. Lars Bendix collaborated with Henrik Bærbak on version models for software development. Three papers document this collaboration.

Paolo Ciancarini and Fabio Vitali, University of Bologna, Italy. Lars Bendix worked with Fabio Vitali on tools for modelling history in collaboration, resulting in two papers, and visited Paolo Ciancarini during a 7 months (July 1997 – January 1998) sabbatical period, working on teaching collaboration.

5.4.2 Collaborations with State Agencies

Poul Daugbjerg and Arne Simonsen, Kort & Matrikelstyrelsen. Nectaria Tryfona collaborates with these researchers on geodata modeling. Anders Friis-Christiansen is a Ph.D. student working on this topic with N. Tryfona as primary advisor and P. Daugbjerg and C. S. Jensen as secondary advisors.

Arie Shoshani and Junmin Gu, Lawrence Berkeley National Laboratory, USA. Torben B. Pedersen and Christian S. Jensen collaborate with these researchers from the Scientific Data Management Group, Lawrence Berkeley National Laboratory on federation of OLAP and object databases. Torben B. Pedersen visited the Lawrence Berkeley National Laboratory during January 1999 to August 1999, Two papers document the collaboration.

5.4.3 Industry Collaborations

John Bair and Michael D. Soo, Amazon.com. When with the company *if...*, John Bair collaborated with Michael H. Böhlen and Christian S. Jensen on the definition of new compatibility requirements to temporally extended query languages. Michael D. Soo was a visiting associate professor for 1 month during Summer 1998. He worked on temporal join performance studies in a collaboration that also encompasses researchers from the University of Arizona.

Yannis Theodoridis and Thanasis Hadzilacos, Computer Technology Institute, Athens, Greece. Christian S. Jensen, Dieter Pfoser, and Nectaria Tryfona collaborate with these researchers from the Applied Information Systems Group, CTI, Athens on issues related to data modeling and trajectories on moving point objects. Two journal papers, two conference papers and a European Union IST Project proposal is the outcome of the collaboration so far.

Theodore Johnson, AT&T. Theodore Johnson hosted Michael Akinde during the first half of 2000 at the At&T research lab in Florham Park, NJ. They collaborated on the extended multi-feature version of SQL suitable for OLAP. Two papers have resulted. The collaboration continues with the exploration of strategies for distributed aggregate processing.

Bjørn Skjellaug, SINTEF. Bjørn Skjellaug has collaborated with Michael H. Böhlen and Christian S. Jensen on the extension of temporal query languages based on statement modifiers to also encompass spatial dimensions. One paper has resulted. This collaboration continues, the focus now being to understand the fundamental differences among different types of “dimensional” query languages.

Nykredit Data. Nykredit is the largest sponsor of research activities in the unit. A number of staff collaborate with Nykredit on data warehousing, metadata management, and e-business. Most recently, the company has sponsored three two-year scholarships for international master’s students.

Kommunedata. Torben Bach Pedersen was an employee at Kommunedata in Århus during his Industrial Ph.D. (Danish: erhvervsforsker) project in the period 1997–2000, where Preben Etzerodt acted as his industrial advisor. During the project, Torben B. Pedersen collaborated with several Kommunedata employees on data warehousing, especially for clinical purposes. The collaboration resulted in Torben B. Pedersen’s Ph.D. thesis as well as seven papers.

Mindpass. Torben B. Pedersen and Christian S. Jensen collaborate with Mindpass on issues related to data warehousing for web portals, and serve as advisors on more general database issues. This collaboration involves patenting. The company has sponsored two two-year scholarships for international master’s students.

TimeChain. This company, founded by three former database students, is a spin-off of the temporal database research in the unit. Various members of the unit have informal interactions with this company.

Netnord and EUMAN. Christian S. Jensen serves as an advisor to these companies on issues related to the deployment of database technology.

Eastfork Object Space. Kurt Nørmark collaborates with Eastfork Object Space on documentation and Elucidative Programming.

5.5 Ph.D. Projects

During the evaluation period, 21 Ph.D. students have been associated with the database and programming technologies unit. Of these students, Eydun Eli Jacobsen, Palle Nowack, and Torben Worm completed their studies at the University of Southern Denmark. Five others, Chaouki Daassi, Arturas Mazeika, Kazimieras Mickus, Mindaugas Pelanis, and Rosanne Price, are enrolled at other

universities, but have substantial interactions with the unit. Of the remaining thirteen Ph.D. students, five have received their degrees from the university, and eight are enrolled at the university.

In the following, each Ph.D. project of a student enrolled at Aalborg University is described. The completed projects are listed first, ordered by the date the degree was awarded; then the projects in progress are listed, ordered by last name.

5.5.1 Project Descriptions

Configuration Management and Version Control Revisited

Name: Lars Bendix
Education: Cand.scient. (M.Sc.) in Computer Science, 1986
Duration: February 1988 – December 1995
Status: Degree awarded April 1996
Funding: Faculty of Natural Science, University of Århus
Advisor: Bent Bruun Kristensen

The dissertation addresses problems within configuration management and version control. An analysis shows shortcomings in present systems with respect to version control support in workspaces and proper integration of configuration management and version control. We provide full version control support in workspaces by moving version control from the level of whole repositories to the individual version groups, and by changing the checkout operation such that it creates a new version group in the workspace instead of a single component. Proper integration of configuration management and version control is done by considering dependency information as local to each single version of a component instead of as a global property of a version group, as is traditionally done. Therefore, we can model dynamically changing software architectures within one and the same model without suffering the traditional added complexity of performing explicit version control of configuration descriptions.

Temporally Enhanced Database Design

Name: Heidi Gregersen
Education: Cand.scient. (M.Sc.) in Computer Science, 1995
Duration: February 1996 – April 1999
Status: Degree awarded May 1999
Funding: Grant 9502695 from the Danish Technical Research Council
Advisor: Christian S. Jensen

The work concerns temporally improved design of databases managing time-varying information. The thesis surveys all temporally enhanced ER models known to the author. With the existing proposals, an ontological foundation, and novel requirements as its basis, a graphical, temporally extended ER model

is formally defined. The ontological foundation serves to aid in ensuring a maximally orthogonal design, and the requirements aim, in part, at ensuring a design that naturally extends the syntax and semantics of the regular ER model. The result is a novel model that satisfies an array of properties not satisfied by any single previously proposed model.

Because commercial database management systems support neither the ER model nor any temporal ER model as a model for data manipulation this thesis provides a two-step transformation from temporal ER diagrams, with built-in support for lifespans and valid and transaction time, to relational schemas. The first step of the algorithm translates a temporal ER diagram into relations in a surrogate-based relational target model; and the second step further translates this relational schema into a schema in a lexically-based relational target model.

Based on a framework from information systems engineering, the thesis evaluates the ontological expressiveness of three different temporal enhancements to the ER model. The evaluation considers the uses of the models for both analysis and design, and the focus is on how well the models capture temporal aspects of reality as well as of relational database designs.

Aspects of Data Modeling and Query Processing for Complex Multi-dimensional Data

Name: Torben Bach Pedersen
Education: Cand.scient. (M.Sc.) in Computer Science, 1994
Duration: February 1997 – March 2000
Status: Degree awarded April 2000
Funding: Danish Academy of Technical Sciences and Kommunedata A/S
Advisor: Christian S. Jensen

The project concerns aspects of data modeling and query processing for complex multidimensional data. Multidimensional databases are increasingly being used in Data Warehousing, On-Line Analytical Processing (OLAP), and other Business Intelligence applications. However, in many application domains, e.g., health care, where the cases in this project are taken from, the multidimensional data is too complex to be handled satisfactorily with current technology, creating a need for more powerful data management techniques.

Four topics are addressed. First, some of the new challenges that the area of health care informatics poses to Data Warehousing and OLAP technologies are presented. Second, based on real-world case studies, an extended multidimensional data model that improves over previous models by handling irregular dimension hierarchies, many-to-many relationships between facts and dimensions, imprecise data, change and time, and aggregation semantics, is defined. Third, a novel technique that supports the practical use of pre-aggregated data for query speed-up even when the dimension hierarchies are irregular is presented. Finally, the concepts and techniques underlying a system that allows queries over multidimensional data are presented.

mensional OLAP databases to easily reference data in external object databases are presented.

Issues in the Management of Moving Point Objects

Name: Dieter Pfoser
Education: Magister in Business Informatics, University of Linz, Austria, 1996.
Ph.D. student, University of Maine, USA, 1996–97
Duration: August 1997 – July 2000
Status: Degree awarded September 2000
Funding: TMR project Chorochronos
Advisor: Christian S. Jensen

The spatiotemporal application domain is vast, consequently many different types of prototypical database applications exist. This thesis focuses on issues related to applications involving moving point objects. We present a method to assess moving point objects data, propose a suitable representation in the form of trajectories of this data, and suggest prototypical queries. Trajectory data can stem from different movement scenarios such as unconstrained and constrained movement. The latter considers "infrastructure" to hinder the movement. For efficient query processing, we propose adapted and new access methods, as well as a new query processing technique that considers infrastructure in a filter step. An important query type is the join. In exploiting properties of the data, we propose a technique of how to incrementally compute this operation. The way we assess and represent moving point objects introduces uncertainty. Part of the thesis describes an augmented representation that considers uncertainty. We adapt the query processing techniques accordingly to consider uncertainty. A topic across chapters is the generation of data for the various experiments. We propose modifications to an existing spatiotemporal data generator and show how to create datasets that correspond to real-life situations.

Implementation Aspects of Temporal Databases

Name: Kristian Torp
Education: Cand.polyt. (M.Sc.) in Computer Science, 1994
Duration: August 1995 – July 1998
Status: Degree awarded October 1998
Funding: Department of Computer Science, Aalborg University
Advisor: Christian S. Jensen

This Ph.D. study concerned temporal relational database management systems. The temporal support for handling multiple versions of data is today typically implemented in an ad-hoc fashion in the application code. This support is implemented anew for each application being developed. Implementing temporal support in the application is time consuming and difficult using SQL-92

and conventional relational DBMSs. In contrast, a temporal DBMS significantly eases the management of temporal data: the query languages for such DBMSs have been enriched with high-level constructs that substantially simplify application code. With the clear benefits of temporal DBMSs, these systems should already be commercially available. However, this is not the case, in part because the temporal database research community has mostly focused on the design of temporal data models and query languages.

This Ph.D. study presented novel techniques for the implementation of significant aspects of temporally enhanced DBMSs. It is a fundamental requirement of these techniques that they be easily integrated into existing DBMS architectures.

The study considers how to extend a query language such as SQL-92 with temporal functionality with minimal implementation efforts, reusing the services of an existing DBMS. Three different meta-architectures for a layered implementation of temporal DBMSs are proposed. Each meta-architecture contains several specific layered architectures. The study has proposed techniques for how to correctly and efficiently time-stamp versions of data in the presence of transactions. The techniques are relevant to both layered and integrated architectures. The use of the temporal variable *now* makes it easier to model now-relative facts. The study defines the semantics of modifications involving *now*. Finally, an efficient implementation of the time-slice operator has been developed. This operator restores a previously current state of a temporal table and is used extensively in temporal DBMSs.

Data Warehouse Configuration

Name: Michael Akinde
Education: Cand.polyt. (M.Sc.) in Computer Science, 1998
Duration: August 1998 – July 2001 (expected)
Status: In progress
Funding: STVF
Advisor: Michael H. Böhlen

This project addresses the problem of data warehouse configuration. More specifically, the project investigates how to select a set of views to materialize in a data warehouse to achieve the optimal balance between query response time and update maintenance costs, given information about the base data and a set of queries to run on it. The existing body of work on this subject considers mostly simple aggregate queries over a single table; in this project we consider the problem for more complex queries involving multiple base tables, joins and selections in addition to aggregation. In addition, the existing work considering the data warehouse configuration problem for this body of queries is primarily theoretical. The project is concerned with creating models and algorithms to solve the data warehouse configuration problem for such complex queries, and to develop a prototype for these algorithms to prove that these algorithms work in

practice.

During the first half of the project, the model (view graphs) for materialized view selection have been developed and algorithms for the construction of these view graphs investigated. As the running time of the optimal view-selection algorithm is exponential with regards to the size of the graph, we have investigated methods to prune the size of the view graphs. This can be done by using heuristic algorithms to remove nodes from the graph that are unlikely to be used in the final view-selection. A substantial part of the view graph construction algorithm has been implemented into a working prototype. Currently, the project is focused on developing further optimization heuristics for the view graphs and view selection algorithms and implementing them in the prototype for performance studies.

Database Technology for Advanced Visual Data Mining

Name: Linas Bukauskas
Education: Cand.scient. (M.Sc.) in Computer Science, 1999
Duration: September 1999 – August 2002 (expected)
Status: In Progress
Funding: 3DVDM
Advisor: Michael H. Böhlen

Data mining and knowledge discovery in databases is a rapidly growing area that builds on the theories and technologies from many fields, including uncertainty modeling, pattern recognition, data warehousing, high performance computing, data visualization, and online analytical processing. Data mining aims to discover the relationship between the facts recorded in a database. Often information to be analyzed is represented in two dimensions (2D), which requires partial human interpretation. Data mining in three dimensions (3D) helps to better understand how the data changes in time and space. 3D visual data mining targets engineers, scientists, and companies with large databases, where the data can be viewed in many dimensions.

Most graphical systems employ 2D techniques, including geometric, icon-based, pixel oriented, hierarchical, and graph-based methods. Such tools use several logical screens to look at the data. To get an overall picture, the data analyst has to interpret partial information which is spread over several 2D screens. A dynamic 3D approach is a much better match for the interpretative power of our visual system. As of August, 1999 the Virtual Reality (VR) facilities that are available at the VR-Center Nord will make it possible to work with this kind of technology at Aalborg University. To create an immersive 3D visual world, which is suitable for data mining, we have planned to use the Cave, which is a part of the VR-Center Nord.

Database technology selects and provides the desired data subsets from the large database. The main challenge in designing a 3D data mining system is to

design and develop methods to dynamically extract the required 3D model from the data. New database structures and methods that map the database content to the 3D data model have to be designed. Developing and investigating such structures, methods, and algorithms is the goal of this Ph.D. project.

Geodata Modeling

Name: Anders Friis-Christensen
Education: Cand.scient. (M.Sc.) in Computer Science and Geography, 2000
Duration: June 2000 – June 2003 (expected)
Status: In progress
Funding: Danish Research Agency and the Danish Survey and Cadastre
Advisor: Christian S. Jensen and Nectaria Tryfona

This Ph.D. study aims at developing a framework needed for conceptual and logical modeling of geographic data. The study is to be carried out at The Danish National Survey and Cadastre (KMS), which is responsible for various different datasets and registers e.g. the Danish Cadastre, the Danish Topographic Map and the Danish Charts. At KMS these datasets are maintained separately and treated as different products and there is a need to work towards a method which has basis in the real world phenomena and objects rather than in specific products. This is because the same real world objects from different products are handled and maintained differently in each product. KMS has launched a research project: 'General Reference for Spatial Information', the objective of which is to establish a standardised framework for spatial registration in general.

The Ph.D. study is a component of this project and will focus on different aspects of geodata modeling based upon the requirements for modeling KMS data. These aspects could include the following: Interoperability issues, quality information of objects, unique object identifiers and generalisation of objects. The main goal of the Ph.D. study is to develop methods and rules for modeling geographic data, which could be implemented in an application modeling tool.

Tuple Versioning

Name: Ole Guttorm Jensen
Education: Cand.polyt. (M.Sc.) in Software Engineering, 1998
Duration: August 1998 – July 2001 (expected)
Status: In progress
Funding: STVF
Advisor: Michael H. Böhlen

The project investigates the problem of schema evolution in populated databases. Databases are often modified and many modifications result in changes to the database structure. In stark contrast, applying schema changes to a populated database is still an open problem. The main difficulty is that after a

schema change some data no longer fits the schema. Resolving this mismatch is a challenge, and a semantically clean and simple solution has yet to emerge.

The project takes a general and practical approach to the problem of schema evolution. During the first half of the project, the foundation for schema evolution was developed, including selective schema changes and a new formal framework for tuple versioning. In tuple versioning, each tuple in the database is associated with a conceptual schema and a recorded schema, respectively. With tuple versioning, mismatches are faithfully recorded and selectively resolved during query processing, data manipulations, and schema modifications. Moreover, tuple versioning avoids data migration, where data is migrated from an old schema version to a new at a significant performance cost. To implement tuple versioning, evolution trees were proposed. Currently, a prototype implementing evolution trees is under development to evaluate and measure the performance of different storage solutions and query processing mechanisms for tuple versioning. On going work also includes a generalization of tuple versioning for semi-structured data.

Indexing Techniques for Continuously Evolving Phenomena

Name: Simonas Šaltenis
Education: Cand.scient. (M.Sc.) in Computer Science, 1998
Duration: August 1998 – July 2001 (expected)
Status: In progress
Funding: Nykredit Center for Database Research
Advisor: Christian S. Jensen

Recent years have shown an increase in the amount of data stored in databases as well as an increase in the diversity of the data and applications that employ database systems technology. Temporal, spatial, and spatiotemporal data are prominent examples of new kinds of data that modern DBMSs should contend with. This project addresses indexing of such data. More specifically, the project focuses on applications where data is continuously evolving. In conventional databases, data is assumed to be constant unless it is explicitly modified, whereas in new application areas, there is often a need to support continuously evolving data - data that changes continuously as a function of time, even without being explicitly updated.

In spatiotemporal applications, there is often a need to record continuously changing spatial information about objects that move and/or change their shape. Examples of such objects include vehicles, mobile computers, mobile telephones, PDA's (personal digital assistants), military equipment, people, forest fires, hurricanes, to name but a few. Specifically, the project focuses on how to index the current positions of continuously moving spatial objects, with the ability to extrapolate these positions into the future. So called time parameterized R-trees are proposed for this purpose. In addition, we plan to investigate the indexing of the history of the evolution of spatial objects. To address continuous motion,

such index must support interpolation between explicit updates of the database.

In temporal and spatiotemporal databases, continuously evolving data occurs naturally when we want to support data related to the continuously progressing current time, termed now-relative data. In such applications, there is a need to index time-intervals that change continuously. An index was developed to index now-relative spatiotemporal data, where now-relative time intervals are associated with the spatial extents of objects. For example, the history of changing boundaries of land parcels can be indexed using the proposed technique.

Managing Aging Data

Name: Janne Skyt
Education: Cand.scient. (M.Sc.) in Computer Science, 1995
Duration: June 1998 – May 2001 (expected)
Status: In progress
Funding: Department of Computer Science, Aalborg University
Advisor: Christian S. Jensen

The Ph.D. project aims to make it possible for append-only databases to support physical deletion, termed vacuuming. Doing this the project specifically moves to develop concepts and techniques for retaining, as much as possible, the accountability and trace-ability requirements of such databases. A wide range of real-world database applications, including financial and medical applications, are faced with these requirements leading to the use of transaction-time databases where logical deletions are implemented as insertions at the physical level, append-only databases that retain all previously current states and are ever-growing.

Many solutions have been proposed for a variety of challenges regarding transaction-time databases. However, the support for vacuuming has received precious little attention, even though it is called for by, e.g., the laws of many countries. Although necessary, with vacuuming, the database's previously reliable recollection of the past may be manipulated via, e.g., selective removal of records pertaining to past states. Based on these opposite requirements this project seeks to define the concept of vacuuming for append-only databases to serve as a solid foundation for the correct and user-friendly processing of queries and updates against vacuumed databases.

Append-only databases are also prominent in data warehousing where, by Inmon's definition, all its data exhibit temporal dimensions. This provides an environment where the concept of vacuuming will be applied. Additional challenges appear with the data's change in "value" and validity, with limited access patterns, and with integrating different storage media.

Temporal Query Optimization

Name: Giedrius Slivinskas
Education: Cand.scient. (M.Sc.) in Computer Science, 1998
Duration: August 1998 – July 2001 (expected)
Status: In progress
Funding: STVF
Advisor: Christian S. Jensen

The project investigates how to add support for temporal (time-referenced) data to database management systems (DBMSs). The main focus is set on optimization and processing of temporal queries. Such queries do not run efficiently in current DBMSs and are difficult to write, therefore temporal data is managed in user applications, which become hard to develop and maintain. A temporal DBMS can be implemented as a layer between user applications and the conventional DBMS, e.g., Oracle. The layer translates temporal query language statements to SQL and sends them to the DBMS; it is also enhanced by query processing capabilities, since complex temporal operations such as temporal aggregation often run faster in the layer than in the DBMS. To efficiently and effectively divide the processing between the layer and the DBMS, temporal query optimization and processing mechanisms are needed.

During the first half of the project, the foundation for temporal query optimization was developed, including an algebra for temporal query representation and a comprehensive set of transformation rules. The algebra enhances existing relational algebras based on multisets by integrating the handling of order and adding temporal support, and the transformation rules are divided into different types according to how they deal with duplicates, order, and time periods. By capturing duplicate removal and retention and order preservation for all queries, as well as coalescing for temporal queries, the foundation formalizes and generalizes existing approaches. A reasonable number of operations and transformation rules have been implemented in Volcano extensible query optimizer, which is being used for performance studies. The current task of the project is to investigate optimization-related issues such as costing and selectivity of temporal operations, and optimization heuristics.

Documentation of Program Understanding – Tools, Experiments and Patterns

Name: Thomas Vestdam
Education: Cand.polyt. (M.Sc.) in Computer Engineering, 1999
Duration: 2000 – 2003 (expected)
Status: In progress
Funding: Aalborg University
Advisor: Kurt Nørmark

Through documentation of program understanding, we wish to maintain the

thoughts and rationales behind a program. These are present during the development process but as time passes they will disappear. Without this knowledge it is difficult for future users to maintain or reuse software. Instead of developing reverse engineering tools, good tools supporting the programmer in maintaining program explanations and in addressing program-code have been created (and will be further developed). Hence, future users can re-gain program understanding by reading program-code addressing explanations written in human language.

DOPU (Department of Computer Science, Aalborg University¹) is aimed at improving the quality of software by means of better habits, techniques and tools for documentation of program understanding. The main outcome so far is a documentation tool called the Elucidator. The tool supports users as described above. However, the tool has not yet been tested in realistic situations and consequently it is unknown whether it meets the actual needs of potential users.

We wish to test and further develop the Elucidator. This includes observing users applying the tool in various software projects, observing maintainers using documentation, and adapting the tool to the needs of the users. We intend to discover and develop a broader application of the tool throughout the entire software development process. Furthermore, through observation of the production and use of documentation we envision that it is possible to identify reoccurring situations that call for a specific solution. Hence we hypothesise that, documentation patterns can be constructed and used by programmers, as guidelines of how to document program understanding.

5.6 Services

This section lists services performed in full or in part during the evaluation period.

5.6.1 Program Committees

- Tutorial program committee, *2001 ACM SIGMOD International Conference on the Management of Data*, Madison, WI, May, 2001 (Christian S. Jensen).
- *Nordic Workshop on Programming Environment Research*, 2000 (Lars Bendix, Kurt Nørmark).
- *International Conference on Extending Database Technology*, Konstanz, Germany, March 27–21, 2000 (Michael Böhlen).
- *Twenty-sixth International Conference on Very Large Data Bases*, Cairo, Egypt, September, 2000 (Christian S. Jensen).

¹<http://dopu.cs.auc.dk/>

- *International Workshop on Temporal, Spatial and Spatio-Temporal Data Mining*, Lyon, France, September, 2000 (Christian S. Jensen).
- *International Workshop on Advanced Spatial Data Management*, Greenwich, UK, September, 2000 (Christian S. Jensen).
- *Twelfth International Conference on Scientific and Statistical Database Management*, Berlin, Germany, July, 2000 (Christian S. Jensen).
- *Seventh International Workshop on Temporal Representation and Reasoning*, Cape Breton, Nova Scotia, Canada, July, 2000 (Christian S. Jensen).
- *Twelfth International Conference on Software Engineering and Knowledge Engineering*, Chicago, IL, June, 2000 (Christian S. Jensen).
- *Twenty-sixth International Conference on Very Large Data Bases*, Cairo, Egypt, September, 2000 (Christian S. Jensen).
- *ACM International Symposium on Advances in Geographic Information Systems*, 2000 (Nectaria Tryfona).
- *International Database Engineering and Applications Symposium*, 2000 (Nectaria Tryfona).
- *International Workshop on Emerging Technologies for Geo-Based Applications*. 2000 (Nectaria Tryfona).
- *International Conference on Spatial Information Theory*, 2000 (Nectaria Tryfona).
- *Geographical Domain and Information Systems - EuroConference on Ontology & Epistemology for Spatial Data Standards*, France, September 2000 (Nectaria Tryfona).
- *International Database Engineering and Applications Symposium*, Montreal, Canada, August 1999 (Michael Böhlen).
- *Twenty-fifth International Conference on Very Large Data Bases*, Edinburgh, September 1999 (Michael Böhlen).
- *Workshop on Databases in Telecommunications*, Edinburgh, September 6, 1999 (Michael Böhlen).
- *First International Workshop on Evolution and Change in Data Management*, Versailles, France, November, 1999 (Christian S. Jensen).
- *Workshop on Spatio-Temporal Data Models and Languages*, Florence, Italy, September, 1999 (Christian S. Jensen).

- *Eleventh International Conference on Scientific and Statistical Database Management*, Cleveland, Ohio, USA, July, 1999 (Christian S. Jensen).
- *Eleventh Conference on Advanced Information Systems Engineering*, Heidelberg, Germany, June 1999 (Christian S. Jensen).
- *Fifteenth IEEE International Conference on Data Engineering*, Sydney, Australia, March 1999 (Christian S. Jensen).
- Scientific committee member, *Fourth International Conference on Spatial Information Theory*, Stade, Germany, August 1999 (Christian S. Jensen).
- International scientific committee member, *3rd IMACS International Multi-conference on Circuits, Systems, Communications and Computers*, Athens, Greece, July 1999 (Christian S. Jensen).
- *ACM International Symposium on Advances in Geographic Information Systems*, 1999 (Nectaria Tryfona).
- *Workshop on Spatio-Temporal Database Management*, 1999 (Nectaria Tryfona).
- *International NSF Workshop on Integrated Spatial Databases*, 1999 (Nectaria Tryfona).
- *Nordic Workshop on Programming Environment Research*, 1998 (Lars Bendix, Kurt Nørmark).
- *Thirteenth Brazilian Symposium on Databases*, Maringa, Brazil, October 1998 (Christian S. Jensen).
- *Twenty-fourth International Conference on Very Large Data Bases*, New York, New York, August 1998 (Christian S. Jensen).
- *Tenth Conference on Advanced Information Systems Engineering*, Pisa, Italy, June 1998 (Christian S. Jensen).
- *Sixth International Conference on Extending Database Technology*, Valencia, Spain, March 1998 (Christian S. Jensen).
- *ACM International Symposium on Advances in Geographic Information Systems*, 1998 (Nectaria Tryfona).
- *Third International Conference on Spatial Information Theory*, Pittsburgh, PA, October 1997 (Christian S. Jensen).
- *Twenty-third International Conference on Very Large Data Bases*, Athens, Greece, August 1997 (Christian S. Jensen).

- *Thirteenth IEEE International Conference on Data Engineering*, Birmingham, England, April 1997 (Christian S. Jensen).
- Scientific committee member, *Dagstuhl Seminar on Temporal Databases*, Schloss Dagstuhl, Germany, June 23–27, 1997 (Christian S. Jensen).
- *ACM International Symposium on Advances in Geographic Information Systems*, 1997 (Nectaria Tryfona).
- *Nordic Workshop on Programming Environment Research*, 1996 (Lars Bendix, Kurt Nørmark).
- *Fifth International Conference on Information and Knowledge Management*, Rockville, MD, USA, November 1996 (Michael Böhlen).
- *International Workshop on Logic in Databases*, San Miniato, Italy, July 1996 (Michael Böhlen).
- *Fifth International Conference on Extending Database Technology*, Avignon, France, March 1996 (Christian S. Jensen).
- *Twelfth IEEE International Conference on Data Engineering*, New Orleans, Louisiana, February/March 1996 (Christian S. Jensen).
- *ACM International Symposium on Advances in Geographic Information Systems*, 1996 (Nectaria Tryfona).
- Member of the executive committee of ECOOP'95, August 1995 (Kurt Nørmark).

5.6.2 Conference Related Services

- Organizer of subworkshop, *Nordic Workshop on Programming Environment Research*, 2000 (Lars Bendix).
- European Coordinator, *2000 ACM SIGMOD International Conference on the Management of Data*, Dallas, TX, May 14–19, 2000 (Christian S. Jensen).
- Program committee co-chair, with Michel Scholl, *Workshop on Spatio-Temporal Database Management*, Edinburgh, Scotland, September 10–11, 1999 (a VLDB'99 co-located event) (Christian S. Jensen).
- Technical chair, *Workshop on Spatio-Temporal Database Management*, Edinburgh, Scotland, 1999 (G. Slivinskas).
- Proceedings chair, *Workshop on Spatio-Temporal Database Management*, Edinburgh, Scotland, 1999 (Michael Böhlen).

- Organizer of subworkshop, *Nordic Workshop on Programming Environment Research*, 1998 (Lars Bendix).
- Organizers, *First Chronos Intensive Workshop*, Aalborg, June 19–20, 1998 (Christian S. Jensen, Michael Böhlen, Dieter Pfoser, and Nectaria Tryfona).
- Organizers, *Tenth National Danish Database Workshop*, Aalborg, June 15–16, 1998 (Christian S. Jensen, Janne Skyt).
- Program committee vice chair, *Fourteenth IEEE International Conference on Data Engineering*, Orlando, FL, April 1998 (Christian S. Jensen).
- Best paper and student paper awards committee member, *Fourteenth IEEE International Conference on Data Engineering*, Orlando, FL, April 1998 (Christian S. Jensen).
- Tutorial chair, *1997 ACM SIGMOD International Conference on the Management of Data*, Tucson, AZ, May 1997 (Christian S. Jensen).
- Proceedings Chair, *Nordic Workshop on Programming Environment Research*, 1996 (Lars Bendix).
- Organizers, *Nordic Workshop on Programming Environment Research 1996*, Aalborg, (Kasper Østerbye, Lars Bendix, and Kurt Nørmark).
- Panel participant, Panel 1: Challenges and Opportunities in Scientific and Statistical Databases (moderator: Yannis Ioannidis), *Eleventh International Conference on Scientific and Statistical Database Management*, Cleveland, Ohio, July 1999 (Christian S. Jensen).
- Panel participant, panel on “Future Research Directions for Spatial, Temporal, and Spatiotemporal Databases,” Dagstuhl Seminar on *Integrating Spatial and Temporal Databases*, Schloss Dagstuhl, Germany, November 1998 (Christian S. Jensen).
- Panel participant, panel on “The State-of-the-Art in Temporal Data Management: Perspectives from the Research and Financial Applications Communities”, *International Workshop on Temporal Databases*, Zürich, Switzerland, September 1995 (Michael Böhlen).
- In addition, members of the unit have served as session chairs at nine conferences.

5.6.3 Editorial Services

- Editor, *IEEE Transactions on Knowledge and Data Engineering* (subject area: temporal databases). May 1997–present (Christian S. Jensen).
- Editorial Board Member, *ACM SIGMOD Digital Review*, May 1999–present (Christian S. Jensen).
- Coordinator and head of three editorial boards for the ongoing Consensus Temporal Database Glossary Initiative, March 1992–present. (Christian S. Jensen; co-editors include J. Clifford, C. E. Dyreson, R. Elmasri, S. K. Gadia, P. Hayes, S. Jajodia, A. Segev, R. Snodgrass, and A. Tansel).
- Editor for the Temporal Database Benchmark Initiative, March 1993 onwards. More than twenty researchers are contributing (Christian S. Jensen).

5.6.4 Reviewing

Journals/Publishers

- *Software - Practice and Experience* (Lars Bendix).
- *Nordic Journal of Computing* (Lars Bendix).
- *Journal of Intelligent Information Systems* (Michael Böhlen).
- *Data & Knowledge Engineering Journal* (Heidi Gregersen).
- *ACM SIGMOD Record* (Christian S. Jensen).
- *ACM Transactions on Database Systems* (Michael Böhlen, Heidi Gregersen, Christian S. Jensen).
- *The Computer Journal* (Michael Akinde, Christian S. Jensen).
- *IEEE Transactions on Knowledge and Data Engineering* (Michael Böhlen, Heidi Gregersen, Christian S. Jensen, Nectaria Tryfona).
- *Information and Software Technology* (Christian S. Jensen).
- *Information Systems* (Christian S. Jensen).
- *The International Journal on Very Large Data Bases* (Michael Böhlen, Christian S. Jensen).
- *Journal of Database Management* (Christian S. Jensen).
- Morgan Kaufmann, San Mateo, California, USA (Christian S. Jensen).

- Prentice Hall, Englewood Cliffs, New Jersey, USA (Christian S. Jensen).
- *Theory and Practice of Object Systems* (Christian S. Jensen).
- *IEEE Personal Communications* (Dieter Pfoser).
- *Journal of Information and Software Technology* (Simonas Šaltenis).
- *Communications of the ACM* (Nectaria Tryfona).
- *Geoinformatica* (Nectaria Tryfona).
- *International Journal on Geographic Information Systems* (Nectaria Tryfona).
- *Journal of Database Management* (Nectaria Tryfona).

External Reviews

Members of the unit have served as external reviewers for more than two dozen conferences.

5.6.5 Other Academic Services

Invited Talks/Demos

- “Temporal Databases,” Center for Linguistics, Aalborg University, October 31, 2000 (Christian S. Jensen).
- “Themes and Challenges in Temporal Databases,” MobIS’2000, Siegen, Germany, October 12, 2000 (Christian S. Jensen).
- “Mobile E-Services and Business Intelligence,” ITudveksling, Nouhauz, Aalborg, September 28, 2000. (Christian S. Jensen).
- “Management of Temporal Data—with Focus on the Support for ‘now’,” Seminar on Temporal and Spatial Data Management, Kort & Matrikelstyrelsen, Copenhagen, September 7, 2000 (Christian S. Jensen).
- “Indexing Moving Objects: A Database Perspective,” National Science Foundation/Army Research Office Workshop on Motion Algorithms, Duke University, August 7, 2000 (Christian S. Jensen).
- “ChoroChronos Research,” Workshop on High Performance Object Databases, Cardiff, Wales, July 20, 2000 (Christian S. Jensen).
- “Business Intelligence,” MMT Program, Aalborg University, Aalborg, Denmark, June 17, 2000 (Christian S. Jensen).

- “Indexing Structures for Temporal and Spatio-Temporal Data,” DHI Water & Environment, Hørsholm, Denmark, May 30, 2000 (Christian S. Jensen).
- “Temporal Data Management,” Oracle, Ballerup, Denmark, May 30, 2000 (Christian S. Jensen).
- “Indexing of On-Line Moving Objects,” University of Cardiff, Wales, May 26, 2000 (Christian S. Jensen).
- “Indexing of On-Line Moving Objects and Other M-Challenges,” Keynote Speech, Korean Database Conference, Korea, May 12, 2000 (Christian S. Jensen).
- “Research in TimeCenter,” Chungbuk National University, Cheongju, Chungbuk, Korea, May 9, 2000 (three talks) (Christian S. Jensen).
- “Data Warehousing,” Nykredit Corporation, Copenhagen, Denmark, April 10, 2000 (Christian S. Jensen).
- “Data Warehousing,” Nykredit Data Corporation, Aalborg, Denmark, April 4, 2000 (Christian S. Jensen).
- “Indexing the Present and Future Positions of Moving Objects,” IT University of Copenhagen, Copenhagen, Denmark, March 15, 2000 (Christian S. Jensen).
- “Designing and Developing Geo-Applications: Models and Tools,” September 7, 2000. Kort & Martikelstyrelsen (Nectaria Tryfona).
- “The Role of Ontologies in Spatial Data Mining,” Geographical Domain and Geographical Information Systems EuroConference on Ontology and Epistemology for Spatial Data Standards La Londe-les-Maures, France, September 2000, (Nectaria Tryfona).
- “Indexing Continuous Movement,” Department of Electrical Engineering and Computer Science, University of Illinois at Chicago, IL, USA, December 15, 1999 (Christian S. Jensen).
- “Temporal Databases,” The Society of Danish Engineers, Copenhagen, Denmark, November 11, 1999 (in Danish) (Christian S. Jensen).
- “Management of 1-D Sequence Data—From Discrete to Continuous” (an opponent’s presentation at the Ph.D. defense of Ling Lin), Linköping University, Department of Computer and Information Science, Linköping, Sweden, March 11, 1999 (Christian S. Jensen).

- Demo of Tiger and Jungle, NATO Technical Visit, Aalborg University, October 6, 1998 (Michael Böhlen).
- “A New Index for General Bitemporal Data,” Norwegian University of Science and Technology, Department of Computer and Information Science, Trondheim, Norway, December 3, 1998 (Christian S. Jensen).
- “Database Support for Time-Related Data,” Selskab for Informationsteknologi, Ingeniørforeningen i Danmark, and Nordjysk Informatikforum, Aalborg, Denmark, November 5, 1998 (in Danish) (Christian S. Jensen).
- “A New Index and its Datablade,” Informix user group meeting, Copenhagen, Denmark, June 9, 1998 (Christian S. Jensen).
- “Experiences of a Researcher,” CIT-organized Ph.D. workshop, Radisson SAS Hotel, Århus, Denmark, May 12, 1998 (in Danish) (Christian S. Jensen).
- “Spatiotemporal Database Design: A Methodological Approach,” Department of Computer Science, University of Genova, Italy, October 1998 (Nectaria Tryfona).
- “Methods and Tools for Efficient GIS Modeling,” Informatics Laboratory, Agricultural University of Athens, Greece, December 1998 (Nectaria Tryfona).
- “Designing Databases With Time-oriented Data,” Computer Laboratory, University of Cambridge, UK, October 22, 1997 (Christian S. Jensen).
- “Designing Temporal Relational Databases—Normal Forms and Decomposition Rules based in Attribute Semantics,” Department of Management Information Systems, University of Arizona, November 8, 1996 (Christian S. Jensen).
- “Temporal Database Design,” National Database Workshop, Department of Mathematics and Computer Science, Aalborg University, June, 1996 (Christian S. Jensen).
- “Database Design for Spatial Applications: Tools and Techniques,” School of Engineering and Applied Science, The George Washington University, Washington DC, USA, October 1996 (Nectaria Tryfona).
- Demo, with A. Steiner, of the TimeDB System, International Workshop on Temporal Databases, Zürich, Switzerland, September 1995 (Michael Böhlen).

External Courses/Tutorials

- Presenter of a half-day tutorial on Temporal Databases, at the *Twelfth Brazilian Symposium on Databases*, Fortaleza, Brazil, October 1997 (Christian S. Jensen).
- Presenter, with R. Snodgrass, of a half-day tutorial on Temporal Databases, at the *Twelfth International Conference on Data Engineering*, New Orleans, Louisiana, February/March 1996 (Christian S. Jensen).

Evaluations for Funding Organizations

- *The Danish Center for IT Research* (Christian S. Jensen).
- Long Term Research proposals, *European Commission, DG III - F2* (Christian S. Jensen).
- The National Science Foundation (USA) (Christian S. Jensen).
- The Swedish Research Council for Engineering Sciences (Christian S. Jensen).

Evaluation Committees

- External examiner for Marlon Dumas' Ph.D. examination, Computer Engineering Laboratory, IMAG Institute, University of Grenoble, 2000 (Christian S. Jensen).
- External opponent for Ling Lin's Ph.D. examination, Linköping University, Department of Computer and Information Science, 1999 (Christian S. Jensen).
- Member of an evaluation committee for a full professor position, 1998, University of Bergen, Department of Information Science (Christian S. Jensen).
- Review board member (w. Arie Shoshani and Gio Wiederhold), External Quality Assessment of Scientific Research, evaluation of the Database Laboratory, Department of Computer Science, Swiss Federal Institute of Technology Lausanne, 1998 (Christian S. Jensen).
- External examiner for Thomas Zurek's Ph.D. examination, University of Edinburgh, Department of Computer Science, 1997 (Christian S. Jensen).
- Member of an evaluation committee in charge of appointing one Associate Professor in Databases and Data Structures, Odense University, Department of Mathematics and Computer Science, 1997 (Christian S. Jensen).

Professional Organizations

Researchers in the unit are members of various national and international organizations, such as ACM, SIGMOD, SIGKDD, SIGPLAN, IEEE, IEEE TCDE, DANIT, ERCIM, EDRG, AGILE, ZobIS, Datalogiforeningen, and the Greek Chamber of Engineers.

5.7 The Unit's Own Evaluation

Separate self-evaluations are provided for the Database Systems (DB) and Programming Systems groups. These groups formulated separate research plans for the evaluation period and were only merged into the present DP unit towards the end of the evaluation period.

5.7.1 Database Systems

At the start of the evaluation period, the DB group consisted of one (tenured) associate professor, two (untentured) assistant professors, one teaching staff awaiting entry into a Ph.D. position, and one staff already in a Ph.D. position. Below, we revisit the overall and specific goals set by this group of staff during Fall 1995, and we consider their fulfillment.

The group aimed at establishing a strong international focus, obtaining good visibility in the general database community and very high visibility in temporal databases. Three means were identified: publication in recognized international outlets, participation in activities in the international research community, and cooperation with established research institutions.

The group has exceeded its own expectations in relation to all three means. Much of the research activity in the group has indeed concerned aspects of temporal databases, and it is felt that the TIMECENTER project has played a significant role in exposing the group's research in temporal databases. It is very pleasing that Richard Snodgrass—known as “Mr. Temporal Databases” and a prominent and highly regarded figure in the database community—in his most recent book on temporal databases predicates the group as the world's strongest in temporal databases².

As indicated in the chapter's introduction, the research in the database group has a substantial constructive and experimental component. This means that much of the research involves the development of very substantial software, which perhaps is not immediately clear from the descriptions elsewhere in this chapter. It is felt the the rigor imposed by prototype implementation of artifacts and,

²“[...] the world's strongest temporal database research group at the University of Aalborg in Denmark,” p. xxi in R. T. Snodgrass, *Developing Time-Oriented Applications in SQL*, in J. Gray, series editor, Morgan Kaufmann Series in Data Management Systems, Morgan Kaufmann Publishers 2000, 504+xxiii pages.

frequently, the subsequent experimental evaluation has had a very positive effect on the group's research.

The group is particularly proud that it has succeeded in publishing much of its research in journals such as *ACM Transactions on Database Systems*, *IEEE Transactions on Knowledge and Data Engineering*, and the *VLDB Journal* (nine papers in total), as well as in conferences such as *ACM SIGMOD*, *EDBT*, *IEEE ICDE*, and *VLDB* (fourteen papers in total). In addition to these and other well-reputed general outlets, the group has published in pertinent specialized outlets. Also some of the research that relates to geographic information systems has been published in prominent outlets in that area. This record places the group well among other European and American database groups.

Further, the group takes pride in its record of service to the research community, which has been quite comprehensive. The specific projects undertaken during the evaluation period have fostered quite substantial cooperation.

The next goal was based on the observation that database research was in its infancy in Denmark. This was in contrast to the state of affairs in the surrounding countries—e.g., Germany, Norway, and Sweden—and the widespread use of database technology in industry. The group declared its commitment to take active part in activities aimed at changing this state of affairs.

This goal is a prominent part of the objectives of the Nykredit Center for Database Research (NDB; please refer to its description in Section 5.2.2). NDB has so far produced four Ph.D.'s. The database group currently consists of one professor, three associate professors, two assistant professors, seven Ph.D. students, two visiting scientists, several external Ph.D. students with close ties to the group, and frequent short-term visitors. This tripling of the size of the group represents in itself a significant step towards a stronger database research presence at Danish universities. On the negative side, the production of ten Ph.D.'s in NDB, many of them international, far from constitutes any guarantee that other database groups will emerge at other Danish universities. In addition, the larger size has enabled the group to interact more with industry and students.

The presence of a good database student population, in terms of quality and quantity, was recognized as an important condition for the functioning of the group. Because the database topics in the curriculum were reduced towards the end of the previous evaluation period, the group anticipated a need for other activities to counter this development.

The group has mainly attempted to maintain its student population by offering inspiring and relevant project topics. It is found that the group has attracted sufficient numbers of students, many of which have been very capable and well motivated.

The group has also taken active part in establishing two new master's degree programs, each with a significant database component. Specifically, the group has taken the initiative to establish an international master's degree in knowledge and data engineering. This was done jointly with the Decision Support Systems

unit. Next, the group has taken part in establishing the master's in advanced data analysis.

The final overall goal was to ensure critical mass in the group. The group was given an upper bound on its size of two permanent faculty. As a result, it was felt that it was important to establish research collaboration among the group members, thus ensuring that the group would not split into subgroups. Given the growth of the group that has occurred, the issue of critical mass has been less pressing during the evaluation period than anticipated.

In addition to the overall goals, specific research agendas were outlined in three related areas: data warehousing, temporal data management, and spatio-temporal database management. The member of staff behind the agenda for the first area left the group early on. However, the remaining and new staff have pursued research in data warehousing. These activities are characterized in several of the Ph.D. program descriptions. The research in temporal databases is described in several Ph.D. program descriptions and in the description of the TIMECENTER project. Finally, the research in spatio-temporal databases is covered in the several Ph.D. program descriptions and the description of the CHOROCHRONOS project.

5.7.2 Programming Systems

According to the plans made at the beginning of the period, we would organise our work in groups of one or two persons, allowing for external collaborations. We would continue to publish our research results in high-quality international conferences and (to some extent) international periodicals covering our research areas. We should also continue the international interaction and exchange of ideas with researchers working with similar subjects, and extend this further in the direction of cooperation on experiments, joint articles, and prolonged visits to international research groups. Furthermore, we would aim to establish better contacts with the industry, to obtain a better interaction with the industry, in order to promote our theoretical results and to get inspiration and feedback from practical reality. Several specific individual or two-person plans were sketched for the future.

Bent Bruun Kristensen, Kasper Østerbye, and Uffe Kock Wiil left the Department of Computer Science during the period. For this reason most of the specific plans were not carried out—or carried out after the people had left the group. Bent Bruun Kristensen and Kasper Østerbye both contributed with factual information to this report, but contributed no qualitative evaluations. Uffe Kock Wiil's work in the period followed the directions that were initially planned: Design, development and deployment of hypermedia system prototypes and development of standards and references models for open hypermedia systems.

Kurt Nørmark's completed his work on dynamic models in object-oriented design during the period. The work on hypertext and software engineering has

been redirected to research on documentation of program understanding and elucidative programming.

Lars Bendix' work on versioning and history models within configuration management has led to strengthened research collaboration with other groups in Denmark, Sweden and Italy. After some initial difficulties his work on a taxonomy for change management is starting to create some response, both from academia for its theoretic aspects and from industry for its practical aspects.

The remaining elements of the plan have all been met. We have had fruitful exchange of ideas internationally—and one occasion of a prolonged visit. Several of these collaborations have resulted in co-authored articles. Contact to industry has improved during the period. Collaborations are going on at both a formal level in the form of joint research projects and at a more informal level as contacts and exchange of ideas, problems, and results.

5.8 Research Plan 2001–2005

The research plan for the coming period includes an umbrella project that aims to integrate the research activities in the unit, as well as other projects to which the different members of the unit have made commitments. We proceed to cover the general research profile, the umbrella project, and then the non-umbrella projects.

5.8.1 Research Profile

The DP unit is the result of merging two quite different research groups. In part because of this, but also because of its increased size, the unit is quite diverse with respect to many aspects of research.

An important theme for the next period will be to manage the diversity imposed by the merger of two previously separate groups and the emerging diversity caused by factors such as the increased size of the unit and its members' needs for individual identities. The goals are to ensure that the members of the unit benefit from the diversity and that the diversity becomes a strength to the unit as a whole.

On the one hand, it is important to ensure that the members representing different approaches interact in a manner that leads to cross-fertilization and the creation of the synergism that characterize critical mass. On the other hand, it is desirable to preserve diversity in the unit. Diversity is a strength in itself and is also important in relation to the broad area of teaching assigned to the unit.

The members of the unit consider the coming period an experiment that will reveal whether or not it is a good idea to have database and programming researchers in the same unit.

As an attempt to meet these goals, the general research theme of *Mobile Internet Services* has been identified as a single theme within which most or all staff members can conduct some of their research. The theme in itself is seen as having no inherent biases towards database or programming aspects; rather, it offers ample challenges for all. The idea is that all staff members will conduct some of their research within this theme, in addition to other research. Some staff members may at first work on this theme with students.

5.8.2 Mobile Internet Services

Continued advances in hardware technologies combine to create a new research area, termed Mobile Internet Services (hereinafter “M-Services”). Below, a general description of this area is given first. Then follows descriptions of four sample challenges in the area.

General Description Several trends in hardware technologies combine to provide the enabling foundation for M-Services. These trends include continued advances in *miniaturization* of electronics technologies, *display devices*, and *wireless communications*. Other trends are the improved *performance* of general computing technologies and the general improvement in the *performance/price ratio* of electronics hardware. Perhaps most importantly, positioning technologies, including GPS (global positioning system) and those embedded in the wireless service infrastructure, are becoming increasingly accurate.

It is expected that the coming years will witness very large quantities of on-line (i.e., Internet-worked), position-aware, wireless objects capable of continuous movement. Examples of such objects include consumers using WAP-enabled mobile-phone terminals and PDAs (personal digital assistants), tourists carrying on-line and position-aware “cameras” and “wrist watches,” vehicles with computing and navigation equipment, etc.

The development described above paves the way to a range of qualitatively new Internet-based services, including *traffic coordination and management*, *location aware advertising*, *integrated tourist services*, *safety related services*, and *position varying information in industrial environments*.

One single generic scenario may be envisioned for such services. Moving objects use services that involve location information. The objects disclose their positional information (position, speed, velocity, etc.) to the services, which in turn use this and other information to provide specific functionality.

Each service maintains a log of the requests made to it, a so-called “click-stream,” and uses this for analyzing user interaction with the service. The service accumulates data derived from the click-streams and integrates this with other customer data in a so-called data warehouse, which is a very large repository of integrated information that may be used for data analysis.

The data in the data warehouse is used for mass-customization of the service, so that each user receives a service customized to the user's specific situation, preferences, and needs, e.g., dynamic, user-specific web-page content. In addition, the warehouse is used for delayed modification of the services provided, and for longer-term strategic decision making. Business intelligence techniques such as on-line analytical processing and data mining are used for these purposes.

The integration of location information into this scenario has received little attention and offers a number of fundamental challenges. Common to these challenges is the task of extending techniques that work well for static data to support dynamic, continuously evolving data. We proceed to describe four specific challenges in M-Services.

Data Mining for Continuously Changing Data Data mining techniques provide automatic knowledge discovery in large databases, e.g., by finding associations between data values, clustering data into high-level groups, classifying data, and predicting unknown data values. Data mining has been applied with great success in both industrial and scientific settings. Current data mining techniques are applicable only to static data, and cannot easily be applied to dynamic, constantly changing data. It is a significant research challenge to identify methods for data mining that can be applied to dynamic data.

Data mining techniques have traditionally performed all computations directly on base data, but a recent trend is the use of precomputed data to speed up the mining process. The extension of data mining techniques to use precomputed data for faster processing represents a non-trivial research challenge.

Indexing Continuous Movement Indexing is a fundamental technique in data management, as it makes it possible to locate desired data items in very large databases efficiently. Indexing techniques are becoming increasingly important, as the improvement in the rate of transfer of data between disk and main memory cannot keep pace with the growth in disk capacities and processor speeds.

Traditional indexing techniques work only for static data, meaning that the indices have to be explicitly updated when changes occur to the data. For large, continuous datasets, e.g., those capturing the positions of moving objects, the constant updating of the indices would require very large computing resources, rendering the use of indices either impractical or totally impossible. Or, alternatively, the large volumes of updates and the mechanisms that regulate the concurrent use of the indices would combine to block the querying of these structures, also rendering them useless.

It should be investigated how it is possible to obtain the normal benefits of indexing when the data being indexed change continuously.

Internet Programming M-Services are expected to call for a high degree of programming flexibility at both the server side and the client side. In particular, it must be expected that server-side as well as client-side dynamics will become important characteristics of M-Services. The LAML approach to Internet programming holds the potential to support both sides using a single and very flexible functional programming language. As a natural consequence, the umbrella project will integrate LAML-based solutions, and hereby serve as an additional testbed for the uniform use of functional programming across the boundary between a Web server and its clients.

Precomputation and Distributed Querying in Data Warehousing Precomputation is an essential technique in warehousing, where it is used to give fast answers to queries that involves aggregating very large amounts of data. For example, the total use of a service by county and month could be precomputed and stored. Reusing this result enables fast answers to queries that ask for the number of uses of the service, e.g., by month alone, by county alone, or by quarter and county in combination. This is because the answers may be derived from the precomputed result alone; access to the bulks of data in the data warehouse is not needed.

However, precomputation has traditionally assumed static data, meaning that precomputation has problems that are to some extent similar to those of traditional indexing techniques. A fundamental problem is then how to apply precomputation when continuous change occurs in the data.

Next, it is a serious impediment to the effective analysis of data available from Internet sources that existing tools for analytical processing require the data to be available in a single, centralized data warehouse. The inherently distributed nature and potentially huge size of data renders this centralized approach impractical.

The natural solution is to maintain a distributed data warehouse, where data is retained in a local data warehouse at its natural location. In order for such a solution to be practical, technologies for the efficient distributed processing of analytical queries are needed.

5.8.3 Non-Umbrella Projects

This section contains research plans for areas that are not directly connected to the umbrella project. This research is primarily a continuation of research performed in the current evaluation period.

Configuration Management The activities carried out by Lars Bendix under the theme “Aspects of Configuration Management” will continue. Emphasis will be on the practical aspects of applying concepts and principles that have been

established. Feedback for improvement will be obtained through experiments. Furthermore, the subject is broadened to cover not just configuration management, but also change management in general. Finally, these aspects will be integrated in the curriculum of the courses taught by the unit.

Elucidative Programming and LAML Technology Kurt Nørmark wants to continue his work on elucidative programming. The emphasis will be on both educational and industrial applications of the tools and the ideas. The aim is to extend the current industrial collaboration in the area of elucidative programming in the coming period. In addition Kurt Nørmark wants to continue the work on LAML technology for Web authoring and Web server-side programming. Existing international contacts in the area of Web programming are expected to be used actively during the coming period.

gbeta The **gbeta** language and implementation will be developed towards a level of maturity where it can be used for large-scale programming. This is very important in order to be able to unfold and experience the large-scale consequences of the fundamental generalization in basic language concepts that characterizes the transition from BETA to **gbeta** (as mentioned on page 40, **gbeta** was developed as a generalized version of BETA). There was also a profound increase in expressive power in the transition from Simula to BETA, and it has given rise to many years of very productive work to explore the richer universe of BETA, compared with the previously known possibilities and techniques of Simula.

There are numerous non-trivial research challenges in the development of such a mature **gbeta** system; for instance, it requires novel techniques to support dynamic specialization of objects and representation of virtual attributes in a time and space efficient manner. Finally, it would be valuable to enter into research collaboration with people who are focusing on prototyping, information systems, or other topics where the language could be put into various usage contexts, and these other people could benefit from the malleability of a language with an in-house implementation and design.

Temporal, Spatial and Spatio-Temporal Databases The research activities that have been going on during the reporting period concern temporal, spatial, or spatio-temporal databases, but not M-Services. Many of these quite diverse activities will continue during the coming period.

In conceptual modeling, activities will continue to focus on the provision of temporally and spatially enhanced data models. Goals include the development of better means of capturing temporal aspects of constraints and spatial relationships. A Ph.D. project funded in part by the Danish National Survey and Cadastre aims to improve the support for geographic data modeling in UML.

Issues include supporting both field-based and entity-based data, multiple and evolving object roles, attached data-quality information, and generalization.

In relation to languages, ongoing research explores the provision of language facilities for physical deletion in append-only databases that otherwise only support logical deletions. We also study means of implementing such facilities.

In query processing, focus is on the support for temporally enhanced query languages via light-weight and adaptable so-called stratum and middleware architectures. In these architectures, special “temporal” software accepts temporal queries and co-processes these jointly with an underlying DBMS that is treated as a black box. This software tries to discover the performance capabilities of the underlying system for various types of queries by monitoring its response times. This enables the software to co-process queries increasingly better.

3D Visual Data Mining Visual data mining provides methods for accessing, analyzing, and visualizing large amounts of data. We want to develop new data analysis methods that exploit the human faculties in the search for unknown structures in large data sets. Specifically, we want to investigate the potential of immersive 3D visualization facilities as provided by a Panorama or Cave. We want to develop new database technology to support observer-relative data processing and want to develop navigation tools to support data mining in immersive 3D environments.

5.8.4 Additional Considerations

This section covers considerations in relation to the research plan that were not covered in the previous sections.

Some of the research conducted in the unit may be of interest to industrial partners, and may benefit from the involvement of such partners. In the coming period, the unit does not aim to increase the quantity its collaboration with industry, but will instead focus on improving the quality of its collaboration with industry. Specifically, the unit will prioritize longer-term and substantial collaborations, and the unit will try to seek out industrial partners that may be able to use the unit’s research results in their products.

The umbrella project represents an attempt at establishing more collaboration with local industry in the “mobile information management” area, e.g., the many “mobile” companies in the area. The majority of the funding will continue to come from traditional research projects, but the unit will try to attract some funding directly from industry.

The unit anticipates a level of research for the coming period that is similar to that of the previous period. Partly as a result of this, the volume of publication per staff member is also expected to remain at the same level as in the previous period. The unit will continue to publish in the prominent journals in its area; and

it aims to publish some parts of its research in prestigious, general conferences, while publishing other parts in pertinent, specialized conferences and workshops.

An important aspect of doing research is good laboratory facilities. To support research and teaching, the unit would like to establish several laboratories. These include a “Mobile Information Management” and a “Business Intelligence” laboratory that both support research in the umbrella project, and a possible “Programming and Program Documentation” laboratory. All of the laboratories include dedicated hardware and software, as well as support personnel. Sufficient space for the laboratories is critical to their success.

The unit believes that it will be a big and essential challenge, and perhaps a problem, to attract staff and Ph.D. students in the coming period. The main reason for this is probably the good job market in both academia and industry for individuals with M.Sc. or Ph.D. degrees in computer science, which makes it easy for such individuals to find interesting and well-paid jobs. The salary gap between university and industry makes it challenging to attract new staff members, especially from industry. It should be noted that Danish Ph.D. salaries are competitive seen in a global perspective, for which reason the lack of funding is likely the factor that limits the recruiting of international Ph.D. students.

After the drain of people from the previous programming systems group, an assistant professor and a Ph.D. student have been hired in 2000. If funding can be secured, an additional Ph.D. student is ready to start in mid-2001. This secures some coverage of programming. (In comparison, database topics are enjoying better coverage at the moment.) The unit is committed to hiring additional staff members with profiles within all the areas spanned by database and programming technologies. This includes new areas as well as those covered by the current staff. The unit will be particularly happy to receive more applicants for positions within the programming area.

An important means of attracting new staff members is to offer a first-class research environment. It is very important to put focus on consolidating and further improving the existing research environment, for at least the following to specific reasons:

- Performing research must be attractive to the staff in comparison to other activities.
- There will be increased competition for staff with other Danish and international universities and industry.

Some problems with providing a good research environment must be overcome. Specifically, we are worried about the consistent teaching overload. The imbalances in the assignments of resource-intensive courses and less intensive project advising is also seen as a problem for some faculty.

5.9 Bibliography

The aim in compiling the bibliography was to document as best as possible the work conducted and results obtained during the evaluation period. As a result, papers that document work done fully or in part during the evaluation period are reported. This includes papers in submission and to appear. Consequently, there is some overlap with papers reported in the evaluation report for 1991–1995, which also followed this policy.

Extended versions of some conference publications were later published in journals, which implies some overlap between the two categories of papers. In addition, a substantial portion of the technical reports (including departmental, TimeCenter, and Chorochronos reports) are extended versions of papers that were later published in conferences and journals.

Finally, it is noted that very substantial program code has been developed that documents the results obtained during the evaluation period. This code is currently not included in the bibliography.

Theses

- [1] L. Bendix. *Configuration Management and Version Control Revisited*. Ph.D. thesis, Department of Computer Science, Aalborg University, 1996. Available as Technical Report R-96-2016, Department of Computer Science, Aalborg University.
- [2] H. Gregersen. *Temporally Enhanced Database Design*. Ph.D. thesis, Department of Computer Science, Aalborg University, 1999.
- [3] C. S. Jensen. *Temporal Data Management*. Dr.Techn. thesis, Aalborg University, 1328+ xxxviii pages, August 1999. Available from <http://www.cs.auc.dk/~csj/Thesis>.
- [4] T. B. Pedersen. *Aspects of Data Modeling and Query Processing for Complex Multidimensional Data*. Ph.D. thesis, Department of Computer Science, Aalborg University, 2000.
- [5] D. Pfoser. *Issues in the Management of Moving Point Objects*. Ph.D. thesis, Department of Computer Science, Aalborg University, 2000.
- [6] K. Torp. *Implementation Aspects of Temporal Databases*. Ph.D. thesis, Department of Computer Science, Aalborg University, 1998.

Edited Proceedings

- [7] L. Bendix, K. Nørmark, and K. Østerbye, editors. *Proceedings of the 1996 Nordic Workshop on Programming Environment Research*. Department of Mathematics and Computer Science, Institute of Electronic Systems, Aalborg University, May 1996. Published as Technical Report R-96-2019.
- [8] M. Bernstein, L. Carr, and K. Østerbye, editors. *The ACM Hypertext'97 Conference*, University of Southampton and Association for Computing Machinery, Southampton, UK, April 1997. ACM Press.
- [9] M. H. Böhlen, C. S. Jensen, and M. Scholl, editors. *Spatio-Temporal Database Management: Proceedings of the International Workshop on Spatio-Temporal Database Management*, volume 1678 of *LNCS*, Edinburgh, Scotland, September 1999. Springer Verlag.

Refereed Journal Articles

- [10] J. Bair, M. Böhlen, C. S. Jensen, and R. T. Snodgrass. Notions of Upward Compatibility of Temporal Query Languages. *Wirtschaftsinformatik*, 39(1):25–34, February 1997.
- [11] M. Böhlen, C. S. Jensen, and R. T. Snodgrass. Temporal Statement Modifiers. *ACM Transactions on Database Systems*, 25(4), December 2000, 47 pages, to appear.
- [12] J. Clifford, C. Dyreson, T. Isakowitz, C. S. Jensen, and R. T. Snodgrass. On the Semantics of ‘Now’ in Databases. *ACM Transactions on Database Systems*, 22(2):171–214, June 1997.
- [13] H. Gregersen and C. S. Jensen. Temporal Entity-Relationship Models—a Survey. *IEEE Transactions on Knowledge and Data Engineering*, 11(3):464–497, May 1999.
- [14] R. H. Güting, M. Böhlen, M. Erwig, C. S. Jensen, N. Lorentzos, M. Schneider, and M. Vazirgiannis. A Foundation for Representing and Querying Moving Objects. *ACM Transactions on Database Systems*, 25(1), March 2000, 37 pages, to appear.
- [15] T. Hadzilacos and N. Tryfona. Evaluation of Database Modeling Methods for Geographic Information Systems. *Australian Journal of Information Systems*, 6(1):15–26, 1998.
- [16] C. S. Jensen and R. T. Snodgrass. Semantics of Time-Varying Information. *Information Systems*, 21(4):311–352, 1996.

- [17] C. S. Jensen and R. T. Snodgrass. Temporal Data Management. *IEEE Transactions on Knowledge and Data Engineering*, 11(1):36–44, January/February 1999.
- [18] C. S. Jensen, R. T. Snodgrass, and M. D. Soo. Extending Existing Dependency Theory to Temporal Databases. *IEEE Transactions on Knowledge and Data Engineering*, 8(4):563–582, August 1996.
- [19] W. Kreutzer and K. Østerbye. BetaSIM—a Framework for Discrete Event Modelling & Simulation. *Simulation Practice and Theory*, 6(6):573–599, 1998.
- [20] B. B. Kristensen and K. Østerbye. Roles: Conceptual Abstraction Theory and Practical Language Issues. *Theory and Practice of Object Systems*, 2(3):1–18, 1996.
- [21] B. B. Kristensen and K. Østerbye. Roles: Conceptual Foundation and Practical Usage. *Theory and Practice of Object Systems*, 2(3):143–160, 1996.
- [22] K. Nørmark. Elucidative Programming. *Nordic Journal of Computing*, 7(2):87–105, 2000.
- [23] K. Nørmark, L. Iversen, and P. Madsen. Animation and Presentation Tools for Object-Oriented Design. *Nordic Journal of Computing*, 6(1):110–128, 1999.
- [24] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. A Foundation for Capturing and Querying Complex Multidimensional Data. *Information Systems* (special issue on data warehousing), 2001, 42 pages, to appear.
- [25] D. Pfoser and Y. Theodoridis. Generating Semantics-Based Trajectories of Moving Objects. *Computers, Environment and Urban Systems*, 2001. To appear.
- [26] R. Price, N. Tryfona, and C. S. Jensen. Extended Spatiotemporal UML: Motivations, Requirements, and Constructs. *Journal of Database Management*, 11(4):13–27, October/December 2000.
- [27] G. Slivinskas, C. S. Jensen, and R. T. Snodgrass. A Foundation for Conventional and Temporal Query Optimization Addressing Duplicates and Ordering. *IEEE Transactions on Knowledge and Data Engineering* (special section with extended versions of the best papers from ICDE’2000), 2001, 37 pages, to appear.

- [28] K. Torp, L. Mark, and C. S. Jensen. Efficient Differential Timeslice Computation. *IEEE Transactions on Knowledge and Data Engineering*, 10(4):599–611, July/August 1998.
 - [29] K. Torp, C. S. Jensen, and R. T. Snodgrass. Effective Timestamping in Databases. *The VLDB Journal*, 8(4):267–288, February 2000.
 - [30] N. Tryfona. Constraint-based Spatial Data Mining. *International Journal of Computers, Environment, and Urban Systems*, 2000. To appear.
 - [31] N. Tryfona and C. S. Jensen. Conceptual Modeling for Spatiotemporal Applications. *Geoinformatica*, 3(3):245–268, October 1999.
 - [32] T. Vestdam. Documentation Threads – Presentation of Fragmented Documentation. *Nordic Journal of Computing*, 7(2):106–125, 2000.
 - [33] Uffe K. Wiil and John J. Leggett. Hyperform: A Hypermedia System Development Environment. *ACM Transactions on Information Systems*, 15(1):1–31, January 1997.
 - [34] K. Østerbye and W. Kreutzer. Synchronization Abstraction in the BETA Programming Language. *Computer Languages*, 25(3):165–187, 2000.
- In submission:**
- [35] J. Chomicki, D. Toman, and M. H. Böhlen. Querying ATSQL Databases with Temporal Logic. Submitted to *ACM Transactions on Database Systems*, 26 pages, March 1999.
 - [36] H. Gregersen and C. S. Jensen. Conceptual Modeling of Time-Varying Information. Awaiting submission.
 - [37] H. Gregersen, L. Mark, and C. S. Jensen. Mapping Temporal ER Diagrams to Relational Schemas. Awaiting submission.
 - [38] D. Pfoser and C. S. Jensen. Querying the Trajectories of On-Line Mobile Objects. Submitted to *IEEE Personal Communications Magazine*.
 - [39] S. Šaltenis and C. S. Jensen. R-Tree Based Indexing of General Spatio-Temporal Data. Submitted to *The VLDB Journal*, 32 pages.
 - [40] J. Skyt, C. S. Jensen, and L. Mark. A Foundation for Vacuuming Append-Only Databases. Submitted to *Transactions on Data and Knowledge Engineering*, 35 pages.
 - [41] K. Torp, C. S. Jensen, and R. T. Snodgrass. Modification of Now-Relative Databases. Submitted to *The VLDB Journal*.

Conference Articles

- [42] M. Agesen, M. H. Böhlen, L. O. Poulsen, and K. Torp. A Split Operator for Now-Relative Bitemporal Databases. In *Proceedings the 17th International Conference on Data Engineering*, 10 pages, April 2000. To appear.
- [43] M. Akinde and M. Böhlen. Constructing GPSJ View Graphs. In *Proceedings of the International Workshop on Design and Management of Data Warehouses*, pages 8 1–12, June 1999.
- [44] M. Akinde, D. Chatziantoniou, T. Johnson, and S. Kim. The MD-join: An Operator for Complex OLAP. In *Proceedings the 17th International Conference on Data Engineering*, 10 pages, April 2001. To appear.
- [45] M. O. Akinde, O. G. Jensen, and M. H. Böhlen. Minimizing Detail Data in Data Warehouses. In *Advances in Database Technology - EDBT'98 Sixth International Conference on Extending Database Technology*, volume 1377 of *LNCS*, pages 293–307, March 1998.
- [46] J. Andersen, A. Giversen, A. H. Jensen, R. S. Larsen, T. B. Pedersen, and J. Skyt. Analyzing Clickstream Using Subsessions. In *DOLAP 2000, ACM Third International Workshop on Data Warehousing and OLAP*, pages 25–32, November 2000.
- [47] U. Asklund, L. Bendix, H. B. Christensen, and B. Magnusson. The Unified Extensional Versioning Model. In *Proceedings of the 9th International Symposium on System Configuration Management*, volume 1675 of *LNCS*, pages 100–122. Springer Verlag, September 1999.
- [48] L. Bendix. Fully Supported Recursive Workspaces. In *Proceedings of the 6th International Workshop on Software Configuration Management*, volume 1167 of *LNCS*, pages 256–261, Berlin, Germany, March 1996. Springer Verlag.
- [49] L. Bendix. How To Obtain Version Control in Workspaces Too. In *Proceedings of the CONF'96 Software Conference*, Paris, France, June 1996.
- [50] L. Bendix. Experiences with an Object-Based approach to Configuration Management. In *Position papers for the 7th International Workshop on Software Configuration Management*, pages 90–99. May 1997.
- [51] L. Bendix. General Object-Based Environments: An Engine for Software Development Environments. In *Proceedings of the Eighth Nordic Workshop on Programming Environment Research*, pages 111–126, June 1998.

- [52] L. Bendix. Continuous Improvement of the Configuration Management Process. In *Proceedings of the conference Views on Software Development in the New Millennium*, pages 50–55, August 2000.
- [53] L. Bendix, P. N. Larsen, A. I. Nielsen, and J. L. S. Petersen. CoEd - A Tool for Versioning of Hierarchical Documents. In *Proceedings of the 8th International Symposium on System Configuration Management*, volume 1439 of *LNCS*, pages 174–187. Springer Verlag, July 1998.
- [54] L. Bendix and F. Vitali. VTML for fine-grained change tracking in editing structured documents. In *Proceedings of the 9th International Symposium on System Configuration Management*, volume 1675 of *LNCS*, pages 139–156. Springer Verlag, September 1999.
- [55] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. R-tree-based Indexing of Now-Relative Bitemporal Data. In *Proceedings of the 24th International Conference on Very Large Databases*, pages 345–356, New York City, NY, August 1998.
- [56] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. Light-Weight Indexing of Bitemporal Data. In *Proceedings of the Twelfth International Conference on Scientific and Statistical Database Management*, pages 125–138, Berlin, Germany, July 2000.
- [57] R. Bliujūtė, S. Šaltenis, G. Slivinskas, and C. S. Jensen. Systematic Change Management in Dimensional Data Warehousing. In *Proceedings of the Third International Baltic Workshop on DB and IS*, pages 27–41, Riga, Latvia, April 1998.
- [58] R. Bliujūtė, S. Šaltenis, G. Slivinskas, and C. S. Jensen. Developing a DataBlade for a New Index. In *Proceedings of the Fifteenth IEEE International Conference on Data Engineering*, pages 314–323, Sydney, Australia, March 1999.
- [59] M. Böhlen, J. Chomicki, R. Snodgrass, and D. Toman. Querying TSQL2 Databases with Temporal Logic. In P. Apers, M. Bouzeghoub, and G. Gardarin, editors, *Advances in Database Technology — EDBT'96, 5th International Conference on Extending Database Technology*, volume 1057 of *LNCS*, pages 325–341. Avignon, France, March 1996.
- [60] M. Böhlen, C. S. Jensen, and B. Skjellaug. Spatio-Temporal Database Support for Legacy Applications. In *Proceedings of the 1998 ACM Symposium on Applied Computing*, pages 226–234, Atlanta, Georgia, February/March 1998.

- [61] M. H. Böhlen, L. Bukauskas, and C. E. Dyreson. The Jungle Database Search Engine. In *Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data*, pages 584–586, Philadelphia, Pennsylvania, June 1999 (prototype description).
- [62] M. H. Böhlen, R. T. Snodgrass, and M. D. Soo. Coalescing in Temporal Databases. In T. M. Vijayaraman, A. Buchmann, C. Mohan, and N. L. Sarda, editors, *Proceedings of the International Conference on Very Large Data Bases*, pages 180–191, Mumbai (Bombay), India, September 1996.
- [63] R. Busatto, M. Böhlen, and C. S. Jensen. Point versus Interval-based Temporal Data Models. In *Proceedings of the Fourteenth IEEE International Conference on Data Engineering*, pages 192–200, Orlando, Florida, February 1998.
- [64] L. Bussard, L. Carver, E. Ernst, M. Jung, M. Robillard, and A. Speck. Safe Aspect Composition. In *Object-Oriented Technology: ECOOP'2000 Workshop Reader*, LNCS, Springer Verlag, 2000, 6 pages, to appear.
- [65] C. E. Dyreson, M. H. Böhlen, and C. S. Jensen. Capturing and Querying Multiple Aspects of Semistructured Data. In *Proceedings of the Twentyfifth International Conference on Very Large Databases*, pages 290–301, Edinburgh, Scotland, UK, September 1999.
- [66] E. Ernst. Separation of Concerns and Then What? In *Proceedings of the Workshop Aspects and Dimensions of Concern at ECOOP'2000*, Cannes, France, June 2000, 8 pages. Electronic form.
- [67] E. Ernst. Syntax Based Modularization: Invasive or Not? In *Proceedings of the Workshop Advanced Separation of Concerns at OOPSLA'2000*, Minneapolis, Minnesota, October 2000, 7 pages.
- [68] H. Gregersen and C. S. Jensen. On the Ontological Expressiveness of Temporal ER Models. In *Advances in Conceptual Modeling*, volume 1727 of LNCS, Springer Verlag, pages 110–121, Paris, France, November 1999.
- [69] H. Gregersen, C. S. Jensen, and L. Mark. Evaluating Temporally Extended ER Models. In *Proceedings of the Second CAiSE/IFIP8.1 International Workshop on Evaluation of Modeling Methods in Systems Analysis and Design*, pages E1–E12, June 1997.
- [70] J. Gu, T. B. Pedersen, and A. Shoshani. OLAP++: Powerful and Easy-to-Use Federations of OLAP and Object Databases. In *Proceedings of the Twenty-Sixth International Conference on Very Large Data Bases*, pages 599–602, Cairo, Egypt, September 2000 (prototype description).

- [71] L. Iversen, Per Madsen, and K. Nørmark. Tools for Presentation and Animation of Dynamic Models in Object-oriented Design. In Mughal and Opdahl, editors, *Proceedings of NWPER'98 - Nordic Workshop on Programming Environment Research*, pages 127–142, June 1998. Also available from <http://www.cs.auc.dk/~normark/dyn-models/tools-2/anitools-njc.pdf>.
- [72] O. G. Jensen and M. H. Böhlen. Evolving Relations. In *Proceedings of the Ninth International Workshop on Foundations of Models and Languages for Data and Objects, Database Evolution and Meta-Modeling*, 17 pages, Schloss Dagstuhl, Germany September 2000.
- [73] B. B. Kristensen. Architectural Abstractions and Language Mechanisms. In *Proceedings of the Asia Pacific Software Engineering Conference '96*, Seoul, Korea, 1996.
- [74] B. B. Kristensen and D. C. M. May. Activities: Abstractions for Collective Behavior. In *Proceedings of the European Conference on Object-Oriented Programming*, volume 1098 of *LNCS*, Springer-Verlag, Linz, Austria, 1996.
- [75] B. B. Kristensen and D. C. M. May. Component Composition and Interaction. In *Proceedings of International Conference on Technology of Object-Oriented Languages and Systems*, Melbourne, Australia, 1996.
- [76] B. B. Kristensen and J. Olsson. Roles & Patterns in Analysis, Design and Implementation. In *Proceedings of the 3rd International Conference on Object-Oriented Information Systems*, London, England, 1996.
- [77] N. Lorentzos, N. Tryfona, and J. R. R. Viquiera. Relational Algebra for Spatial Data Management. In *International Workshop on Integrated Spatial Databases: Digital Images and GIS*, volume 1737 of *LNCS*, Springer Verlag, pages 192–210, 1999.
- [78] N. Lorentzos, J. R. R. Viquiera, and N. Tryfona. Quantum-based Spatial Extension to the Relational Model. In *Proceedings of the 7th Panhellenic Conference on Informatics*, pages 37–44, 1999.
- [79] K. Nørmark. Tools for Exploration of Dynamic Models in Object-oriented Design. In L. Bendix, K. Nørmark, and K. Østerbye, editors, *Proceedings of the Nordic Workshop on Programming Environment Research, NWPER'96*, pages 185–195. Department of Computer Science, Institute for Electronic Systems, Aalborg University, R-96-2019, May 1996. <http://www.cs.auc.dk/~normark/NWPER96/proceedings/proceedings.html>.
- [80] K. Nørmark. Towards an Abstract Language for Dynamic Modelling in Object-oriented Design. In *Proceedings of Tools 23*, pages 120–131. IEEE Computer Society Press, 1997.

- [81] K. Nørmark. An Elucidative Programming Environment for Scheme. In *Proceedings of the 2000 Nordic Workshop on Programming Environment Research*, pages 109–126, May 2000. Also available from <http://www.cs.auc.dk/~normark/elucidative-programming/>.
- [82] K. Nørmark. Requirements for an Elucidative Programming Environment. In *Proceedings of the Eighth International Workshop on Program Comprehension*, pages 119–128. IEEE, June 2000. Also available from <http://www.cs.auc.dk/~normark/elucidative-programming/>.
- [83] K. Nørmark. A Suite of WWW-based Tools for Advanced Course Management. In *Proceedings of the 5th Annual SIGCSE/SIGCU Conference on Innovation and Technology in Computer Science Education*, pages 65–68. ACM Press, July 2000. Also Available from <http://www.cs.auc.dk/~normark/laml/>.
- [84] K. Nørmark, M. R. Andersen, C. N. Christensen, V. Kumar, S. Staun-Pedersen, and K. L. Sørensen. Elucidative Programming in Java. In *The Proceedings on the Eighteenth Annual International Conference on Computer Documentation*, pages 483–495, September 2000.
- [85] T. B. Pedersen and C. S. Jensen. Clinical Data Warehousing—A Survey. In *Proceedings of the VIII Mediterranean Conference on Medical and Biological Engineering and Computing*, Section 20.3, Lemesos, Cyprus, July 1998. CD-rom proceedings.
- [86] T. B. Pedersen and C. S. Jensen. Clinical Data Warehousing—A Survey. In *Conference on Healthcare Computing 1998*, 4 pages, Harrogate, United Kingdom, March 1998.
- [87] T. B. Pedersen and C. S. Jensen. Research Issues in Clinical Data Warehousing. In *Proceedings of the Tenth International Conference on Scientific and Statistical Database Management*, pages 43–52, Capri, Italy, July 1998.
- [88] T. B. Pedersen and C. S. Jensen. Multidimensional Data Modeling for Complex Data. In *Proceedings of the Fifteenth IEEE International Conference on Data Engineering*, pages 336–345, Sydney, Australia, March 1999.
- [89] T. B. Pedersen and C. S. Jensen. Advanced Implementation Techniques for Scientific Data Warehouses. In *Proceedings of the First European Media Lab Workshop on Management and Integration of Bio-Chemical Data*, Heidelberg, Germany, September 2000, 9 pages.
- [90] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Extending Practical Pre-Aggregation in On-Line Analytical Processing. In *Proceedings of the*

Twentyfifth International Conference on Very Large Databases, pages 663–674, Edinburgh, Scotland, UK, September 1999.

- [91] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Supporting Imprecision in Multidimensional Databases Using Granularities. In *Proceedings of the Eleventh International Conference on Scientific and Statistical Database Management*, pages 90–101, Cleveland, Ohio, July 1999.
- [92] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. The TreeScape System: Reuse of Pre-Computed Aggregates over Irregular OLAP Hierarchies. In *Proceedings of the 26th International Conference on Very Large Databases*, pages 595–598, Cairo, Egypt, September 2000 (prototype description).
- [93] T. B. Pedersen, A. Shoshani, J. Gu, and C. S. Jensen. Extending OLAP Querying to Object Databases. In *Ninth International Conference on Information and Knowledge Management*, pages 405–413, Washington, DC, November 2000.
- [94] D. Pfoser and C. S. Jensen. Capturing the Uncertainty of Moving-Object Representations. In *Proceedings of the Sixth International Symposium on Spatial Databases*, pages 111–132, Hong Kong, PRC, July 1999.
- [95] D. Pfoser and C. S. Jensen. Incremental Join of Time-Oriented Data. In *Proceedings of the Eleventh International Conference on Scientific and Statistical Database Management*, pages 232–243, Cleveland, Ohio, July 1999.
- [96] D. Pfoser, C. S. Jensen, and Y. Theodoridis. Novel Approaches in Query Processing for Moving Objects. In *Proceedings of the 26th International Conference on Very Large Databases*, pages 395–406, Cairo, Egypt, September 2000.
- [97] D. Pfoser and Y. Theodoridis. Generating Semantics-Based Trajectories of Moving Objects. In *Proceedings of the International Workshop on Emerging Technologies for Geo-Based Applications*, pages 59–76, 2000.
- [98] D. Pfoser and N. Tryfona. Requirements, Definitions, and Notations for Spatiotemporal Applications. In *Proceedings of the 6th International Symposium on the Advances of Geographic Information Systems*, pages 124–130, 1998.
- [99] R. Price, N. Tryfona, and C. S. Jensen. Modeling Part-Whole Relationships For Spatial Data. In *Eighth International Symposium of ACMGIS*, pages 1–8, Washington DC, November 2000.
- [100] S. Šaltenis, C. S. Jensen, S. Leutenegger, and M. Lopez. Indexing the Positions of Continuously Moving Objects. In *Proceedings of the 2000 ACM*

SIGMOD International Conference on the Management of Data, pages 331–342, Dallas, TX, USA, May 2000.

- [101] J. Skyt and C. S. Jensen. Persistent Views—A Mechanism for Managing Aging Data. In *Proceedings of the Fifth IFCIS International Conference on Cooperative Information Systems*, volume 1901 of *LNCS*, Springer Verlag, pages 132–137, Eilat, Israel, September 2000.
- [102] G. Slivinskas, C. S. Jensen, and R. T. Snodgrass. Query Plans for Conventional and Temporal Queries Involving Duplicates and Ordering. In *Proceedings of the Sixteenth IEEE International Conference on Data Engineering*, pages 547–558, San Diego, CA, USA, February/March 2000.
- [103] R. T. Snodgrass, M. Böhlen, C. S. Jensen, and A. Steiner. Transitioning Temporal Support in TSQL2 to SQL3. In *Proceedings of the Dagstuhl Seminar on Temporal Databases*, pages 150–194, 1997.
- [104] K. Torp, C. S. Jensen, and M. Böhlen. Layered Temporal DBMSs—Concepts and Techniques. In *Proceedings of the Fifth International Conference on Database Systems for Advanced Applications*, pages 371–380, Melbourne, Australia, April 1997.
- [105] K. Torp, C. S. Jensen, and R. T. Snodgrass. Supporting Temporal Data Management Applications via Stratum Approaches. In *Proceedings of the 1998 International Database Engineering and Applications Symposium*, pages 4–13, Cardiff, Wales, U.K., July 1998.
- [106] N. Tryfona. Modeling Phenomena in Spatiotemporal Applications: Desiderata and Solutions. In *9th International Conference on Database and Expert Systems Applications*, volume 1460 of *LNCS*, Springer Verlag, pages 155–165, 1998.
- [107] N. Tryfona. A Framework for Constraint-based Spatial Data Mining. In *International Workshop on Emerging Technologies for Geo-Based Applications*, pages 309–328, Springer Verlag, 2000.
- [108] N. Tryfona, S. Andersen, S. R. Mogensen, and C. S. Jensen. A Methodology and a Tool for Spatiotemporal Database Design. In *Proceedings of the Seventh Hellenic Conference on Informatics*, pages III 53–60, Ioannina, Greece, August 1999.
- [109] N. Tryfona, F. Busborg, and J. B. Christiansen. starER: A Conceptual Model for Data Warehouse Design. In *ACM Second International Workshop on Data Warehousing and OLAP*, pages 3–8. ACM Press, 1999.

- [110] N. Tryfona and T. Hadzilacos. Logical Data Modeling for Spatio-Temporal Applications: Definitions and a Model. In *International Database Engineering and Applications Symposium*, pages 14–23. IEEE Press, 1998.
- [111] N. Tryfona and C. S. Jensen. Using Abstractions for Spatio-Temporal Conceptual Modeling. In *Proceedings of the 2000 ACM Symposium on Applied Computing*, pages 313–322, Villa Olmo, Como, Italy, March 2000.
- [112] N. Tryfona and M. Kavouras. A Structured Methodology to Support Interoperable Geographic Applications: The case of the Greek Cadastral System. In *Association of Geographic Information Laboratories in Europe*, 1998. CD-rom proceedings.
- [113] Uffe K. Wiil and John J. Leggett. The HyperDisco Approach to Open Hypermedia Systems. In *Proceedings of the Seventh ACM Conference on Hypertext*, pages 140–148, Washington DC, March 1996.
- [114] Uffe K. Wiil and John J. Leggett. Workspaces: The HyperDisco Approach to Internet Distribution. In *Proceedings of the Eighth ACM Conference on Hypertext*, pages 13–23, Southampton, England, April 1997.
- [115] U. K. Wiil and K. Østerbye. Using the Flag Taxonomy to Study Hypermedia System Interoperability. In *Proceedings of the Ninth ACM Conference on Hypertext*, pages 188–197, 1998.
- [116] K. Østerbye. Minimalist Documentation of Frameworks. In *The 3rd ECOOP'99 Workshop on Experiences in Object-Oriented Reengineering*, 1999.
- [117] K. Østerbye and U. K. Wiil. The Flag Taxonomy of Open Hypermedia Systems. In *Proceedings of the Seventh ACM Conference on Hypertext*, pages 129–139, 1996.

In submission:

- [118] M. Akinde, T. Johnson, L. Lakshmanan, and D. Srivastava. Efficient OLAP Query Processing in Distributed Data Warehouses. Submitted to the *ACM SIGMOD International Conference on the Management of Data*, 25 pages, Santa Barbara, CA, May 2001.
- [119] E. Ernst. Family Polymorphism. Submitted to *ECOOP'2001*.
- [120] G. Slivinskas, C. S. Jensen, and R. T. Snodgrass. Adaptive Query Optimization and Evaluation in Temporal Middleware. Submitted to the *ACM SIGMOD International Conference on the Management of Data*, 25 pages, Santa Barbara, CA, May, 2001.

Book Chapters

- [121] M. Breuning, T. Can, Böhlen, M., S. Dieker, R. H. Güting, C. S. Jensen, L. Relly, P. Rigaux, H.-J. Schek, and M. Scholl. Architecture and Implementation of Spatio-Temporal DBMS. In T. Sellis et al., editors, *CHOROCHRONOS: Research Issues in Spatiotemporal Applications*, Springer Verlag, 59 pages, to appear.
- [122] A. Di Pasquale, L. Forlizzi, C. S. Jensen, Y. Manolopoulos, E. Nardelli, D. Pfoser, G. Proietti, S. Šaltenis, Y. Theodoridis, T. Tzouramanis, and M. Vassilakopoulos. Access Methods and Query Processing. In T. Sellis et al., editors, *CHOROCHRONOS: Research Issues in Spatiotemporal Applications*, Springer Verlag, 61 pages, to appear.
- [123] R. H. Güting, M. Böhlen, M. Erwig, L. Forlizzi, C. S. Jensen, N. Lorentzos, E. Nardelli, M. Schneider, and M. Vazirgiannis. Spatiotemporal Models and Languages: An Approach Based on Data Types. In T. Sellis et al., editors, *CHOROCHRONOS: Research Issues in Spatiotemporal Applications*, Springer Verlag, 52 pages, to appear.
- [124] C. S. Jensen. Temporal Database Concepts. In A. Kent and J. Williams, editors, *Encyclopedia of Computer Science and Technology*, Marcel Dekker, 27 pages, to appear (same as [125]).
- [125] C. S. Jensen. Temporal Database Concepts. In A. Kent and J. Williams, editors, *Encyclopedia of Microcomputers*, Marcel Dekker, 27 pages, to appear (same as [124]).
- [126] C. S. Jensen and C. E. Dyreson, editors (with multiple other contributors). The Consensus Glossary of Temporal Database Concepts—February 1998 Version. In *Temporal Databases: Research and Practice*, volume 1399 of *LNCS*, Springer Verlag, pages 367–405, 1998.
- [127] C. S. Jensen. Temporally Enhanced Database Design. Chapter 7, pp. 163–193, in M. P. Papazoglou, S. Spaccapietra, and Z. Tari, editors, *Advances in Object-Oriented Data Modeling*, MIT Press, 2000.
- [128] N. Lorentzos, J. R. R. Viqueira, and N. Tryfona. Quantum-based Spatial Extensions to the Relational Model. In, D. I. Fotiadis and, S. D. Nikolopoulos, editors, *Advances in Informatics*, World Scientific, to appear.
- [129] N. Tryfona and B. Theodoulidis. Spatiotemporal Conceptual Modeling. In T. Sellis et al., editors, *CHOROCHRONOS: Research Issues in Spatiotemporal Applications*, Springer Verlag, 2000, to appear.

Lightly Refereed or Unrefereed Journal Articles

- [130] U. Asklund, L. Bendix, and Persson J. Summary of the Subworkshop on Change Management for Open Source Software. *Nordic Journal of Computing*, 7(2):143–147, 2000.
- [131] L. Bendix and U. Asklund. Summary of the NWPER'98 - Subworkshop on Change Management. *Nordic Journal of Computing*, 6(1):129–133, 1999.
- [132] A. Frank, R. H. Güting, C. S. Jensen, M. Koubarakis, N. Lorentzos, Y. Manolopoulos, E. Nardelli, B. Pernici, H.-J. Schek, M. Scholl, T. Sellis, B. Theodoulidis, and P. Widmayer. Chorochronos—A Research Network for Spatiotemporal Database Systems. *ACM SIGMOD Record*, 28(3):12–21, September 1999.
- [133] B. B. Kristensen and K. Østerbye. A Conceptual Perspective on the Comparison of Object-Oriented Programming Languages. *SIGPLAN Notices*, 31(2):42–54, February 1996.
- [134] K. Nørmark. Programming World Wide Web Pages in Scheme. *SIGPLAN Notices*, 34(12):37–46, December 1999. Also available from <http://www.cs.auc.dk/~normark/laml/>.
- [135] V. Tsotras, C. S. Jensen, and R. T. Snodgrass. An Extensible Notation for Spatiotemporal Index Queries. *ACM SIGMOD Record*, 27(1):47–53, March 1998.

Miscellaneous

- [136] M. H. Böhlen. Computing Review for Handling Infinite Temporal Data by F. Kabanza, J.-M. Stèvenne, and P. Wolper in Journal of Computer and System Sciences, 51, 1995. *ACM Computing Review*, October 1996.
- [137] M. H. Böhlen. Computing Review for Tail Recursion Elimination in Deductive Databases by K. A. Ross in ACM Transactions on Database Systems, Vol. 21, Nr. 2, 1996. *ACM Computing Review*, August 1997.
- [138] M. H. Böhlen. Computing Review for A First Course in Database Systems by J. D. Ullman and J. Widom, Prentice Hall, 1997. *ACM Computing Review*, April 1998.
- [139] M. H. Böhlen, W. Jonker, and M. Revett. Data Management Technologies in Telcos. EURESCOM P817 Booklet, May 2000.
- [140] C. S. Jensen. Databasen er strategisk. *Computerworld*, page 27, November 1999. in Danish.

- [141] C. S. Jensen. Review - Equivalence of Relational Algebra and Relational Calculus Query Languages Having Aggregate Functions. *ACM SIGMOD Digital Review*, 1, 1999.
- [142] C. S. Jensen. Review - Multi-Step Processing of Spatial Joins. *ACM SIGMOD Digital Review*, 1, 1999.
- [143] C. S. Jensen. Review - R-Trees: A Dynamic Index Structure for Spatial Searching. *ACM SIGMOD Digital Review*, 1, 1999.
- [144] C. S. Jensen. Review - The Logical Access Path Schema of a Database. *ACM SIGMOD Digital Review*, 2, 2000.
- [145] C. S. Jensen. Themes and Challenges in Temporal Databases. In H. Schmidt, editor, *Modellierung Betrieblicher Informationssysteme. (Proceedings der MOBIS-Fachtagung 2000)*, pages 175–176, Universität Siegen, Germany, October 2000.
- [146] C. S. Jensen and R. T. Snodgrass. Tutorial on Temporal Databases. In *Twelfth International Conference on Data Engineering—Tutorial Notes*, New Orleans, Louisiana, February/March 1996.
- [147] K. Nørmark. Dynamic Modelling With Dynamo. Poster Synopsis from OOPSLA 1996, 1996. Available from <http://www.cs.auc.dk/~normark/dyn-models/oopsla-96-poster/synopsis.pdf>.
- [148] K. Nørmark. Deriving Classes from Scenarios in Object-oriented Design, 1997. Preliminary version available on <http://www.cs.auc.dk/~normark/dynamo.html>.
- [149] K. Nørmark. DYNAMO: A set of Tools for Dynamic Modelling, 1997. Available from <http://www.cs.auc.dk/~normark/dyn-models/tools-2/tools2.pdf>.
- [150] K. Nørmark. Synthesis of Program Outlines from Dynamic Models in DYNAMO, 1997. Preliminary version available on <http://www.cs.auc.dk/~normark/dynamo.html>. Abbreviated version of [148].
- [151] K. Nørmark. Synthesis of Program Outlines from Scenarios in DYNAMO. Available from <http://www.cs.auc.dk/~normark/dyn-models/static-models/abstract.html>, 1998.
- [152] K. Nørmark. Using Lisp as a markup language—The LAML approach. In *European Lisp User Group Meeting*. Franz Inc., 1999. Also available from <http://www.cs.auc.dk/~normark/laml/>.

- [153] R. T. Snodgrass, M. H. Böhlen, C. S. Jensen, and A. Steiner. Adding Transaction Time to SQL/Temporal. Technical Report ANSI X3H2-96-501r1, International Organization for Standardization, 44 pages, November 1996.
- [154] R. T. Snodgrass, M. H. Böhlen, C. S. Jensen, and A. Steiner. Adding Valid Time to SQL/Temporal. Technical Report ANSI X3H2-96-502r2, International Organization for Standardization, 77 pages, October 1996.

Technical Reports

- [155] L. Bendix. How to Introduce Maturity in Software Change Management. Technical Report R-98-5012, Department of Computer Science, Aalborg University, July 1998.
- [156] L. Bendix. A Disciplined Approach to Change Management. Technical Report R-99-5001, Department of Computer Science, Aalborg University, March 1999.
- [157] L. Bendix and Vitali F. CoEd meets VTML. Technical Report R-98-5007, Department of Computer Science, Aalborg University, June 1998.
- [158] L. Bendix, P. N. Larsen, A. I. Nielsen, and J. L. S. Petersen. CoEd - A Tool for Cooperative Development of Hierarchical Documents. Technical Report R-97-5012, Department of Computer Science, Aalborg University, September 1997.
- [159] M. Böhlen and C. S. Jensen. Seamless Integration of Time into SQL. Technical Report R-96-2049, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, December 1996.
- [160] H. Gregersen and C. S. Jensen. Temporal Entity-Relationship Models—a Survey. Technical Report R-96-2039, TimeCenter, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, September 1996.
- [161] C. S. Jensen and R. T. Snodgrass. Semantics of Time-Varying Information. Technical Report R-96-2008, Department of Mathematics and Computer Science, Aalborg University, Department of Mathematics and Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, March 1996.
- [162] B. B. Kristensen. Component Composition and Interaction in BETA: The Customer Servicing System Example. Technical Report R 97-5001, Department of Computer Science, Aalborg University, 1997.

- [163] B. B. Kristensen and D. C. M. May. Modeling with Activities: Abstractions for Collective Behavior. Technical Report R 96-2001, Department of Computer Science, Aalborg University, 1996.
- [164] N. C. Larsen, M. K. Molz, and K. Nørmark. Graph Abstractions as the basis of an Extensible Graph Editing Tool. Available from <http://www.cs.auc.dk/~normark/ginger.html>, 1998.
- [165] K. Nørmark. Dynamic Models in Object-Oriented Design. Technical Report R-96-2005, Department of Mathematics and Computer Science, Institute of Electronic Systems, Aalborg University, February 1996.
- [166] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Extending Practical Pre-Aggregation in On-Line Analytical Processing. Technical Report R-99-5004, Aalborg University, Department of Computer Science, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, September 1999.
- [167] T. B. Pedersen, C. S. Jensen, and C. E. Dyreson. Supporting Imprecision in Multidimensional Databases Using Granularities. Technical Report R-99-5003, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, May 1999.

Technical Reports, Chorochronos

- [168] R. H. Güting, M. Böhlen, M. Erwig, C. S. Jensen, N. Lorentzos, M. Schneider, and M. Vazirgiannis. A Foundation for Representing and Querying Moving Objects. Technical Report CH-98-15, Chorochronos Technical Report, September 1998.
- [169] D. Pfoser, C. S. Jensen, and Y. Theodoridis. Novel Approaches in Query Processing for Moving Objects. Technical Report CH-00-03, Chorochronos Technical Report, February 2000.
- [170] D. Pfoser, Y. Theodoridis, and C. S. Jensen. Indexing Trajectories of Moving Point Objects. Technical Report CH-99-3, Chorochronos Technical Report, June 1999.
- [171] R. Price, N. Tryfona, and C. S. Jensen. A Conceptual Modeling Language for Spatiotemporal Applications. Technical Report CH-99-20, Chorochronos Technical Report, December 1999.
- [172] S. Šaltenis, C. S. Jensen, S. Leutenegger, and M. Lopez. Indexing the Positions of Continuously Moving Objects. Technical Report TR-44, Time-Center, November 1999. Chorochronos Technical Report CH-99-19.

- [173] N. Tryfona and C. S. Jensen. A Component-Based Conceptual Model for Spatiotemporal Application Design. Technical Report CH-98-10, Chorochronos, November 1998.
- [174] N. Tryfona and C. S. Jensen. Conceptual Modeling for Spatiotemporal Applications. Technical Report CH-98-8, Chorochronos, November 1998.

Technical Reports, TimeCenter

- [175] J. Bair, M. Böhlen, C. S. Jensen, and R. T. Snodgrass. Notions of Upward Compatibility of Temporal Query Languages. Technical Report R-96-2038, TimeCenter, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, September 1996.
- [176] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. Light-Weight Indexing of General Bitemporal Data. Technical Report TR-30, TimeCenter, September 1998.
- [177] R. Bliujūtė, C. S. Jensen, S. Šaltenis, and G. Slivinskas. R-tree Based Indexing of Now-Relative Bitemporal Data. Technical Report TR-25, TimeCenter, March 1998.
- [178] R. Bliujūtė, S. Šaltenis, G. Slivinskas, and C. S. Jensen. Developing a DataBlade for a New index. Technical Report TR-29, TimeCenter, September 1998.
- [179] R. Bliujūtė, S. Šaltenis, G. Slivinskas, and C. S. Jensen. Systematic Change Management in Dimensional Data Warehousing. Technical Report TR-23, TimeCenter, January 1998.
- [180] M. Böhlen, R. Busatto, and C. S. Jensen. Point versus Interval-based Temporal Data Models. Technical Report TR-21, TimeCenter, January 1998.
- [181] M. Böhlen, C. S. Jensen, and B. Skjellaug. Spatio-Temporal Database Support for Legacy Applications. Technical Report TR-20, TimeCenter, July 1997.
- [182] C. E. Dyreson, M. H. Böhlen, and C. S. Jensen. Capturing and Querying Multiple Aspects of Semistructured Data. Technical Report TR-36, TimeCenter, November 1998.
- [183] H. Gregersen and C. S. Jensen. Conceptual Modeling of Time-Varying Information. Technical Report TR-35, TimeCenter, September 1998.

- [184] H. Gregersen, L. Mark, and C. S. Jensen. Mapping Temporal ER Diagrams to Relational Schemas. Technical Report TR-39, TimeCenter, December 1998.
- [185] C. S. Jensen and R. T. Snodgrass. Temporal Data Management. Technical Report TR-17, TimeCenter, June 1997.
- [186] T. B. Pedersen and C. S. Jensen. Multidimensional Data Modeling for Complex Data. Technical Report TR-37, TimeCenter, November 1998.
- [187] D. Pfoser and C. S. Jensen. Incremental Join of Time-Oriented Data. Technical Report TR-34, TimeCenter, September 1998.
- [188] S. Šaltenis and C. S. Jensen. R-Tree Based Indexing of General Spatio-Temporal Data. Technical Report TR-45, TimeCenter, December 1999.
- [189] J. Skyt and C. S. Jensen. Vacuuming Temporal Databases. Technical Report TR-32, TimeCenter, September 1998.
- [190] G. Slivinskas, C. S. Jensen, and R. T. Snodgrass. A Foundation for Conventional and Temporal Query Optimization Addressing Duplicates and Ordering. Technical Report TR-49, TimeCenter, February 2000.
- [191] R. T. Snodgrass, M. Böhlen, C. S. Jensen, and A. Steiner. Transitioning Temporal Support in TSQL2 to SQL3. Technical Report TR-8, TimeCenter, April 1997.
- [192] K. Torp, C. S. Jensen, and M. Böhlen. Layered Temporal DBMSs—Concepts and Techniques. Technical Report R-96-2037, TimeCenter, Aalborg University, Department of Computer Science, Fredrik Bajers Vej 7E, DK-9220 Aalborg, September 1996.
- [193] K. Torp, C. S. Jensen, and R. T. Snodgrass. Stratum Approaches to Temporal DBMS Implementation. Technical Report TR-5, TimeCenter, March 1997.
- [194] K. Torp, C. S. Jensen, and R. T. Snodgrass. Modification of Now-Relative Databases. Technical Report TR-43, TimeCenter, September 1999.
- [195] K. Torp, R. T. Snodgrass, and C. S. Jensen. Correct and Efficient Timestamping of Temporal Data. Technical Report TR-4, TimeCenter, March 1997.
- [196] V. Tsotras, C. S. Jensen, and R. T. Snodgrass. A Notation for Spatiotemporal Queries. Technical Report TR-10, TimeCenter, April 1997.

Chapter 6

Decision Support Systems

6.1 Profile of the Unit

The research of the unit concerns development of theories and methods for computer-based support of decision making under uncertainty. The characteristics of this kind of decision making are that decisions have to be made on the basis of insufficient or uncertain information, and that usually the consequences of the decisions are subject to uncertainty.

This kind of research is part of *Artificial Intelligence* (AI). However, rather than aiming at constructing systems that “think” as humans, the research of the unit aims at constructing systems that can advise humans to act rationally. By *acting rationally* we mean making decisions which maximize the expected utility (the so-called *normative* approach).

Normative decision support systems are based on models of the problem domain rather than on models of the decision maker’s line of reasoning. A normative decision support system is characterized by

- (graphical) representation of causal relations between quantities
- probabilistic representation of the strengths of causal relations
- representation of needs and preferences through utility measures
- application of the principle of maximum utility for decision making

The most widely used types of normative systems are known as *Bayesian networks* and *influence diagrams*.

The research unit is engaged in the development of theories and methods for the modelling and inference in normative systems along with the development of theories and methods for systems that adapt to specific users and contexts. Since the mid-80s, when a pronounced renewed interest in normative systems arose, the unit has had a central position in the world-wide development of this field.

The results have been tested through the development of systems for e.g. medical diagnosis, troubleshooting in computer equipment, and agricultural planning.

6.2 Activities and Results

A research and development project between the research unit and Hewlett-Packard Company was initiated in November 1997 and has played a dominant role for the research activities of the unit, and has enabled a significant increase in staffing of the unit (see Section 6.3).

6.2.1 Research Results

As a background for the specification of research results listed below we should mention that our research in normative systems has been focusing on graphical models. A graphical model consists of a graph with various types of nodes and with directed and/or undirected links between them. The most prominent examples are Bayesian networks (directed acyclic graphs with only chance nodes) and influence diagrams (Bayesian networks extended with decision nodes and utility nodes — meeting certain structural constraints). Quantitative specifications consisting of (conditional) probability potentials and utility functions are attached to the graph, and the structure of the graph determines uniquely the types and dimensions of the potentials and functions attached. In this respect, “graphical models” is a formal language for specification and communication of models — interhuman as well human-computer. When a model specification has been communicated to a computer, the model is compiled into a structure fit for computing answers to user requests like optimal strategies, posterior beliefs, expected benefit, optimal next action, analysis of the evidence for internal conflicts, etc. Our favourite computational structure is the so-called junction tree and variants thereof. The method used in the algorithms is message passing between nodes in a junction tree.

The following summarizes the research conducted within various categories. Please note that some of the references appear in several sections, as the work reported in these references fall under several categories.

Inference

The basic belief updating task for Bayesian networks is NP-hard, and therefore there is a constant need for improving the updating algorithms (exact and approximate) in order to expand the scope of tractable models. Contributions from the unit has been

- lazy propagation [26, 27, 4] which is a lazy evaluation method for Bayesian networks and influence diagrams,

- use of binary decision diagrams for Boolean parts of a model [30],
- exploitation of the junction tree architecture to sample large numbers of variables simultaneously [3, 17],
- exploitation of deterministic relations in influence diagrams [53],
- exploitation of scenario specific independences through nested junction trees [24, 10] or lazy propagation [18, 4],
- an analysis of various approaches to distributed computation of belief updating in junction trees [52],
- a method for inference in mixed inheritance models [17],
- a method for maximal subgraph decomposition of Bayesian networks enabling divide and conquer triangulation and hybrid propagation [55], and
- a method for treating likelihood evidence as hard evidence [57].

Analysis tools

Apart from basic belief updating and calculation of optimal strategies, the models are used to give answers to many other types of queries. In this period, the unit has contributed with methods for

- computing the m most probable configurations of unobserved variables [5, 19],
- computing the probabilities of future decisions [11],
- performing linkage analysis in large and complex pedigrees [12] based on a sampling-based inference algorithm [3],
- computing the impact of knowing the state of a variable (value of information) for influence diagrams [22],
- performing sensitivity analysis with respect to the parameters in a Bayesian network model: how sensitive is the conclusion to the specific values of the parameters [25, 47], and more specifically, in Bayesian troubleshooting models [49],
- multivariate updating of genotypes in a mixed inheritance model [38],
- exploiting the typical structure of medical systems [56],
- finding a configuration of positive probability given evidence in complex Bayesian networks [13],

- performing gradient descent training of Bayesian networks [36], and
- parameter estimation for conditional Gaussian networks with missing data [44].

Extending the scope of graphical models

Traditionally, influence diagrams have rather severe structural constraints. First of all, a linear temporal ordering of the decisions is required. In [28] it is shown that this constraint can be relaxed to a partial order, and necessary and sufficient structural conditions are given to ensure the specification in the model to be well-defined. Also, as opposed to decision trees, influence diagrams require the decision scenario to be symmetric: future decision options may not depend on past decisions and observations. In [54, 29] a graphical specification language is given which has almost the same expressive power as decision trees. The specifications are much more compact than decision trees, and the solutions algorithm is substantially faster than the algorithm for decision trees. Also, [53] presents an adjusted representation of influence diagrams to exploit deterministic relations.

Knowledge representation

Although graphical models is a very handy and compact way of specifying structure, it has shortcomings with respect to repetitive structures and models with many (almost) identical submodels. The unit has extended Y. Xiang's work on multiply-sectioned Bayesian networks [32, 21], and [33, 46] present a graphical language for object-oriented specification of Bayesian networks and for specification of networks with repetitive (sub)structures.

Knowledge acquisition

Efficient methods for acquiring domain knowledge in building normative systems have also been in the focus of the unit's research.

A method has been developed for easing the specification of probabilities for Bayesian-network troubleshooters through acquisition of likelihoods [41]. Later, some of these ideas were incorporated into a knowledge acquisition tool [31].

Applications

The joint research and development project with Hewlett-Packard has played a major role in the application-oriented research of the unit, but also applications within agriculture, medicine, and games have been developed.

The list of application-oriented research includes

- troubleshooting systems for networked HP-printer systems [42, 41, 48, 40],

- a framework for generic object recognition with Bayesian networks [37],
- a system for management of mildew in winter wheat [23],
- methods for improving the performance of the MUNIN system (diagnosis of muscle and nerve diseases) [20, 56, 45],
- a model for the problem of laboratory prevalences versus population prevalences [20],
- a minor extension in a system for treatment of bacterial infections [6], and
- a system for calculation of probabilities in the card game Bridge [39].

Foundations

Some of the research conducted has dealt with more fundamental problems in normative systems. This work includes

- problems with the determination of the noncommunicating classes for Markov chain Monte-Carlo applications in pedigree analysis [14],
- exploration of the classes of problems solvable by local computations in a junction tree [16],
- proof that troubleshooting is NP-hard even under a number of different sets of constraining assumptions [43], and
- exploration of the conditions for troubleshooting data to be missing at random [51] (required by algorithms for learning/updating parameters in a Bayesian-network troubleshooter).

Tutorials and surveys

A member of the unit has written and co-authored a number of tutorials and surveys on Bayesian networks and influence diagrams, including

- two textbooks introducing the theory and algorithms for building normative systems and for making inference in them [1, 2] and
- tutorials and/or surveys on Bayesian networks and influence diagrams [15, 35, 7, 8, 9].

6.2.2 Projects

DINA

The unit has been associated with the Danish Informatics Network in the Agricultural Sciences (DINA) from 1995 onwards. In particular, it has been involved on contracting basis in one of the projects running within DINA.

Members of the unit have participated in several research workshops and seminars in DINA, and two of the members of the unit have played active roles as members of the Network Committee of DINA, and one as a member of the board of DINA.

Furthermore, as a result of the association with DINA, one the members of the unit currently acts as a part-time supervisor for a Ph.D. student affiliated with a DINA research unit at the Royal Veterinary and Agricultural University of Denmark.

SACSO

The SACSO project is a collaboration between the research unit and Hewlett-Packard. The project was initiated in November 1997 and is still running. The project is funded jointly by Hewlett-Packard (HP) and the Danish National Centre for IT Research (CIT).

The aim of the project is development of systems for automated customer support operations (SACSO) through the use of decision theoretic methodologies, notably Bayesian networks and influence diagrams.

The project, called SACSO, has been running in three phases. In Phase-I, lasting one year, the aim was to establish a proof of concept for the application of Bayesian network technology in the area of troubleshooting electro-mechanical systems, specifically networked printer systems. By the end of Phase-I a complete prototype troubleshooter successfully demonstrated the applicability of the Bayesian network technology in this area.

Phase-II, which ended in the spring of 2000, had as its two main objectives to

- continue the Phase-I activity, including completing, maturing, and validating the Phase-I prototype system, and refining and extending the knowledge acquisition process (the HP-side), and to
- carry out supporting scientific work, involving development of methodologies for validation and maintenance of troubleshooters, including methods for manual as well as automatic adaptation of model structure, probabilities and costs.

Concretely, the first two phases of the project have focused on development of methods for interactive troubleshooting of complex electro-mechanical systems,

construction of a generic troubleshooting environment, construction of causal models of concrete electro-mechanical systems, and construction of a tool supporting the construction of these models.

The causal models (Bayesian networks) represent causal relationships among system components, system configuration, possible causes of system faults, and possible observations of system behaviour. Given the models and observations of system behaviour, the developed methods provide means for computing cost-effective sequences of troubleshooting actions, eventually solving the problems.

Complete troubleshooters for four HP printers have been developed, and, as a spin-off, HP expects troubleshooters for a broad variety of products and services to be developed and offered to end-users as web-based e-service applications. The introduction of such e-services is expected to provide quicker, higher quality service to customers, and thereby significantly reduce the need for customer support operations provided by human call agents.

The research and development is being carried out in a collaboration involving a combination of 3 full-time and 2-3 part-time researchers from the unit and 1-2 full-time and 3-5 part-time research and development engineers of Hewlett-Packard.

The research activities in the SACSO project have so far resulted in 3 patent applications, 15 reviewed scientific papers, 11 technical reports, 10 working papers, and 13 presentations.

Further information about the SACSO project can be found at the web pages of the project (www.cs.auc.dk/research/DSS/SACSO).

6.2.3 Patents and Products

As a result of the research and development activities in the SACSO project Hewlett-Packard and researchers from the unit have jointly filed three patents:

- A method for knowledge acquisition for diagnostic Bayesian networks
- Automated diagnosis of printer systems using Bayesian networks
- An authoring tool for Bayesian-network troubleshooters

The research conducted in the SACSO project has enabled Hewlett-Packard to develop a number of products, including

- complete troubleshooting systems for four HP laser printers (HP are currently preparing the introduction of these as e-service applications to be offered to end-users)
- a knowledge acquisition tool, called BATS Author, making it possible for domain experts to build troubleshooting models without prior exposure to Bayesian-network technology

- a run-time troubleshooting environment, called BATS Troubleshooter

6.2.4 External Funding

In the period from 1996–2000, the unit has received the following external funding.

- The Danish Natural Science Research Council has funded a Ph.D. programme with 1M DKK.
- The Danish National Centre for IT Research (CIT) has sponsored Phase-I of the SACSO project with 1.2M DKK.
- The Danish National Centre for IT Research (CIT) has sponsored Phase-II of the SACSO project with 3.5M DKK (including funding of a Ph.D. programme).
- Over the 5-year period 1998–2003, Hewlett-Packard sponsors the unit with 1M DKK per year.
- In response to the donation from Hewlett-Packard, the Ministry of Education has supported the unit with 2.4M DKK (1.6M DKK in the 5-year period 1998–2003).

6.3 Organization and Staff

6.3.1 Staff

The tables below give the numerical information on the staffing (both on teaching and pure research) in the period of this evaluation. Numerical information is given for each year with more detailed half-year information indicated by left- or right-shifted numbers.

Teaching Positions					
	1996	1997	1998	1999	2000
Full Professors			1	1	1
Associate Professors	1	1 2	2 1	2	2
Assistant Professors	1	1			

Research Positions					
	1996	1997	1998	1999	2000
Visiting Professors					
Visiting Researchers			1	1	1
Research Assistant			3	3	3
Ph.D. students	3 2	3 2	2	2	2

The following people have contributed to the research results of the unit during the period. Only Finn V. Jensen and Uffe Kjærulff have been employed throughout the period.

Finn V. Jensen

Academic degrees

- 1977 Pædagogikum (Secondary School Teacher's Certificate)
- 1974 Dr. math. in Mathematical Logic, Warsaw University
- 1970 Cand. scient. (M.Sc.) in Mathematics, Århus University

Positions

- 1998– Professor, Aalborg University
- 1989–98 Reader (docent), Aalborg University
- 1986–89 Project Manager, JUDEX A/S, Aalborg
- 1974–86 Associate Professor, Aalborg University
- 1972–74 Research Fellow (Kandidatstipendiat), Århus University
(studies in Warsaw)
- 1970–72 Scientific Assistant (Kandidatinstruktør), Århus University
- 1969–70 Scientific Assistant (Videnskabelig Assistent), Oslo University

Key Publications: [1, 16, 18, 29]

Uffe Kjærulff

Academic degrees

- 1993 Ph.D. in Computer Science, Aalborg University
- 1985 Cand. polyt. (M.E.) in Computer Engineering, Aalborg University

Positions

- 1997– Associate Professor, Aalborg University
- 1994–97 Assistant Professor, Aalborg University
- 1993–94 Research Assistant, Aalborg University
- 1990–93 Ph.D. Student, Aalborg University
- 1989–90 Research Fellow (erhvervsforsker), JUDEX A/S, Aalborg
- 1985–89 Research Assistant, JUDEX A/S, Aalborg

Key Publications: [10, 25, 40]

Kristian G. Olesen

Academic degrees

- 1992 Ph.D. in Computer Science, Aalborg University
- 1985 Cand. scient. (M.Sc.) in Computer Science and Mathematics,
Aalborg University

Positions

- 1999– Associate Professor, Aalborg University
- 1996–98 Associate Research Professor, Aalborg University
- 1991–95 Assistant Professor, Aalborg University
- 1990–91 Research Assistant, Aalborg University
- 1988–89 Project Manager, Nordjysk Udviklingscenter, Aalborg
- 1986–88 Ph.D. Student, Aalborg University
- 1985–86 Teaching Assistant, Aalborg University

Key Publications: [20, 53, 55, 56]

Other Staff

- Claus Skaanning (1996–1997 as Ph.D. student, 1997–1998 as Assistant Prof.)
- Dennis Nilsson (1996–1998 as Ph.D. student)
- Anders L. Madsen (1997–1999 as Ph.D. student)
- Thomas D. Nielsen (1998– as Ph.D. student)
- Olav Bangsø (1998–1999 as Research Assistant, 1999– as Ph.D. student)
- Pierre-Henri Wuillemin (1998–1999 as Research Assistant)
- Johan Myhre Andersen (1998–1999 as Research Assistant)
- Helge Langseth (1999–2000 as Research Assistant)
- Jiří Vomlel (1999–2000 as Research Assistant)
- Marta Sochorová (1999–2000 as Research Assistant)

6.3.2 Research organization

In the entire period, the unit has organised series of research seminars taking place every week in the months February to June and September to January. Since the autumn term of 1998, lists of seminars (including title, abstract, etc.) have been maintained at the web-pages of the unit (see www.cs.auc.dk/research/DSS/). At the seminars, presentations are given by internal as well as external researchers.

Starting in the spring of 1998, the research of the unit has been supported by workshops held every week, lasting approximately 1.5 hours. These workshops form the forum in which ideas and results are exchanged and discussed, and have proven to be very valuable in coordinating and initiating research activities of subgroups.

6.4 Collaborations

Within the Unit

As mentioned above, the research collaboration within the unit has been organised through weekly workshops, which have greatly contributed to fruitful discussions and exchange of ideas. In most cases, the research has been carried out in smaller teams of 2-4 people, involving researchers of the unit as well as research engineers of Hewlett-Packard.

Within the Department

The collaboration with other research units of the department has been limited to a couple of presentations, one given to members of the Research Unit of Distributed Systems and Semantics and one given by the same unit. There seems to be basis for further cross fertilisations, as the techniques and methods under study in the two units seem to have some common traits.

With Other Danish Institutions

As described above, the unit is associated with the Danish Informatics Network in the Agricultural Sciences (DINA), and has been involved in a DINA project. Furthermore, members of the unit have participated in several research workshops and seminars in DINA.

As a result of the association with DINA, Finn V. Jensen co-supervises a Ph.D. student of the Royal Veterinary and Agricultural University of Denmark.

At a couple of occasions, two members of a research group at the Technical University of Denmark have visited the research unit, and presented work on normative approaches to diagnosis and decision making.

Also, several members of the Danish Institute of Agricultural Sciences have visited the unit, and presented work on statistical and decision theoretic methods applied in the agricultural sciences. In particular, Claus Skaanning and Ph.D. student Mogens Lund from the institute have collaborated on developing methods for applying Gibbs sampling techniques for inference in mixed inheritance models.

Internationally

Below, we list, in chronological order, the international research collaborations that members of the unit have been involved in during the period of the evaluation.

At University of Chicago, Claus Skaanning worked six months in 1995–96 with Professor Augustine Kong on applying the method of blocking-Gibbs sampling for performing linkage analysis in large pedigrees, used for determining the loci

of specific genes. Until the invention of the blocking-Gibbs algorithm, linkage analysis was considered a computationally extremely hard problem.

Ph.D. student Antonio Salmeron from University of Granada visited the unit for two months in 1996.

Nuala Sheehan from the Department of Mathematical Sciences at Loughborough University, UK, paid the research unit a two-week visit, in which she mainly collaborated with Claus Skaanning on developing methods based on Markov-Chain Monte-Carlo techniques for pedigree analysis (inference in complex normative systems). Later, Claus Skaanning visited Loughborough University for a week.

At University of Kansas, Dennis Nilsson worked six months in 1997 with Professor Prakash Shenoy on methods for finding optimal strategies in Influence Diagrams. An efficient algorithm for finding the m best strategies was developed. The method is a divide-and-conquer algorithm inspired by previous work for finding the m most likely configurations in Bayesian networks.

At Oregon State University, Anders L. Madsen worked for six months in 1998 closely together with Professor Bruce D'Ambrosio on developing methods for improving the representational and computational efficiency of Bayesian networks and influence diagrams. In particular, efficient methods for representation of and computation with independence of causal influence models were developed.

Finn V. Jensen was a Visiting Professor for a month in 1998 at The University of Pierre and Marie Curie, Paris.

Kristian G. Olesen has participated in the MUNIN research group, a group of researchers at the Department of Medical Informatics and Image Analysis at Aalborg University and The University Hospital of Turku, Finland.

In the spring of 1998, Associate Professor Yang Xiang from the Department of Computer Science at the University of Saskatchewan, Canada, visited the unit for a period of two months. During his visit, the work on developing an object-oriented methodology for the construction of Bayesian-network models was initiated. He also presented work on multiply-sectioned Bayesian networks, which later resulted in joint publications with members of the research unit.

Uffe Kjærulff has collaborated with Associate Professor Linda C. van der Gaag from Utrecht University, The Netherlands, on developing efficient methods for making sensitivity analyses in Bayesian networks. The collaboration was facilitated through a one-week visit in January 1999 at the research unit by Linda C. van der Gaag.

Senior Researcher Jaap Suermondt of HP Labs, Palo Alto, paid a one-week visit to the research unit in February 1999.

In the spring of 1999, Ph.D. student Veerle Coupé from the Erasmus University in Rotterdam, The Netherlands, visited the unit for a period of two months. During her stay, she worked with members of the unit on developing a general method for performing sensitivity analysis in Bayesian networks.

In the autumn of 1999 and the spring of 2000, Associate Professor Marco

Valtorta from the Department of Computer Science at the University of South Carolina, USA, visited the unit for a total period of 2.5 months. During his stay, he participated in the SACSO project as well as presented work on agent communication in Bayesian networks. He also held a series of lectures on data mining in the field of genetics, in which the used techniques are supposedly general enough to be applied in more general settings (e.g., text mining).

Thomas D. Nielsen worked for six months in 2000 at Université de Paris, together with Professor Jean-Yves Jaffray on an operational approach for rational decision making based on non-expected utility. In particular, the work considered decision aiding where the non-expected utility criterion was that of Rank Dependent Utility, and the rationality requirement was that no first order stochastically dominated strategy should be proposed to the decision maker.

With Industry

The research in normative systems at Aalborg University (involving researchers from the Department of Computer Science as well as researchers from the Department of Mathematics), which began in 1985, resulted in the establishment of the company Hugin Expert A/S in 1989, which produces and sells the market-leading tool for constructing and running Bayesian networks and influence diagrams. Most of the research engineers and programmers employed at Hugin have earned their M.Sc. or Ph.D. degrees from the Department of Computer Science, and thus the ties between the research unit and Hugin are quite close. For example, several employees at Hugin participate in the workshops and seminars held in the unit, through which a natural channel for the transfer of research results to Hugin has been established.

As a result of the successful collaboration with the unit, Hewlett-Packard has established a separate research and development laboratory at Aalborg University, which currently employs 3 full-time research engineers and 3 part-time programmers, and HP expects to continue to expand its presence there. Also, as a spin-off of the success of the SACSO project, Hewlett-Packard has purchased 45% of the shares in the Hugin company.

In 1998, Hewlett-Packard hosted a 3-day seminar on normative systems in Aalborg, attended by researchers and engineers from various organisations within HP and by members of the research unit. The participants learned about the Bayesian-network research going on at the university and about the SACSO project, as well as shared some of their experiences with other reasoning technologies and discussed how Bayesian-network technology and the work going on in Aalborg might be incorporated into their own work and projects.

Uffe Kjærulff has acted as a consultant to the local company MindPass A/S in Aalborg, which is a provider of technology for search engines at the World Wide Web.

In August 1999 the research unit gave three lectures as part of the Life-long

Education Programme at Aalborg University, covering results of the research conducted in the SACSO project. The lectures were attended by approximately 10 computer engineers and computer scientists from industry.

Finn V. Jensen, Kristian G. Olesen, and Uffe Kjærulff have acted as consultants to the local company TargIT A/S in Hjørring, which is a provider of data-mining technology.

6.5 Ph.D. Projects

6.5.1 Overview

In the period from 1996 to 2000, the unit has supervised five Ph.D. students, out of which three have graduated in the period, one is expected to graduate in 2001, and one in 2002.

6.5.2 Project Descriptions

Blocking Gibbs Sampling for Inference in Large and Complex Bayesian Networks with Applications in Genetics

Name: Claus Skaanning
Education: Cand. polyt. (M.E.) in Computer Engineering,
Aalborg University, 1993
Duration: February 1994 – May 1997
Status: Degree awarded September 1997
Funding: ODIN and Department of Mathematics and Computer Science
Advisor: Uffe Kjærulff

Due to the NP-hardness of inference in Bayesian networks, it has been necessary to use approximate inference methods when dealing with large and complex networks. One such method is Gibbs sampling, which obtains dependent realisations from a Bayesian network by sampling one component at a time, conditional on the remaining components. One such component can consist of one or more variables of the network.

Single-site Gibbs sampling is a variant of Gibbs sampling, where only one variable is sampled at a time conditional on the remaining variables. This method is simple to implement and has been used extensively in research. It is, however, prone to very slow (or zero) convergence as the complexity of the networks increases. However, if strongly correlated variables are sampled simultaneously, much faster convergence can be attained. This principle is used in blocking Gibbs sampling, which combines Gibbs sampling with exact local computations in junction trees. Blocking Gibbs sampling at the same time attempts to minimise the time spent at simulating the components with exact local computation,

and maximise the rate of convergence — two opposed criteria. The study is aimed at analysing, optimising, and implementing this algorithm.

Gibbs sampling has been used very extensively in the area of genetics for very computationally demanding methods of pedigree and linkage analysis. Here, the problems of slow (or zero) convergence have been present as in many other fields, and, in addition, due to the near-logical conditional probability tables used in genetics, it is hard to determine when the induced Markov chain will be irreducible. The study is also aimed at solving this and additional problems experienced in genetics.

Finding Sets of Best Responses for Normative Systems

Name: Dennis Nilsson
Education: Cand. scient. oecon. (M.Sc.), University of Copenhagen
Duration: July 1995 – August 1998
Status: Degree awarded March 1999
Funding: The Faculty of Engineering and Science, Aalborg University
Advisor: Finn V. Jensen

The thesis deals with the well-known problem of finding not only the most probable configuration, but the top m most probable configurations in probabilistic expert systems. An algorithm is given which utilises the standard max-propagation method in junction trees and clique tables constructed through a so-called max-propagation. The resulting algorithm is more efficient than other known algorithms for the same problem. Furthermore, a technique for forecasting future decisions in decision problems represented as traditional influence diagrams is given. It is shown how this technique can be used to analyse the strategies in more depth. Several examples for further use of the established Bayesian network are given (probabilities of past observations and decisions, risk profile, variance of total utility, etc.). Finally, the technique and algorithm mentioned above is combined to produce an algorithm for finding the m best strategies in decision problems represented as influence diagrams.

All Good Things Come to Those Who Are Lazy

Name: Anders L. Madsen
Education: Cand. scient. (M.Sc.) in Computer Science and Mathematics, Aalborg University, 1996
Duration: February 1997 – September 1999
Status: Degree awarded January 2000
Funding: Danish Natural Science Research Council, grant #9601649
Advisor: Finn V. Jensen

The framework of Bayesian networks is an increasingly popular knowledge representation framework for reasoning under uncertainty. A Bayesian network is a

compact representation of interactions among a set of uncertain entities (random variables) in a domain.

In this thesis a junction tree based inference architecture exploiting the structure of the original Bayesian network and independence relations induced by evidence to improve the efficiency of inference is presented. The efficiency improvements are obtained by maintaining a multiplicative decomposition of clique and separator potentials. Maintaining a multiplicative decomposition offers a tradeoff between off-line constructed junction trees and on-line exploitation of barren variables and independence relations induced by evidence. The architecture is called *lazy propagation*.

The efficiency improvements offered by lazy propagation are emphasised via empirical evaluations involving large real-world Bayesian networks. The time and space performance of lazy propagation is compared with non-optimised implementations of the Hugin and Shafer-Shenoy inference architectures.

Lazy propagation increases the computational efficiency of performing a number of commonly performed Bayesian network tasks. Furthermore, lazy propagation has been extended to exploit structure within conditional probability distributions to improve the efficiency of inference, to solve Bayesian decision problems, and to perform inference with continuous variables.

Poorly Structured Decision Graphs

Name: Thomas D. Nielsen
Education: Cand. scient. (M.Sc.) in Computer Science and Mathematics,
Aalborg University, 1998
Duration: August 1998 – July 2001 (expected)
Status: In progress
Funding: CIT
Advisor: Finn V. Jensen

Influence diagrams serve as a powerful modelling language for symmetric decision problems with a single decision maker. When formulating a decision problem as an influence diagram, a linear temporal order over the decision variables is required; this constraint ensures that the decision problem is well-defined. However, the structure of a decision problem often renders certain decisions conditionally independent, and it is therefore unnecessary to impose a linear temporal ordering on the decisions. In this thesis we deal with partial influence diagrams i.e. influence diagrams with only a partial temporal ordering specified. We present a set of conditions which are necessary and sufficient to ensure that a partial influence diagram is well-defined. These conditions are used as a basis for an algorithm to determine whether or not a partial influence diagram is well-defined. Based on these conditions we also give two different algorithms for solving influence diagrams. Both of these algorithms ensure that no strategy contains redundant information; for both modelling and computational reasons, it is important not

to deal with such redundancy.

The efficiency of influence diagrams as a modelling language also requires that the decision problem is symmetric. Unfortunately, most real-world decision problems are asymmetric, and modelling an asymmetric decision problem using influence diagrams can be a cumbersome task which usually fails to provide a compact representation of the decision problem in question. We propose a formal framework, termed *asymmetric influence diagrams*, that is based on influence diagrams and allows an efficient representation of asymmetric decision problems. We give an algorithm for solving asymmetric influence diagrams. The algorithm works by decomposing the asymmetric decision problem into a set of symmetric subproblems that can be solved using existing evaluation methods; a solution to the original problem can then be found by combining the solutions of the “smaller” symmetric subproblems.

Finally, we consider an operational approach for decision aiding based on non-expected utility. This type of decision aiding combines a prescriptive model with a descriptive model of the decision maker, in order to accommodate the situations where the decision maker rejects the model being proposed because it can not account for all of his choices (this is usually the problem with models based on expected utility, e.g., influence diagrams).

Object-Oriented Normative Systems

Name: Olav Bangsø
Education: Cand. scient. (M.Sc.) in Computer Science and Mathematics,
Aalborg University, 1998
Duration: August 1999 – July 2002 (expected)
Status: In progress
Funding: Hewlett-Packard
Advisor: Uffe Kjærulff

Bayesian networks for large and complex domains are difficult to construct and maintain. For example, modifying a small network fragment in a repetitive structure might be very time consuming. Top-down modelling may simplify the construction of large Bayesian networks, but methods (partly) supporting top-down modelling have only recently been introduced and tools do not exist. In this thesis a top-down approach to constructing Bayesian networks by using object oriented methods is taken. This provides a new framework, object-oriented Bayesian networks that allows top-down methodologies for the construction of Bayesian networks, provides an efficient class hierarchy, and a compact way of specifying and representing temporal Bayesian networks.

A method for parameter learning in these object-oriented Bayesian networks is constructed as well and it is proven that maintaining the object orientation imposed by the prior model will increase the learning speed in object-oriented domains. This method will also efficiently estimate the probability parameters in

domains that are not strictly object oriented. A method for structural learning given a partial object-oriented specification is defined as well.

The framework is expanded to incorporate object-oriented influence diagrams, and a method to ensure well-definedness is proposed.

6.6 Services

Editorial Service

- Program co-chair for the *Twelfth Conference on Uncertainty in Artificial Intelligence*, Portland, Oregon, 1996 (FVJ)
- Area chair for the *Fifth European Conference on Symbolic and Quantitative Approaches to Uncertainty*, London, 1999 (FVJ)
- Area chair for the *Sixteenth Conference on Uncertainty in Artificial Intelligence*, Stanford, California, 2000 (FVJ)

Conference and Workshop Organisation

Program Committees

- *Twelfth Conference on Uncertainty in Artificial Intelligence (UAI'96)*, Portland, Oregon 1996 (UK)
- *Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI'97)*, Providence, Rhode Island 1997 (FVJ,UK)
- *Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI'98)*, Madison, Wisconsin 1998 (UK)
- *Seventh International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'98)*, Paris, France 1998 (KGO)
- *Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI'99)*, Stockholm, Sweden 1999 (FVJ,UK)
- *International Joint Conference on Artificial Intelligence (IJCAI'99)* (FVJ)
- *Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI'00)*, Stanford, California 2000 (KGO,UK)

Refereeing for Journals and External Refereeing for Conferences

- *Networks* (FVJ, KGO)
- *Biometrika* (UK)
- *Computers and Electronics in Agriculture* (UK)
- *Journal of AI Research* (UK)
- *Journal of Applied Intelligence* (UK)
- *The Artificial Intelligence Journal* (FVJ, KGO, UK)
- *IEEE Transactions on Pattern Analysis and Machine Intelligence* (FVJ, KGO, UK)
- *IEEE Transactions on Systems, Man and Cybernetics* (KGO, UK)
- *Artificial Intelligence Review* (KGO)
- *The Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU'99)* (KGO, UK)
- *The Sixteenth International Joint Conference on Artificial Intelligence (IJCAI'99)* (KGO, UK)
- *The Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI'99)* (KGO)
- *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* (KGO)
- *European Journal of Operational Research* (FVJ)
- *Journal of Applied Stochastic Models in Business and Industry* (FVJ)
- *Computational Intelligence* (FVJ)
- *IEEE Transactions on Knowledge and Data Engineering* (FVJ)
- *The Fourth European Conference on Symbolic and Quantitative Approaches to Uncertainty (ECSQARU'97)* (FVJ)
- *The Ninth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'98)* (FVJ)
- *European Conference on Artificial Intelligence (ECAI'98)* (FVJ, UK)

Refereeing for Funding Organisations

- *The Swedish Research Council for Engineering Sciences* (FVJ)
- *The Netherlands Organisation for Scientific Research* (FVJ)

Other Academic Services for External Institutions

- Member of the board of the *Danish Informatics Network on Agriculture* (FVJ)
- Career track evaluation for a researcher at University of Regina (FVJ)
- Career track evaluation for a researcher at University of California, Los Angeles (FVJ)
- Career track evaluation for a researcher at Technion, Israel (FVJ)
- Ph.D. evaluation committee for Steen M. Kristensen, Aalborg University (Department of Medical Informatics and Image Analysis), 1996 (FVJ)
- Participated as invited lecturer at a summer school and workshop on graphical models, organised by the Italian Research Council, Erice, Italy, 1996 (UK)
- A one week Ph.D. course on normative systems organised by the *European Association for Artificial Intelligence*, Vilnius, 1997 (FVJ)
- Ph.D. evaluation committee for Ross M. Curds, University College London, 1997 (FVJ)
- Participated as invited lecturer at a one week workshop on handling uncertainty, organised by the Italian Research Council, Erice, Italy, 1997 (FVJ)
- Invited talk at the Laboratory of Intelligent Systems, Economic University (VSE), Prague, 1998 (UK)
- Ph.D. evaluation committee for Nathalie Jitnah, Monash University, Melbourne, 1999 (FVJ)
- One week Ph.D. course on normative systems, Fields Institute, Toronto 1999 (KGO, UK)
- Invited talk at *Workshop on Prognostic Models in Medicine — Artificial Intelligence and Decision Analytic Approaches* in conjunction with *The Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making, AIMDM'99*, Aalborg, Denmark, 1999 (KGO)

- A three day Ph.D. course on normative systems for students in the agricultural sciences, Århus, 1999 (FVJ)
- Invited talks at University of Stockholm 1999, University of Oslo 2000, Danish Technical University 1999, The Pharmaceutical Association of Switzerland, Basel 1999 (FVJ)
- Ph.D. evaluation committee for Veerle Coupé Erasmus University, Rotterdam, 2000 (FVJ)

6.7 The Unit's Own Evaluation

Research in decision support has three components

- development of tractable methods for computerised management of uncertainty,
- efficient implementation of general methods, and
- construction of specific systems.

All three components are interdisciplinary in nature. Therefore, the success of the unit is not only dependent on the creativity and industry of the members of the unit, but also on the opportunities given for interdisciplinary work.

At the research evaluation five years ago, we therefore concluded that the unit lacked projects with enthusiastic domain experts taking part, and we would put substantial effort into selecting a project with all three parties participating enthusiastically: domain experts, normative systems specialists, and domain users.

Fortunately, the SACSO project with Hewlett-Packard met the needs, and for a couple of years we have had at our premises all three kinds of participants working on a common project. Also, the project had challenges of all three kinds listed above. This has created a very positive and fruitful working environment. The results of the SACSO project are now being disseminated in the HP organisation — and reported in scientific papers.

Over the five-year period 1998–2003, the unit receives an annual donation of one million DKK from Hewlett-Packard, which — among other things — has enabled the recruitment of three research assistants and the establishment of a Ph.D. programme. Thus, the donation has enabled a significant increase in the research activities of the unit.

The research activities of the unit — and especially in the SACSO project — have been coordinated and stimulated through workshops held every week. These workshops have proven to be very valuable in initiating, coordinating, and stimulating research activities.

The research plans set out in the previous research evaluation five years ago included

- good application projects,
- use of modern information technology (multimedia and multiprocessor systems),
- development of flexible inference methods, and
- expanding of scope to multi-agent decision scenarios.

From the point of view of this plan, we can be fairly satisfied with the results achieved. As explained above, the SACSO project has fulfilled our goal to have an application project with enthusiastic domain experts and users involved. Also, flexible inference methods have been developed, notably lazy evaluation methods for both Bayesian networks and influence diagrams. Regarding the use of modern information technology, a parallel version of the state-of-the-art inference method has been developed, whereas there still is work to be done regarding use of multimedia. Also, we still are facing the challenge of expanding our scope to multi-agent decision scenarios.

In parallel with the SACSO project, the unit has taken up various theoretical challenges. Particularly, we have worked with smooth and efficient graphical representation of decision scenarios, namely representation languages, which can also easily be compiled for computer treatment and the decision problem solved through efficient algorithms.

Conclusively, we are very satisfied with the results achieved, and it is our impression that the unit is well respected in the field of computerised management of uncertainty.

We have over the last five years experienced a rapidly growing scientific as well as industrial interest in normative systems. As a result, the unit has been very fortunate to attract sufficient funding for its work. Actually, the unit has, due to lack of qualified personnel, declined several invitations to participate in externally sponsored projects, and at present the lack of personnel is a pressing problem. Although application projects are a tremendous source of good research problems, it is also vital for the unit to have “free” research resources not engaged in particular needs of specific application projects. To meet the demand, we have attracted several foreign researchers as well as Ph.D. students, and we advertise internationally for more researchers. With current funding the unit could employ at least three more.

The unit’s collaboration outside the department is good, whereas the collaboration with other units at the department should be improved. Also, we have not been sufficiently successful in recruiting Ph.D. and M.Sc. students to our field. Although there is a lot of work to be done on normative systems in research as

well as in industry, the field may be considered too narrow from a student's point of view. However, at least two factors may lead to a change in that situation over the coming five-year period: First, the change of focus of the unit from methodological-oriented research to more application-oriented — and especially Internet-based applications — most probably will lead to an increased interest among students. Second, the increasing awareness in industry of the “Bayesian approach” to data mining and decision support is also likely to be influential on the students' choices.

6.8 Research Plan 2001–2005

Research tasks

The unit intends to continue the work on methods for handling uncertainty in artificial intelligence. Also, the main methodological focus of interest will continue to be on normative methods, but other soft computing methods (fuzzy logic, possibility theory, genetic algorithms, neuro-computing) may be applied.

We are currently considering a new orientation of our research in the light of recent technological developments. We intend to redirect our application focus to the internet, and tentatively we use the term wise and adaptive agents for the internet for the *specialité de la maison* to be. Particularly, the unit will work with agents for e-commerce and e-services. In order to meet this goal we are during the first year undertaking a self-education programme to build up expertise on e-commerce and e-services.

Based on this strategy and the successful collaboration in the SACSO project, the research unit and Hewlett-Packard are currently (autumn 2000) in the process of formulating the basis for a Phase-III project. In this phase, the intention is to expand the research conducted on troubleshooting systems to enable a broad family of intelligent “e-services” for generalised *decision support services*. This will most probably require development of methods for

- extracting key information from semi-structured or even unstructured text documents (*text mining*)
- *service brokering* to select the most appropriate e-service, based upon a set of attributes that characterise the desired service
- *object-oriented construction and learning* of normative models from configuration data and object-oriented model templates
- *learning of models from semi-structured databases*, requiring development of data/text mining techniques

among others.

A side interest, which may flourish, is connected to game theory. Some preliminary theoretical investigations and experiments have been performed on using Bayesian adaptation methods to let competing agents converge to a Nash-equilibrium. The preliminary results seem promising, and we may pursue this further, if we can establish appropriate industrial contacts. The purpose of such contacts are twofold. First, we wish to apply and test our developments in realistic settings and, second, experience has shown that committed resources are needed in order to realise larger program complexes. It seems, unfortunately, that public employment is not sufficiently attractive for a professional programmer.

The unit has started collaboration with the research unit of Database and Programming Technologies. So far, the collaboration is limited to two new educations, International Master in Knowledge and Data Engineering and Cand. scient. in Knowledge Discovery and Data Mining (Avanceret dataanalyse) (the latter also in collaboration with the Department of Mathematics). We expect this collaboration to extend to also comprising research — particularly in data mining.

The unit is a member of the consortium for the ESPRIT project MERIT. The project, which starts up in 2001, aims at developing an integrated water resource management methodology suitable for use at the river basin scale throughout Europe. The approach is to develop a generic Bayesian network with normative decision features, and to instantiate the network to four catchments in the UK, Denmark, Spain, and Italy.

Organisation and funding

The unit seems to be adequately funded for the next three years, where the constraining factor is qualified personnel. We expect to have a steady state of three permanent positions, one assistant professor (adjunkt) and three Ph.D. students. Besides, we hope to have three to four externally funded researchers as members of the unit.

One of the reasons why the HP-collaboration has been successful is that the HP researchers working with the project have been situated at the same place as the unit. This supports a rapid spreading of ideas and considerations to the entire group, and it also prohibits a “them-and-us” situation. We expect this situation to continue, and in case of other industrial collaboration projects we will work persistently on obtaining the same situation.

Goals for the next five years

The five-year goals for the unit are divided into three areas: scientific impact, broadness of scope, and staffing.

Scientific impact With respect to scientific publications we aim at maintaining and possibly improve the current productivity. In raw numbers, we expect in average for the scientific staff yearly to produce two papers out of which one is published in a first class journal or presented at a top class conference, such as UAI , IJCAI or ECSQARU.

We also wish to improve our industrial contacts, and we aim at establishing at least three industrial partnerships involving common projects.

The unit already has scientific contacts with other European groups working with normative systems. We wish to extend these contacts and possibly make them grow into a more formalized corporation with common projects, symposia, and exchange of staff and students.

Broadness of scope The goal is to broaden the scope of research to also deal with areas close to the current research, such as data mining, business intelligence, teams of intelligent agents, games. The unit will continue the activities with respect to establishing corporation with other research groups at Aalborg University, such as Statistics, Databases, Distributed Systems.

Staffing The goal is that by the end of 2005 the unit consists of 6 faculty members (associate and full professors), 5 post-docs (assistant professors or project funded) and 5 PhD students. To meet this goal we will intensify the search internationally for applicants to our positions. Particularly the unit will search for highly qualified researchers in areas contributing to broadening our scope.

6.9 Bibliography

The list of references below constitute the tangible results of the research conducted in the unit in the period 1996–2000.

Among researchers in *uncertainty in artificial intelligence* (UAI), the annual *Conference on Uncertainty in Artificial Intelligence* is regarded as the most preferred place to publish research papers. Consequently, papers published in the proceedings of the UAI conferences are ranking at the same level as journal papers. Therefore, we have devised a special section for our UAI publications below.

Please note that for papers published more than one place, only the more recent references have been included (e.g. a conference paper later published in a book of selected papers is not included under “Refereed conference articles”).

Books

- [1] Finn V. Jensen. *An Introduction to Bayesian Networks*. UCL London, 1996.

- [2] Finn V. Jensen. *Bayesian Networks and Decision Graphs*. Springer-Verlag, 2000. To appear.

Ph.D. Theses

- [3] Claus Skaanning Jensen. *Blocking Gibbs Sampling for Inference in Large and Complex Bayesian Networks with Applications in Genetics*. Phd thesis, The European Doctoral School of Technology and Science, Aalborg University, May 1997.
- [4] Anders L. Madsen. *All Good Things Come to Those Who Are Lazy*. Phd thesis, The European Doctoral School of Technology and Science, Aalborg University, December 1999.
- [5] Dennis Nilsson. *Finding Sets of Best Responses for Normative Systems*. Phd thesis, The European Doctoral School of Technology and Science, Aalborg University, August 1998.

Refereed Contributions to Books

- [6] Steen Andreassen, Leonard Leibovici, Henrik C. Schönheyder, Brian Kristensen, Christian Riekehr, Anders G. Kjær, and Kristian G. Olesen. A decision theoretic approach to empirical treatment of bacteraemia originating from the urinary tract. In W. Horn, Y. Shahar, G. Lindberg, S. Andreassen, and J. Wyatt, editors, *Artificial Intelligence in Medicine — Proceedings the Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making, AIMDM'99*, volume 1620 of *Lecture Notes in Artificial Intelligence*, pages 197–206. Springer-Verlag, 1999.
- [7] Finn V. Jensen. Bayesian graphical models. In *Encyclopedia of Environmetrics*. Wiley & Sons, 2000. To appear.
- [8] Finn V. Jensen. Influence diagrams. In *Encyclopedia of Environmetrics*. Wiley & Sons, 2000. To appear.
- [9] Finn V. Jensen and Steffen L. Lauritzen. Probabilistic networks. In *Handbook of Defeasible and Uncertainty Management*, volume 5, pages 289–320. Kluwer Academic Publishers, 2000.
- [10] Uffe Kjærulff. Inference in Bayesian networks using nested junction trees. In Michael I. Jordan, editor, *Learning in Graphical Models*, volume 89 of *NATO Science Series, Series D: Behavioural and Social Sciences*, pages 51–74. Kluwer Academic Publishers, 1998.

- [11] Dennis Nilsson and Finn V. Jensen. Probabilities of future decisions. In Bernadette Bouchon-Meunier, Ronald R. Yager, and Lotfi A. Zadeh, editors, *Information, Uncertainty and Fusion*, pages 161–171. Kluwer Academic Publishers, 2000.

Refereed Journal Articles

- [12] Claus Skaaning Jensen and Augustine Kong. Blocking Gibbs sampling for linkage analysis in large pedigrees with many loops. *American Journal of Human Genetics*, 65(3):885–902, 1999.
- [13] Claus Skaanning Jensen. A simple method for finding a legal configuration in complex Bayesian networks. *Statistics and Computing*, 8(3):243–251, 1998.
- [14] Claus Skaanning Jensen and Nuala Sheehan. Problems with the determination of the noncommunicating classes for MCMC applications in pedigree analysis. *Biometrics*, 54(2):416–425, 1998.
- [15] Finn V. Jensen. Bayesian network basics. *AISB Quarterly*, 94:9–22, 1996.
- [16] Steffen L. Lauritzen and Finn V. Jensen. Local computation with valuations from a commutative semigroup. *Annals of Mathematics and Artificial Intelligence*, 21:51–69, 1997.
- [17] Mogens Lund and Claus Skaanning. Blocking Gibbs sampling in the mixed inheritance model using graph theory. *Genetics, Selection and Evolution*, 31:3–24, 1999.
- [18] Anders L. Madsen and Finn V. Jensen. Lazy propagation: A junction tree inference algorithm based on lazy evaluation. *Artificial Intelligence*, 113:203–245, 1999.
- [19] Dennis Nilsson. Finding the m most probable configurations in probabilistic expert systems. *Statistics and Computing*, 2:159–173, 1998.
- [20] Marko Suojanen, Steen Andreassen, and Kristian G. Olesen. The EMG diagnosis — an interpretation based on partial information. *Medical Engineering and Physics*, 21:517–523, 1999.
- [21] Yang Xiang, Kristian G. Olesen, and Finn V. Jensen. Practical issues in modeling large diagnostic systems with multiply sectioned Bayesian networks. *International Journal of Pattern Recognition and Artificial Intelligence*, 14(1):59–71, 2000.

Conference on Uncertainty in Artificial Intelligence

- [22] Søren L. Dittmer and Finn V. Jensen. Myopic value of information in influence diagrams. In Geiger and Shenoy [60], pages 142–149.
- [23] Allan Leck Jensen and Finn V. Jensen. MIDAS — an influence diagram for management of mildew in winter wheat. In Eric Horvitz and Finn V. Jensen [59], pages 349–356.
- [24] Uffe Kjærulff. Nested junction trees. In Geiger and Shenoy [60], pages 294–301.
- [25] Uffe Kjærulff and Linda C. van der Gaag. Making sensitivity analysis computationally efficient. In Boutilier and Goldszmidt [63], pages 317–325.
- [26] Anders L. Madsen and Finn V. Jensen. Lazy propagation in junction trees. In Gregory F. Cooper and Serafín Moral [61], pages 362–369.
- [27] Anders L. Madsen and Finn V. Jensen. Lazy evaluation of symmetric Bayesian decision problems. In Laskey and Prade [62], pages 382–390.
- [28] Thomas D. Nielsen and Finn V. Jensen. Well-defined decision scenarios. In Laskey and Prade [62], pages 502–511.
- [29] Thomas D. Nielsen and Finn V. Jensen. Representing and solving asymmetric Bayesian decision problems. In Boutilier and Goldszmidt [63], pages 416–425.
- [30] Thomas D. Nielsen, Pierre-Henri Wuillemin, Finn V. Jensen, and Uffe Kjærulff. Using ROBDDs for inference in Bayesian networks with troubleshooting as an example. In Boutilier and Goldszmidt [63], pages 426–435.
- [31] Claus Skaanning. A knowledge acquisition tool for Bayesian-network troubleshooters. In Boutilier and Goldszmidt [63], pages 549–557.
- [32] Yang Xiang and Finn V. Jensen. Inference in multiply sectioned Bayesian networks with extended Shafer-Shenoy and lazy propagation. In Laskey and Prade [62], pages 680–687.

Other Refereed Conference Articles

- [33] Olav Bangsø and Pierre-Henri Wuillemin. Top-down construction and repetitive structures representation in Bayesian networks. In *Proceedings of the Thirteenth International FLAIRS Conference*, Florida, USA, 2000.

- [34] Søren L. Dittmer and Finn V. Jensen. Tools for explanation in Bayesian networks with application to an agricultural problem. In *Proceedings of the first European conference for information technology in agriculture, Copenhagen, Denmark*, June 1997.
- [35] Finn V. Jensen. Bayesian networks and influence diagrams. In *Proceedings from the EUNITA seminar on Risk management strategies in agriculture: State of the art and future perspectives*, Wageningen, The Netherlands, January 1996.
- [36] Finn V. Jensen. Gradient descent training of Bayesian networks. In *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU)*, Springer Lecture Notes Series in Artificial Intelligence, 1999.
- [37] Liang Jianming, Finn V. Jensen, and Henrik I. Christensen. A framework for generic object recognition with Bayesian networks. In *Proceedings of the First International Symposium on Soft Computing for Pattern Recognition*, Reading, UK, 1996.
- [38] Mogens Lund and Claus Skaanning. Multivariate updating of genotypes in a Gibbs sampling algorithm of the mixed inheritance model. In *Proceedings of the 6th World Congress on Genetics Applied to Livestock Production*, volume 25, pages 521–524, Armidale, Australia, 1997.
- [39] Anders L. Madsen, Lars M. Nielsen, and Finn V. Jensen. Probsy — a system for the calculation of probabilities in the card game Bridge. In *Proceedings of the Eleventh FLAIRS Conference*, Florida, USA, 1998.
- [40] Claus Skaanning, Finn V. Jensen, and Uffe Kjærulff. Printer troubleshooting using Bayesian networks. In *Proceedings of the Thirteenth International Conference on Industrial and Engineering Applications of AI and Expert Systems (IEA/AIE-2000)*, 2000.
- [41] Claus Skaanning, Finn V. Jensen, Uffe Kjærulff, and Anders L. Madsen. Acquisition and transformation of likelihoods to conditional probabilities for Bayesian networks. In *AAAI Spring Symposium*, Stanford, USA, 1999.
- [42] Claus Skaanning, Finn V. Jensen, Uffe Kjærulff, Paul Pelletier, and Lasse Rostrup-Jensen. Printing system diagnosis: A Bayesian network application. Workshop on Principles of Diagnosis, Cape God, Massachusetts, May 1998.
- [43] Marta Sochorová and Jiří Vomlel. Troubleshooting: NP-hardness and solution methods. In *Proceedings of The Fifth Workshop on Uncertainty Processing (WUPES'2000)*, pages 198–212, Jindrichuv Hradec, Czech Republic, 2000.

Technical Reports

- [44] Johan M. Andersen. Parameter estimation for conditional Gaussian Bayesian networks with missing data. Technical report, Department of Computer Science, Aalborg University, Denmark, September 1999.
- [45] Steen Andreassen, Marko Suojanen, Bjørn Falck, and Kristian G. Olesen. Improving the diagnosing performance of MUNIN by remodelling of the diseases. Technical report, Department of Medical Informatics, Aalborg University, 2000. Submitted.
- [46] Olav Bangsø and Pierre-Henri Wuillemin. Object oriented Bayesian networks, a framework for top-down specification of large Bayesian networks with repetitive structures. Technical report, Department of Computer Science, Aalborg University, Denmark, 2000.
- [47] Veerle M. H. Coupé, Finn V. Jensen, Uffe Kjærulff, and Linda C. van der Gaag. A computational architecture for n-way sensitivity analysis of Bayesian networks. Technical report, Department of Computer Science, Utrecht University, 2000.
- [48] Finn V. Jensen, Claus Skaanning, and Uffe Kjærulff. The SACSO system for troubleshooting of printing systems. Technical report, Department of Computer Science, Aalborg University, Denmark, February 2000.
- [49] Uffe Kjærulff and Marta Sochorová. Sensitivity analysis in Bayesian troubleshooting models. Technical report, Department of Computer Science, Aalborg University, Denmark, August 2000.
- [50] Helge Langseth. Conditional cost in the SACSO troubleshooter. Technical report, Department of Computer Science, Aalborg University, Denmark, March 2000.
- [51] Helge Langseth. Conditions for the SACSO data to be missing at random. Technical report, Department of Computer Science, Aalborg University, Denmark, March 2000.
- [52] Anders L. Madsen and Finn V. Jensen. Parallelization of inference in Bayesian Networks. Technical Report R-99-5002, Department of Computer Science, Aalborg University, Denmark, April 1999.
- [53] Anders L. Madsen, Kristian G. Olesen, and Søren L. Dittmer. Practical modeling of Bayesian decision problems — exploiting deterministic relations. Technical report, Department of Computer Science, Aalborg University, Denmark, 2000. Submitted.

- [54] Thomas D. Nielsen and Finn V. Jensen. Representing and solving asymmetric Bayesian decision problems. Technical Report R-99-5010, Department of Computer Science, Aalborg University, Denmark, November 1999.
- [55] Kristian G. Olesen and Anders L. Madsen. Maximal prime subgraph decomposition of Bayesian Networks. Technical Report R-99-5006, Department of Computer Science, Aalborg University, Denmark, June 1999.
- [56] Marko Suojanen, Kristian G. Olesen, and Steen Andreassen. A method for diagnosing multiple disorders in MUNIN. Technical report, Department of Medical Informatics, Aalborg University, 2000. Submitted.
- [57] Marco Valtorta, Y.-G. Kim, and Jiří Vomlel. Soft evidential update for probabilistic multi-agent systems. Submitted to *International Journal of Approximate Reasoning*, 2000.
- [58] Jiří Vomlel. On quality of BATS troubleshooter and other approximative methods. Technical report, Department of Computer Science, Aalborg University, Denmark, July 2000.

Cross references

- [59] Eric Horvitz and Finn V. Jensen, editors. *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, 1996.
- [60] Dan Geiger and Prakash P. Shenoy, editors. *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, 1997.
- [61] Gregory F. Cooper and Serafín Moral, editors. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, 1998.
- [62] Kathryn B. Laskey and Henri Prade, editors. *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, 1999.
- [63] Craig Boutilier and Moisés Goldszmidt, editors. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers, 2000.

Software

Loke is an interactive software tool for investigating Partial Influence Diagrams (PIDs) as described in [28]. Loke supports the specification of PIDs and can be used to determine whether or not a PID is well-defined. If the PID is not well-defined it gives suggestions for further specification of the temporal ordering; the suggested refinement is chosen s.t. it maximises the expected utility of the decision scenario being modelled.

Loke has been implemented using the multiplatform environment wxWindows but so far it has only been tested with Redhat Linux.

Chapter 7

Distributed Systems and Semantics

7.1 Profile of the Unit

The research of the unit concerns modelling, analysis and realization of computer programs, with an emphasis on distributed and embedded systems. This includes the following areas:

- Semantic theories for modelling the behaviour of computer programs and systems.
- Algorithms, methods and tools for verification and validation of programs and systems.
- Design, implementation and models for analysis of distributed systems and networks.

In its present form the Distributed Systems and Semantics unit is no more than two years old. The unit was created by a fusion of the two previous groups of Computer Systems and Formal Systems motivated by a growing collaboration between the two groups. Consequently, the research of the unit spans a wide spectrum ranging from development of the semantic foundations of computing and concurrency to contributions to the technological state of the art within real-time and networking systems. To conduct quality research within such a wide spectrum is a challenging task. However, the common emphasis on distributed and concurrent systems provides a point of focus for the activities within the unit.

Each of the three research areas mentioned above constitutes a subject in its own right. Moreover, the areas are interrelated in a number of ways such as:

- Semantic models offer important guidelines for the development of languages and paradigms for distributed systems. Moreover, they are necessary prerequisites for the development of verification algorithms.
- Experience gained from the development of verification and validation tools provides new insight into the underlying semantic models. Furthermore, the algorithms employed by such tools are applied in environments for the construction and analysis of distributed systems.
- Paradigms and languages for distributed systems provide insight into the strengths and weaknesses of existing semantic models, and serve as an inspiration for the development of new ones. Furthermore, distributed systems truly expose the limits of verification algorithms.

The formation of one single unit has been a first decisive step towards exploiting these interrelationships. The current research of the unit includes the following activities:

Process Calculi: Semantic theories and meta-theories for concurrent processes and their logical properties. Study of non-classic extensions such as hybrid, real-time, and probabilistic processes.

Calculi for Mobility and Objects: Semantic theories for processes whose communication topology changes dynamically, including object-oriented programs and security protocols.

Verification and Validation: Development and implementation of data structures, algorithms and tools for model-checking state/event systems and real-time systems. Applications to communication protocols and control programs.

Real-Time Systems: Modelling and analysis of industrial case studies. Specification and testing methodology. Theoretical foundations of hybrid systems.

Networks: Analysis and construction of services and protocols for computer networks.

In the next section we provide a detailed account of the research of the unit structured along the activities above, while stressing overlaps and mutual links.

7.2 Activities and Results

7.2.1 Research Results

We now proceed to present the main results of the research carried out within the unit in the period of this evaluation. The presentation will be structured along the aforementioned main research activities.

Process Calculi

The Distributed Systems and Semantics research unit has continued its contributions to the development and dissemination of results on the theory of processes. The relevance and visibility of the process algebraic work carried out by the research unit is witnessed, amongst other things, by the three chapters contributed by its members to the forthcoming *Handbook of Process Algebra* [7, 8, 9]. Apart from the continuing development of the theory of behavioural equivalences for process description languages and of their connections with logical specification formalisms (see, e.g., [19, 43, 64]), our main contributions to process algebra may be found in the study of the equational logic of processes, the meta theory of process description languages, and probabilistic, timed and value passing extensions of process algebras.

The Equational Logic of Processes Much of the research on process algebra carried out during the evaluation period has focussed on the study of the expressive power of equational logic in the characterization of behavioural congruences for (fragments of) algebraic process description languages, with particular emphasis on the study of process algebras with iteration operators. Along this line of investigation, we have obtained both positive and negative results. The positive results take the form of complete equational axiomatizations of notions of behavioural congruence for several process description languages (see, e.g., [13, 16, 39, 17, 42, 21]), and are the more interesting the richer the language that is being axiomatized. The negative results are to the effect that no finite equational axiomatization can exist for some notion of behavioural congruence over some given process description language (see, e.g., [38, 14, 40, 17]), and are the more interesting the weaker the studied language.

In the process of developing the aforementioned negative results, we have discovered fruitful and interesting connections between process algebraic results and techniques, and work on formal languages, automata theory and max-plus algebra. In particular, the techniques applied in [14] to prove the non-existence of finite equational axiomatizations over BPA* for many behavioural congruences in the linear time-branching time spectrum stem from formal language theory, and have been re-applied, in a sharpened form, to solve a long-standing problem of Salomaa's in the theory of formal languages (see [15]). Similarly, the problems that motivated the research reported in [12, 38] are rooted in process algebra, but turned out to allow us to solve open questions in the field of max-plus algebra.

Meta Theory The meta theory of process algebras aims at contributing to the systematic development of process theory by offering results that hold for whole classes of process description languages. As these languages are often equipped with a Plotkin-style Structural Operational Semantics (SOS), this way of giving semantics to processes provides a natural handle to establish results that hold

for all languages whose semantics is given by means of inference rules that fit a certain format. Examples of results of this type and surveys of the theory may be found in, e.g., [6, 7, 18]. The research reported in [18] is based on, and has spurred the development of, the studies [19, 20, 66, 25], which explore the connections between domain-theoretic models of processes and behavioural models based upon variations on bisimulation.

Models and Logics for Timed Systems The work on real-time extensions of process algebras has focussed on contributions to the understanding of real-timed versions of bisimulation equivalence for Wang Yi’s TCCS, including characterizations of induced congruences [30]. Additionally, work has been done on the development of a calculus based on durational actions [21]. A second line of research has focused on a real-time extension of Kozen’s modal mu-calculus in terms of expressiveness, (conditional) decidability, complexity [44], as well as quotienting with respect to parallel composition of real-time processes.

These compositionality results have provided the background for the development of the compositional model-checking tool CMC [76] reported on in Sect. 7.2.1. The compositionality results are also exploited in extending the model-checking power of UPPAAL, which solely offers the possibility of checking simple reachability properties. In the papers [37, 36] we provide a logical characterization of the expressive power obtained by checking for reachability of a real-timed system in the context of a testing component.

Work on extensions of the real-time theory towards more general models suitable for the description of hybrid systems is offered in [28, 53].

Finally, much of the unit’s other research on the theory of real-time processes is carried out within the framework of Alur and Dill’s timed automata and the tool UPPAAL, and is reported in Sect. 7.2.1.

Probabilistic and Value Passing Extensions With respect to other extensions of standard process algebras, we have continued our study of their probabilistic and value passing extensions. Previous work on probabilistic and value passing process algebras has found its climax in the surveys [8, 9] and the invited paper [79]. In addition, we have begun new work on verified implementations of value passing process algebras [65].

Calculi for Mobility and Objects

The family of pi-calculi provides a rich mathematical theory for studying the properties of systems whose communication properties change dynamically. Prime examples of such systems are object-oriented programs and security protocols, and these two topics have provided important inspiration for research within the unit.

The research on behavioural theories for object-oriented programs has focused on the family of object calculi proposed by Abadi and Cardelli.

In [61] Hüttel, Kleist, Andersen and Pedersen started out from the labelled transition semantics for the functional object calculus and established a correspondence between a first-order recursive type system for this calculus and a sublogic of the modal mu-calculus.

Abadi and Cardelli have defined a denotational semantics for the aforementioned functional object calculus; in [41], Aceto, Hüttel, Ingólfssdóttir and Kleist established that this denotational semantics is computationally adequate, but not fully abstract, with respect to contextual equivalence.

Encodings into pi-calculi have been an important strategy in reasoning about the behavioural properties of objects.

In [74] Kleist and Sangiorgi have considered an encoding of the imperative object calculus into a label-selective pi-calculus. This encoding enabled them to prove soundness properties of the equational theory of the imperative object calculus.

In the distributed object-oriented language *Obliq*, migration of objects is essential. In a setting without explicit distribution over sites, migration can be understood as surrogation—an object is cloned, and all requests made to the original object are then redirected to the clone. Nestmann, Kleist, Hüttel and Merro have studied the correctness properties of surrogation in *Øjeblik*, a subset of *Obliq*. This involved defining a type system for *Øjeblik* and a number of operational semantics with different migration characteristics in [93]. In [73], Nestmann, Kleist and Merro provided an encoding of *Øjeblik* into a label-selective pi-calculus; using the encoding, they showed that surrogation was safe.

TyCO is a typed calculus for concurrent objects that was much inspired by the original pi-calculus. In fact, it is quite similar to the label-selective pi-calculus that was developed later by Sangiorgi. Nestmann and Ravara investigate the formal relationship between these two calculi; the first results underlining their equally expressive power were presented at EXPRESS'00. Ultimately, this work addresses a non-uniform interpretation of the underlying standard type system, where standard types may dynamically change over time. This extension is only possible in *TyCO*, but not in the label-selective pi-calculus.

Abadi and Gordon proposed the spi-calculus, which extends the pi-calculus with syntactic primitives for encryption and decryption. The aim is to provide a mathematical theory for reasoning about security protocols, and among the important notions of the theory is the notion of framed bisimilarity. In [55], Hüttel, Elkjær, Overgaard Nielsen and Hoehle have established an alternative characterization of framed bisimilarity.

Verification and Validation

One of the most significant research contributions of the unit during the period has been within the area of verification and validation of real-time and embedded software systems. For many of these systems correctness is of great importance as faults are often very costly or impossible to repair, and may lead to great material and human losses. Furthermore, these types of systems tend to be of an extraordinary complex nature, because of real-time aspects, concurrent and hierarchical structure, and interaction with the physical environment.

The main contributions of the unit are associated with the verification tools UPPAAL and VISUALSTATE, and in particular with the development and realization of the underlying algorithmic bases of these tools. In both cases, the developments have led to performance superior to other verification tools within the same application domains (real-time and embedded systems). A strongly related and integrated research activity has been the continuous application of the tool UPPAAL to industrial case studies. For a more detailed description of this research we refer to the section on Real-Time Systems.

A characteristic of the research carried out in this area has been the strong involvement of students. In the period of this evaluation, three Ph.D. theses [101, 102, 104] have been completed within the area and an additional three are in progress. Also, a large number of MSc students has taken part via their projects and as employed programmers. Though the algorithmic techniques all have firm theoretical backgrounds and justification, the inherent computational hardness of the problems considered¹ necessitates experimentation in order to firmly demonstrate *practical* applicability of a suggested technique. Here, the involvement of students as well as access to realistic case studies have been imperative. Another characteristic of the research conducted within this area has been the high degree of external collaboration both at national and international levels.

Below we present the main past achievements in greater detail.

The UPPAAL Tool UPPAAL is a tool for modelling, simulation and verification of real-time systems, developed in collaboration with the Department of Computer Systems at Uppsala University since the beginning of 1995. The tool allows for the modelling of real-time systems as collections of non-deterministic processes with finite control structure and real-valued clocks (so-called timed automata), communicating through channels and shared variables. During the period of this evaluation, the performance of the tool has been subject to substantial improvements: on the full collection of industrial case studies both space- and time-performance have improved by a factor of 10 every nine months². Since

¹In the best cases, the problems are NP-complete, but more often the problems dealt with are either PSPACE- or EXPTIME-complete.

²This factor does *not* include improvement in performance of machines and compilers.

its release in autumn 1999, the newest version, UPPAAL2K, has been downloaded by more than 700 users in 35 countries.

Tool development A main activity has been the continuous development of the UPPAAL tool itself. The effort has been two-fold: constantly working towards improved interfaces while simultaneously allowing experimentation with and giving access to the most recent (and efficient) algorithms and data structures for analysis. In the first versions of the tool [52, 50] only the editing of systems was graphically supported via an adaptation of the existing graphical tool AUTOGRAPH. In the autumn of 1996 a graphical interface for visualizing simulations was implemented [81, 27]. With this addition the tool certainly became interesting and more accessible for research groups not traditionally doing verification, and consequently the number of UPPAAL users increased substantially. The newest version of the tool, UPPAAL2K [51, 10], offers access to the most efficient verification and validation techniques via a common graphical interface to all components of the tool, i.e. system editor, simulator and verifier. To allow for the tool to run on several platforms the new version is designed as a client/server application with a verification server providing efficient C++ services to a Java client over a socket based protocol. This design also makes it possible to execute the server and the GUI on two different machines. We are currently establishing an XML-based interface between the GUI and the various tool components.

Algorithms and Data structures Another main line of research has been the development of data structures and algorithms for (efficient) symbolic representation and manipulation of the state spaces encountered during exploration of system models. In particular, symbolic data structures dealing with the infinite, continuous part of the state space due to the presence of clocks have been of primary concern. One such data structure is that of a *zone*, which allows certain simple subsets of the Euclidean space to be represented by a matrix of bounds on differences between clocks (so-called Difference Bounded Matrices or DBMs). Most operations required for state space exploration benefit from a canonical form obtained by computing the shortest-path closure of the graph associated with the DBM. Moreover, diagnostic traces can be synthesized as a result of a zone-based state-space exploration.

A major obstacle for state-space exploration is the so-called *state-space explosion problem*, i.e. the fact that the number of (symbolic) states grows exponentially in the number of processes, clocks and variables of the underlying model. The use of DBMs adds to this problem as each symbolic state will have size quadratic in the number of clocks. In [80] an alternative data structure for zones based on a shortest path *reduction* of the graph associ-

ated with a DBM is presented. The new data structure remains canonical while resulting in significant space and time savings.

Verification of untimed systems is well established, especially in the field of hardware verification. An important factor behind this success was the introduction of Reduced Ordered Binary Decision Diagrams (BDDs) by Bryant in 1986, which is a data structure allowing for the compact representation and efficient manipulation of boolean functions and (by boolean encoding) large, discrete state spaces. A long-standing problem has been how to best combine this data structure suitable for the discrete part of the state-space with a symbolic representation of the continuous part. In a sequence of papers [84, 29, 47] the data structure Clock Difference Diagrams (CDDs) is introduced as an answer to this problem. It generalizes at the same time DBMs and BDDs, and offers the benefits of a BDD-like data structure for boolean functions over real (clock) values. Experimentations reported in [47] indicate that the CDD data structure yields clear performance improvements over DBMs.

The huge memory requirements of state space exploration require careful consideration of handling of secondary memory. In [85] a technique allowing the user to control traversal of memory-blocks in an advantageous manner *without* implementing an independent memory manager is given.

A distributed version of the state exploration algorithm of UPPAAL has been implemented and experimentally tested on a Sun Enterprise 10000 with 24 333Mhz processors and on a Linux Beowulf cluster of 10 450Mhz Pentium III CPUs [48]. The distributed version has allowed previously unsuccessful verification attempts to be carried out, and demonstrates in some cases super-linear speed-up.

Applications During the period of this evaluation, UPPAAL has been applied to a number of industrial or large scale real-time case studies. For a detailed description of some of them we refer to the section on Real Time Systems. Information on UPPAAL applications conducted outside the unit may be found at UPPAAL's homepage www.uppaal.com.

Verification Methodologies In order to become truly scalable, it is imperative that UPPAAL is complemented by various methodologies allowing, in a safe manner, the users' creativity and insight to play an active role in verification. This fact has also been recognized by other researchers within the model-checking community (e.g. Ken McMillan and Rajeev Alur).

The guiding methodology from [60] offers precisely such a possibility. Here the user essentially prunes redundant and unproductive behaviour of the model by augmenting it with additional variables and guards. The method

is safe in that any behaviour of the guided model is also realizable in the full model.

In the thesis of Jensen [102] the issue of scalability is dealt with within a general framework of *abstraction*. In [71] the framework is instantiated to the modelling language of UPPAAL. Combined with precongruence results, this allows the verification of a large composite system to be dealt with in a compositional manner. The resulting methodology has been applied to obtain a complete analysis of the Bang & Olufsen Power Down protocol [57].

Extensions For purposes of efficiency, the modelling language of UPPAAL was initially designed to be rather limited in expressive power. However, a number of features have been added to it during the period of this evaluation. These extensions (location invariants, committed and urgent locations, arrays, process templates, etc.) were often driven by the needs from particular case studies. Most recently an extension of UPPAAL with stopwatches (clocks that may be stopped occasionally) has been given [53] allowing an approximate analysis of the full class of linear hybrid automata to be carried out using the efficient data structures and algorithms of UPPAAL.

Education UPPAAL has recently been used as the primary software tool in two courses. A course on *Real-Time Systems* for 3rd-4th year students at the Technical University of Denmark and a course on *Computer Engineering* for 1st year students at Aalborg University. In the latter course emphasis was put on modelling in UPPAAL and realization using LEGO MINDSTORMS.

We refer to the UPPAAL home-page www.uppaal.com for more detailed information.

The visualSTATE Tool During the period 1997–99 the unit has been engaged in the project VVS (Verification and Validation of large State-machines). VVS was a joint project between (former) Baan visualSTATE, BRICS at Aalborg University, and the Computer Systems section at the Technical University of Denmark. The project was funded jointly by Baan visualSTATE and the national Centre for IT Research (CIT). The general aim of the project was to expand the applicability of verification tools for checking the correctness of large industrial designs such as embedded controllers, remote controls, simulators, and electronic equipment. More specifically, the challenge put to us was to improve the performance of the verification engine of the already commercially successful tool visualSTATE.

The results of the project were particularly encouraging. A new technique called compositional backward reachability (CBR) was developed, allowing for verification of designs with more than 1400 parallel machines (and state spaces

with more than 10^{500} states) in less than half an hour on a standard PC. In comparison, application of traditional, state of the art BDD-technology only allows for the analysis of systems with 30–50 parallel machines. Compared with the existing verification engine of visualSTATE, the CBR method reduced verification time by several orders of magnitude (from 14 days to 6 seconds). By now the method is patented and an integrated part of the distributed version of the visualSTATE tool.

The interaction between academia and industry was a very fruitful one in the VVS project. As researchers we obtained vital access to industrial case studies in large numbers via the costumers of the company. On the other hand, the company had access to state of the art knowledge on verification techniques. Again, the research benefited from the involvement of several MSc students. The main achievements of the project are:

Compositional Backwards Reachability The first publications reporting on the CBR technique were [86, 31]. In these papers the technique is developed for the state/event model used in the visualSTATE tool. The CBR technique settles reachability properties in a compositional manner: a sequence of backwards state-space explorations is conducted involving an increasing number of parallel components. A key to the success of the method is the constant reuse of results of previous explorations. In addition state-spaces are represented symbolically using BDDs.

In [34] the technique and experimental results are presented in more general terms in an issue of the journal IEEE Computer addressing the state of the Software Engineering profession.

Hierarchical Structure In [46] we investigate techniques for verifying systems expressed in a hierarchical extension of the state/event model of visualSTATE or as results of designs given as StateCharts or in the Unified Modelling Language (UML). The straightforward way of analyzing such systems via a flattening into a non-hierarchical system suffers from the state explosion problem, even when applying the CBR technique. The aforementioned paper develops a new technique which exploits the hierarchical structure by reusing already established properties of super-states. The performance of the resulting method is very encouraging, as it proves to be not only insensitive to the hierarchical depth, but even leads to improved performance as the depth increases.

Quotienting In [124] the quotienting technique mentioned in the section on Process Calculi is developed for the state/event model used in visualSTATE, and for a variant of the modal mu-calculus suitably extended in order to allow quotient formulae to be expressed.

Minimization The work presented in [121] tackles the verification problem for state/event systems by iterative application of composition of part of the overall system followed by a minimization taking the context as well as the overall property into account. The experimental findings during this work (which preceded the CBR technique) gave clear indications of the importance of the use of compositionality: even though an explicit state-space representation was used, the context-dependent minimization technique performed in some non-trivial cases better than a traditional BDD-approach.

Compositionality and Time Given the strong process algebraic background of the group, much of the fundamental research carried out within the area of verification and validation has addressed *compositionality*, i.e., the interplay between the structure of a system and its logical properties. However, as witnessed by the CBR technique, the theoretical results on compositionality may play an important role in practical verification. Further evidence for this claim is given by the following research contributions, where compositionality is employed in the verification of real-time systems:

Constraint Oriented Methodology A constraint oriented, state based proof methodology has been put forward and applied on a number of examples [83]. In particular the applicability of the methodology to real-timed systems is demonstrated in [82].

CMC Together with François Laroussinie (Laboratoire Spécification et Vérification, CNRS URA 2236, Ecole Normale Supérieure de Cachan, Paris), a compositional verification methodology for networks of timed automata and a timed extension of the modal mu-calculus has been developed by applying the quotient technique mentioned in the section on Process Calculi. The paper [80] provides the theoretical foundation of the method together with a first prototype implementation. In [75] model checking via quotienting is experimentally shown to combat state-explosion in an automatic verification of Fischer's protocol. The method is supported by the tool CMC (Compositional Model-Checking); the status of development of this tool, together with a range of verified examples, is discussed in [76].

Real-Time Systems

The unit's research on real-time systems is concerned with methods and tools supporting the software life cycle ranging from specification to implementation and test. The results fall into three main groups: modelling and analysis of industrial and large scale case studies, specification methodology, and theoretical foundations of hybrid systems. Below, we describe the results and activities in more detail.

Industrial and large scale case studies During the period of this evaluation, the unit has applied the tool UPPAAL in the analysis of a substantial number of industrial and large scale case studies.

The first industrial case studies were based on a real-time protocol used in Philips audio equipments. In earlier work of ours, the Manchester-encoding principle of the protocol was modelled and its correctness demonstrated under assumptions on timing delays between signals. In [49] a significant variant of the protocol is examined with the purpose of studying detection of bus collisions. In [58] a similar audio control protocol used in Bang & Olufsen audio/video components is analyzed. The protocol was known to be faulty, but the error had not been found using conventional testing methods. Using UPPAAL the error was revealed, diagnosed and corrected.

As a result of this successful analysis, Bang & Olufsen and the unit decided to collaborate on the development of a power-control protocol, which was going to be used in a new product line. A model of the design was analyzed prior to its implementation [57], whereby the testing phase became very short and straightforward without revealing any design errors. The product line was released in December 1999, and the protocol is functioning without any problems.

A more recent line of case studies addresses how to check properties of actual programs, rather than their abstract models. The platform for these studies has been that of LEGO MINDSTORM RCX systems. In particular, we showed how UPPAAL models can be automatically synthesized from RCX programs, written in the programming language Not Quite C (NQC). To make the modelling faithful, the scheduling policy of the underlying operating system and policies for communication between RCX bricks are part of the model [59]. In [67] the developed translation is applied to modelling and verification of a simple box-sorting scenario using UPPAAL2K. In [56] a translation using a more sophisticated preemptive scheduling policy is modelled and analyzed using UPPAAL2K.

Another line of case studies has its roots in the European project VHS (Verification of Hybrid Systems). In [26] an analysis of an experimental batch plant is carried out. All physical components of the plant, such as the valves, pumps, tanks etc. are modelled as timed automata using integer variables, together with real-valued clocks, for modelling continuous phenomena such as variations of the level of liquid in a tank. In [60] the problem of using UPPAAL for automatic synthesis of production schedules and control programs for batch production plants has been addressed. The study is carried out using a LEGO MINDSTORM plant rather than the actual steel production plant, SIDMAR, in Ghent. To cope with the state-explosion problem, we developed a general way of adding guidance to a model. The use of this technique makes it possible to deal with a plant (concurrently) producing as many as 60 batches.

Specification methodology The unit has contributed to specification methodology of real-time systems through two collaborations on real-time actors and UML respectively. Together with colleagues at the University of Illinois, USA, a new specification formalism has been developed [95, 33]. The basic idea is to separate the functional and timing aspects of a given design. The work demonstrates how this separation of concerns may lead to reusable (component based) specifications.

The success of object-oriented analysis and design methods and, in particular, the widespread use of the UML language makes it interesting to understand how formal methods may be applied within such a framework. Based on our earlier work on refinement based specification and design, we have investigated how the structural notations of UML (case diagrams) reflect refinement or abstraction. It turns out that the refinement notion corresponds well to the UML concept of "realization". However, the structural diagram has to follow certain patterns in order to incorporate simulation relations (links) and assumptions. A first paper [97] has been published on the results of this research, but there is much further work to be done along these lines.

In addition to the above activities on specification methodology, a PhD project has investigated how the design process can support the development of software architectures in more detail [88, 90, 89, 117, 105].

Theories for hybrid systems Hybrid systems, i.e., non trivial compositions of continuous and discrete real-time systems, have been investigated since the early 90's and the fundamental models have been clarified. Currently most work is directed towards analysis and synthesis techniques that utilize model checking technology. However, there is still work to be done when integrating hybrid constructs into existing frameworks. We have been engaged in such an integration into the Action System Framework of Back and Kurki-Suonio. Together with Sere and Rönkko, we have proposed a simple hybrid action and investigated its properties [119, 120, 100].

Testing real-time systems Despite the fact that algorithms and tools for verification of real-time systems can now handle rather large system designs, there is still an urgent need to improve the state of the art with respect to automated analysis of system implementations. This is partly due to the state explosion problem, but also to the fact that the run-time platforms of such systems cannot be modelled in full detail. Work on tools for automated testing has been carried out along two lines: (1) We have developed a tool for generating timed test suites covering relevant parts of an implementation [96]. The test suites are generated on the basis of a design expressed as a timed automaton. (2) Another tool has been implemented, which can extract timing information from a running real time system [68]. Furthermore, the tool can present the observations as timed

sequence diagrams which may serve as an overview of the timing properties. Also a timed automaton model may be derived, which can be analyzed in depth by using the UPPAAL model checker.

Networks

This research activity emphasizes analysis and construction of services and protocols for computer networks. There are two main aspects of our research. The first is to formulate new or improved protocols and services and to implement them in such a way that the desired functionality and quality of service will be provided. The second aspect is to determine and precisely characterize the functionality and quality of service that is provided by protocol and service implementations, mainly through empirical experiments with modern and future network technology. Ultimately our research is driven by the desire to construct services and protocols that correspond to application needs, and applications that use the functionality and quality of services that is actually provided.

High-Speed Networks During the period 1996–1998, we collaborated with a number of Danish research institutions (Department of Computer Science at Copenhagen University (DIKU), Department of Information Technology at the Technical University of Denmark (IT-DTU), and the Informatics and Management Accounting Department at Copenhagen Business School) in a series of projects within the programme “Networks and Paradigms for the Next Generation of Distributed Systems” funded by the Danish Natural Science Research Council (SNF grant no. 9502830). The overall aim was to investigate and demonstrate the applicability of modern high speed data communication technologies for distributed systems. Our focus was on ATM technology which was perceived as the the central future carrier of high speed and integrated traffic (data, audio, video, etc.).

In [111] we analyzed the performance of our ATM network when applied to communication of compressed video and other time sensitive data. We found that ATM networks provide a substantial improvement with respect to throughput, jitter, and sensitivity to concurrent data streams compared to Ethernets. We also investigated the effects on bandwidth and CPU requirements when using MPEG-1 and MPEG-2 compression standards. MPEG-2 offers better compression, but it is too CPU demanding for software encoding of live-video using general purpose CPU's.

In [69, 115] we present a new technique for encoding video traffic based on the idea of splitting the video representation into layers of varying importance, so that the individual receivers of a one-to-many session only need to receive precisely the amount of data their network and host can handle, or they are willing to pay for. The technique is implemented by taking advantage of the dedicated Visual Instruction Set of SUN's Ultra SPARC platform. This gives a

speed improvement of a factor 3-4 compared to optimized C-code, thus enabling real-time encoding/decoding of the video.

When applying high speed networks for video communication, it is normal to accept a certain minimal error rate. However, transmission of many other media types such as high-volume data files must be completely reliable. Unfortunately, the current techniques for reliable transfer (e.g., sliding window protocols) are ill suited for high speed networks like ATM. In [110] we confirm this thesis through experiments with an existing technique (the TCP protocol in SUN OS), and by developing a new reliable protocol. It is shown that this protocol outperforms the TCP—thanks to the use of a modern retransmission technique and to the reduction of buffer copying.

In 2000 the partners (now including Department of Mathematics and Computer Science at SDU/Odense University, but excluding Copenhagen Business School) have engaged in a new three year programme entitled “Support of Network Services for Distributed Applications” also funded by the Danish Natural Science Research Council (SNF grant no. 9902071). This project will emphasize better support for Internet based applications.

IP-Technology During the last few years it has become clear that the Internet will be *the* carrier³ of more and more services offered by companies and public organizations. However, it has also become clear that it has several weaknesses and functional limitations, e.g., poor response times due to network congestion, lack of support of mobility and quality of service guarantees, major security problems etc.

The work described in [54, 125] presents a careful performance analysis of how the choice and configuration of packet queuing policies in Internet routers affects the response times of HTTP requests—the primary source of Web traffic on the Internet. *Ibidem* we empirically show that there is little difference between the random early detection (RED) queuing policy and a traditional first-in first-out (FIFO) policy on under-saturated network links. When the links becomes saturated, a carefully tuned RED policy outperforms the traditional FIFO policy, thus improving response times and reducing congestion. However, as the improvement is fairly small for HTTP traffic, it will not solve the current bandwidth problem of the Internet. This work is carried out in collaboration with the Department of Computer Science at University of North Carolina.

A Mobile Ad Hoc Network consists of a collection of autonomous roaming hosts equipped with wireless network interface cards. Contrary to cell-based networks such as the GSM-network, no existing network infrastructure like base-stations is assumed. Each host therefore also functions as a router forwarding packets from other hosts within its receiver’s range. This work aims at providing

³ATM is still widely in use, but now mainly in backbones, and often as the underlying technology for the Internet.

better support than possible in the current Internet for deployment of wireless mobile networks in environments with limited or non-existent (e.g. damaged) network infrastructure. Potential applications include emergency areas, where rescue-workers equipped with short-range, battery-operated communication devices at the same time function as part of a temporary network infrastructure, as well as military operation in hostile areas.

A main challenge here is to design new routing algorithms that are efficient and robust in this extremely dynamic self organizing environment. This line of research performed in collaboration with INRIA Rocquencourt has resulted in a proposal for a new routing algorithm, called optimized link state routing (OLSR), which has now been adopted as an Internet draft document [114].

Distributed Programming Environments While efficient basic transmission protocols are essential, these alone are insufficient for the development of sophisticated distributed systems. An important research area is therefore that of high level programming languages and environments for distributed applications. Especially important are the coordination languages which offer distributed communication and synchronization services for (arbitrary) sequential languages.

In [108, 109] we propose an object-oriented programming environment called “ActorFoundry” based on the Actor-model of distributed computation. This environment supports transparent object invocation and process migration for Java programs. A key feature of “ActorFoundry” is its modular, flexible and customizable runtime system. This allows the runtime system to be easily optimized towards the actual platform or application needs without modifying the application code. An Actor-based programming language for real-time systems is proposed in [33, 95]. This work is a result of a collaboration with the Department of Computer Science at the University of Illinois.

We have furthermore extended a distributed mutual exclusion protocol to optimally handle a large and dynamic collection of processes, and to increase its fault tolerance [116]. A prototype implementation of this protocol in Java can easily be used as a mutual exclusion service for distributed Java applications.

7.2.2 Projects

BRICS, Basic Research in Computer Science, 1994-2006 BRICS is a large scale research project funded by the Danish National Research Foundation initially for the period 1994-98 (6.5 million DKK per year). During the initial period the centre has been given additional funding for starting an international Ph.D. school at Aarhus University. More recently, the centre has been extended for a new five year period (with a total budget of 40 million DKK). The original aim of BRICS was to establish important areas of basic research in the mathematical foundations of Computer Science in Denmark, notably Algorithmics and Mathematical Logic, alongside existing activities in Semantics of Computation.

In addition to these areas, the centre currently focuses also on verification and validation, distributed systems and networks. The centre is based at the Departments of Computer Science at Aarhus and Aalborg University. The BRICS activities consist of a number of long term projects in combination with an annual theme as well as a series of workshop and conferences (e.g. ICALP in July 1998). The activities are centred around the permanent staff at the two universities, and an average of 15-20, mainly foreign, long term junior and senior associated researchers, 20 Ph.D. students and an intensive programme of visiting researchers and scientific events. The activities in Aalborg form approximately 25% of the total project. Among ongoing projects are: Semantic foundations for concurrency, semantic models for real-time and hybrid systems, automated tools for verification and test of embedded and real-time systems, semantic models for objects and mobile processes, distributed and networking systems. The Aalborg branch of BRICS is directed by Professor Kim Guldstrand Larsen. The total AAU grant in the period is 10 million DKK. More information on BRICS may be found at <http://www.brics.dk/>.

VVS, Verification and Validation of large State-machines. CIT-project no 99, 1997–1999 VVS was a joint project between Baan VisualSTATE, BRICS at Aalborg University, and the Computer Systems section, Technical University of Denmark, supported by the Centre for Information Technology. The project aimed at developing verification tools for checking the correctness of large industrial designs such as embedded controllers, remote controls, simulators, and electronic equipment. The first result of the co-operation was the verification of a design with more than 1400 concurrent state machines in less than an hour on a standard PC. This constituted a breakthrough in the art of automatic verification. The technique used in this verification is patented and is incorporated in new generations of the visualSTATE tool. Total AAU grant in the period 800,000 DKK. (Kim G. Larsen, Kåre J. Kristoffersen, Gerd Behrmann, Arne Skou, Carsten Weise and Klaus Havelund.)

VHS, Esprit Project 26270, Verification of Hybrid Systems, 1998-2001 Hybrid systems comprise systems which depend critically on the interaction between the discrete dynamics of a digital controller and the continuous dynamics of the environment in which it is embedded. The project aims at developing a framework for hybrid systems allowing for modelling, simulation, verification, synthesis and implementation of such systems. The main activity of the project is centred around a number of academic and industrial case studies, mainly from chemical industry, in order to define formal models of plants. These models are used to verify properties concerning their behaviour. The project uses, among others, the models of timed and hybrid automata, invented by computer scientists in order to express hybrid phenomena. Several tools for analyzing systems

expressed in this formalism have been built, and the project uses them for automatic verification. Due to the hardness of the problem, some less ambitious validation methods such as simulation are also investigated. Total AAU grant in the period 600,000 DKK. (BRICS, Aalborg, Denmark; Verimag, France; Weizmann, Israel; KUN, Nijmegen, Netherland; Kiel, Germany; Dortmund, Germany; LAG, Grenoble, France; CWI, Amsterdam, Netherland; Gent, Belgium)

Fireworks, Feature Interaction in Requirements Engineering, 1998-2000 The project investigates the development and application of formal methods for analyzing the interaction between existing and new services in computer systems like modern telecommunication services. Experience has shown that the requirements definition methodology has a major impact on whether or not new services (like e.g. call propagation) can be integrated without implying inconsistencies with existing services.

The project is supported by the EEC as a working group (travelling expenses). Total AAU grant in the period 50,000 DKK. (Kim G. Larsen, Arne Skou)

Networks and Paradigms for the Next Generation of Distributed Systems, 1996-2002 The introduction of networks like ATM, GSM, UMTS and Gigabit introduces new requirements with respect to flexibility and performance of the basic network software. In general, these requirements are not fulfilled, and the research project makes a closer investigation of the state of the art. Also new paradigms and protocols (meeting the requirements in a better way) are designed, implemented and evaluated through experiments. The project is a collaboration project with colleagues at DTU and Copenhagen University, and it is supported by the National Science Foundation in 1996-1998 and 2000-2002. Total AAU grant in the period 380,000 DKK. (Brian Nielsen, Wladyslaw Pietraszek, Mikkel Christiansen, Arne Skou)

Calculi and verification techniques for security and mobility, 1999-2001 The family of pi-calculi provides a rich mathematical theory for studying the properties of systems whose communication properties change dynamically. An important spin-off of this theory is the Mobility Workbench, an automatic software tool for verifying properties of pi-calculus processes. Recently, Abadi, Gordon and others have used the closely related *spi-calculus* to describe the behaviour of security protocols.

The project has as a main goal that of developing the theory of the spi-calculus and to use it in the construction of a software tool for verifying properties of security protocols, and subsequently integrating this tool with the Mobility Workbench, if possible.

The project is a collaborative effort involving colleagues at Uppsala University and is funded by Göran Gustafssons Stiftelse (The Göran Gustafsson Foundation,

a Swedish foundation for basic research in the natural sciences). (Hans Hüttel, Josva Kleist, Uwe Nestmann with Björn Victor of Uppsala University)

API. Autonomous Plant Inspection System 2000-2003 API is a joint project between the Danish Institute of Agricultural Sciences, The Department of Agritechnology of The Royal Veterinary and Agricultural University, and the departments of Control Engineering and Computer Science at Aalborg University. The Project is supported by the Danish Ministry of Food, Agriculture and Fisheries and a national industry consortium.

The goal of the project is to deploy an operating autonomous scout vehicle which can map weed densities in a field with a specific crop. The prototype shall be extensible to later take over field operations, e.g. weeding. Furthermore the strategy software shall be prepared for operating collections of such vehicles.

The Department of Computer Science (Anders P. Ravn, Peter K. Jensen) is responsible for one of 4 major work packages: Systems Architecture. The purpose is to define and document the overall system structure and analyze its requirements with respect to real-time processing. Furthermore, we take active part in the final system's software integration.

The project starts November 2000 and runs till April 2003. Total funding for Computer Science at Aalborg University 1,300,000 DKK.

Centre for Agricultural Technology - Aalborg University This centre is a network of groups within the Engineering and Science Departments which work with problems related to Agriculture. The Centre was formed by the Faculty of Engineering and Sciences in September 1999 and is funded by an annual grant of 100 000 DKK. The Faculty has appointed Anders P. Ravn as director of the Centre.

In the period 1999–2000, the Centre has received an additional local government grant of 600,000 DKK. This additional grant has the purpose of further developing the Centre, its relations to local industries, and curricula for engineering students so that the students have a fundamental understanding of the agricultural application area.

EXPRESS, 1996–1997 EXPRESS was a network funded by the European Human Capital and Mobility initiative. The DSS unit participated in that network with Luca Aceto as representative.

One of the main research objectives of computer science is the development of formal methods for the design and implementation of programming languages. A most prominent feature of this research area is the proliferation of programming concepts which have been developed and studied: we mention programming paradigms such as imperative programming, logic programming, functional programming, concurrent programming, object-oriented programming, and data-

flow, and the various combinations thereof. The formalization of the relation between a specification and its implementation has been a main focus of systematic research. However the issue of the relative expressive power of the various programming concepts, an issue which is directly related to the implementation and use of programming languages, has hardly been addressed in a systematic manner.

Such methods provide a tool to classify the variety of programming languages, and a formal basis of the design principles and implementation of programming languages. The research programme of EXPRESS aimed at providing a general framework for the comparison of the formal methods for specification and verification developed within the various programming paradigms.

The EXPRESS proposal brought together fifteen prominent teams in the area of expressiveness. Besides the coordinator, the consortium consisted of six proposers, four associated proposers, and four external proposers from non-Member States: Centre for Mathematics and Computer Science (CWI), The Netherlands (coordinator); Eindhoven University of Technology (TUE), The Netherlands (proposer); Swedish Institute of Computer Science (SICS), Sweden (proposer); University of Genova, Italy; University of Rome “La Sapienza”, Italy; University of Hildesheim, Germany (proposer); Gesellschaft für Mathematik und Datenverarbeitung (GMD), Bonn, Germany; University of Amsterdam, The Netherlands; INRIA Sophia-Antipolis, France (associated proposer); INRIA Rennes, France (associated proposer); University of Sussex, UK (associated proposer); Weizmann Institute, Israel (external); State University of New York at Stony Brook, USA (external); Cornell University, USA (external). Total funding for AAU 34,000 ECUs.

Mindpass Centre for Distributed Systems 2000-2005 Mindpass Centre for Distributed Systems (Mindpass-CDS) is established as a research and development collaboration between Mindpass A/S and researchers from the DSS unit. Mindpass-CDS is created with the following main objectives: to develop a tradition for research and development activities within distributed systems that is technologically innovative, theoretically well-founded, and based on collaboration and exchanges between industry and academia, and to establish and develop research and development activities within distributed computing systems at Mindpass A/S as an integral part of their business strategy and practice.

During the first year Mindpass-CDS will establish MindsLab: an experimental laboratory facility for conducting large-scale, realistic experiments and simulations, supporting the research and educational activities related to Mindpass-CDS. A first use of the laboratory will be for experiments with the adaptation of existing distributed applications to SCI based clusters. Another topic under investigation is modeling and reasoning on dynamic network properties. The intention is to model, reason about, and further develop mobile ad-hoc routing

protocols, such as OLSR, through application of formalisms and verification tools, such as timed automata and Uppaal.

Mindpass-CDS will have a minimal organization and budget during its first exploratory year. Based on the insights and results gained during the first year, the centre will subsequently consolidate and expand its activities. The ambition is to continue the collaboration for a five year period. The budget for the first year is of 600,000 DKK, jointly financed by Mindpass A/S and the faculty of Faculty of Technical and Natural Sciences, AAU. (Gerd Behrmann, Thomas Heide Clausen, Josva Kleist.)

7.3 Organization and Staff

7.3.1 Staff

The tables below give numerical information about our staffing (both teaching and pure research) in the period of this evaluation. We show the number of staff members within the unit for each year, with more detailed half-year information provided by left- or right-shifting of numbers.

Scientific Staff					
	1996	1997	1998	1999	2000
Full Professors	1	1	1	1	1
Associate Professors	3 4	4	3	3	4
Assistant Professors	2	2 4	3 4	4	5
Teaching Assistants	1		1	3	1
Visiting Professors	1		1	1	1
Visiting Researchers	1	2 3	2 3	4 3	2
Ph.D. Students	6 7	7 6	6 4	5 6	5

Luca Aceto

Academic Degrees

1991 DPhil in Computer Science, University of Sussex
 1986 Laurea (M.Sc.) in Computer Science, University of Pisa

Positions

- 1996- Associate Professor, Department of Computer Science, Aalborg University
- 1994–96 BRICS Senior Research Fellow, Department of Mathematics and Computer Science, Aalborg University
- 1995 Visiting Research Professor (half a year), Department of Mathematics and Computer Science, Aalborg University
- 1993 Visiting Researcher (half a year), Department of Mathematics and Computer Science, Aalborg University
- 1992–96 Lecturer in Computer Science, School of Cognitive and Computing Sciences, University of Sussex
- 1991–92 Researcher, Hewlett–Packard Lab., Pisa Science Center
- 1991 Professeur Invité (half a year), Centre de Math., INRIA Sophia Antipolis
- 1987–90 Research Fellow, Computer Science and Artificial Intelligence Subject Group, University of Sussex

Administrative duties

- 1996–97 Steering committee for the EEC funded HCM project EXPRESS
- 1997– Chairman of the public relations and student recruiting working group, Department of Computer Science, Aalborg University
- 1999– Member of the Travel Funding Committee, Department of Computer Science, Aalborg University
- 1999– Socrates coordinator, Department of Computer Science, Aalborg University

Key Publications: [12, 14, 16, 18]

Hans Hüttel

Academic Degrees

- 1991 Ph.D. (Theory of Computation), University of Edinburgh, Scotland
- 1988 Cand. scient. (M.Sc.) in Computer Science and Mathematics, Aalborg University

Positions

- 1995- Associate Professor in Computer Science, Aalborg University, Institute for Electronic Systems
- 1991-95 Assistant Professor in Computer Science, Aalborg University, Institute for Electronic Systems

Key Publications: [93, 61, 24, 55]

Anna Ingólfssdóttir

Academic Degrees

- 1994 DPhil in Computer Science, University of Sussex
- 1987 Cand. scient. (M.Sc.) in Mathematics/Computer Science, Aalborg University
- 1976 B.Sc. in Mathematics/Physics, The University of Iceland

Positions

- 2000– Associate Professor, Department of Computer Science, Aalborg University
- 1998–99 Research Assistant, Department of Computer Science, Aalborg University
- 1998 Visiting Researcher at LSV, ENS de Cachan, Paris, France and at The Chinese Academy of Sciences
- 1997–98 Researcher, The University of Florence, Italy
- 1995–97 Research Assistant, Department of Computer Science, Aalborg University
- 1991–95 Assistant Professor, Department of Computer Science, Aalborg University
- 1990 Research Fellow, The University of Sussex

Awards

In 1997 she received a two-year postdoctoral grant from the Danish Research Council.

Key Publications: [12, 14, 16, 18]

Henrik Ejersbo Jensen

Academic Degrees

- 1999 Ph.D. in Computer Science, Aalborg University
- 1995 Cand. Polyt. (Master of Science in Engineering, Msc. EE.), Aalborg University

Positions

- 1999– Assistant Professor, Department of Computer Science, Aalborg University
- 1995–99 Ph.D. Student, Department of Computer Science, Aalborg University
- 1996–98 Visiting Research Scholar, Laboratory for Computer Science, Massachusetts Institute of Technology (M.I.T.), Cambridge, U.S.A.
- 1996 Teaching Assistant, Department of Computer Science, Aalborg University

Key Publications: [71, 102, 72, 70]

Ole Høgh Jensen

Academic Degrees

1993 Cand. Polyt. (Master of Science in Engineering, Msc. Eng.), Aalborg University

Positions

2000– Amanuensis, Department of Computer Science, Aalborg University
1997–00 Assistant Professor, Department of Computer Science, Aalborg University
1993–97 Ph.D. Student, University of Edinburgh and Cambridge University, Great Britain
1989–92 Teaching Assistant, Department of Computer Science, Aalborg University

Peter Krogsgaard Jensen

Academic Degrees

1993 Master of Science in Electrical Engineering, Aalborg University

Positions

1999– Assistant Professor in Computer Science, Aalborg University, Department of of Computer Science
1995–99 Ph.D. Scholarship, Aalborg University, Department of of Computer Science
1993–95 Software developer at Terma Elektronik A/S

Josva Kleist

Academic Degrees

2000 Ph.D. in Computer Science, Aalborg University
1994 Cand. Scient. (M.Sc.), Aalborg University

Positions

2000–01 Post-doc INRIA Sophia-Antipolis (sabbatical), France
1998– Assistant Professor, Department of Computer Science, Aalborg University
1995–98 Ph.D. student in the semantics group at Department of Computer Science, Aalborg University
1994–95 System developer at Modern CAD/CAM Systems in Randers, Denmark
1988–89 Technician/programmer at the engineering company Henrik Steffensen in Randers, Denmark

Key Publications: [24, 73, 74, 93]

Kim Guldstrand Larsen

Academic Degrees

- 1986 Ph.D. in Computer Science, Edinburgh University, Scotland
- 1982 Cand. scient. (M.Sc.) in Mathematics, Aalborg University

Positions

- 2000–02 Industrial Professor (part-time), University of Twente, Holland
- 1995 Visiting Professor, Uppsala University, Sweden
- 1993– Full Professor in Computer Science, Aalborg University, Institute for Electronic Systems
- 1992–93 Associate Professor in Computer Science, Aalborg University, Institute for Electronic Systems
- 1989–92 Senior Researcher (Seniorstipendiat) in Computer Science, Aalborg University, Institute for Electronic Systems
- 1986–89 Associate Professor in Computer Science, Aalborg University, Institute for Electronic Systems
- 1985–86 Assistant Professor in Computer Science, Aalborg University, Institute for Electronic Systems

Administrative duties

- 1999– Member of the Steering Committee for European Educational Forum
- 1999– Member of the Steering Committee for the Danish Research Net.
- 1999– Member of the Board of Directors for Institute of Computer Science, Aalborg.
- 1998– Director of Aalborg section of BRICS
- 1996– Chairman for 'Bonusudvalget' (Danish Natural Science Research Council)
- 1996– Administrator of the Computer Science and Engineering Ph.D. Programme, Aalborg University
- 1991–97 Member of Danish Natural Science Research Council.
- 1987–89 Co-chairman of the Department of Mathematics and Computer Science, Aalborg University

Awards

In 1996 and 1999 he received awards for outstanding research contributions from the Faculty of Technical and Natural Sciences, Aalborg University. In 1999, he became Honorary Doctor (Honoris causa) at Uppsala University. In 2000, he became member of Royal Danish Academy of Sciences and Letters (Det Kongelige Danske Videnskabernes Selskab), Copenhagen.

Key Publications: [22, 27, 29, 34]

Brian Nielsen

Academic Degrees

- 2000 Ph.D. in Computer Science, Aalborg University
- 1993 Cand. Polyt. (Master of Science in Engineering, MSc. EE.), Aalborg University

Positions

- 2000- Teaching Assistant at the Department of Computer Science, Aalborg University
- 1997–00 Assistant Professor at the Department of Computer Science, Aalborg University
- 1994-97 Ph.D. student at the Department of Computer Science, Aalborg University
- 1995-96 Visiting scholar at the University of Illinois at Urbana-Champaign, Illinois, U.S.A.
- 1993-94 Teaching Assistant at the Department of Computer Science, Aalborg University

Anders Peter Ravn

Academic Degrees

- 1995 dr. techn. Technical University of Denmark
- 1973 cand. scient. in Computer Sc. and Mathematics, Univ. of Copenhagen

Positions

- 1999- Res. Prof., Dept. Comp. Sc., Aalborg University
- 1998–99 Visit. Prof., Dept. Comp. Sc., Aalborg University
- 1989–99 Reader IT-DTU, Dept. Inf. Technology
- (1994) Visit. Prof., Informatik, Univ. of Kiel, Germany
- 1984-89 Lecturer, ID-DTH, Dept. Comp. Sc., Techn. Univ. of Denmark
- 1985-89 Adjunct Lect., Dept. Math., Royal Vet. and Agric. Univ.
- 1982-83 Visit. Sc., IBM T. J. Watson Res. Cen., Yorkt. Hgths., USA
- 1980-84 Associate Prof., Dept. Comp. Sc. Univ. of Copenhagen
- 1976-80 Assistant Prof., Dept. Comp. Sc. Univ. of Copenhagen
- 1972-76 Systems programmer, A/S Regnecentralen

Key Publications: [23, 97, 98]

Arne Skou

Academic Degrees

- 1990 Ph.D. in Computer Science, Aalborg University
- 1975 Cand. scient. (M.Sc.) in Computer Science, Copenhagen University

Positions

- 1989- Associate Professor in Computer Science, Aalborg University, Institute for Electronic Systems
- 1986-89 Ph.D. Scholarship, Århus University, Computer Science Department. Visiting University of Sussex, UK, for 7 months during the Ph.D. study.
- 1986 Research assistant, Copenhagen University.
- 1975-86 Member of the senior scientific staff at the computing centre of Aalborg University. Head of the computer networking group.
- 1971-75 Teaching assistant, Department of Computer Science, Copenhagen University.

Administrative duties

- 1998–00 Course planning for unemployed engineers
- 1996–00 Board member, Institute of Electronic Systems
- 1999–00 Head of computer resources committee
- 1996–98 Head of Department

Key Publications: [57, 58, 71, 96]

7.3.2 Research Organization

During the second half of the evaluation period, the unit has organized research seminars on a regular basis, either each week or every second week depending on the precise flow of visitors. The full list of seminars (including titles and abstracts) since autumn 1998 is available from the web-page of the unit (see <http://www.cs.auc.dk/research/FS>). In addition a series of more focused, informal seminar meetings has been held on the topics of Calculi for Mobility and Objects and Reactive, Real-time and Hybrid Systems.

Each year the unit has been on a two day retreat aiming at discussing both scientific and organizational matters. These retreats have been particularly useful in the transition from two separate groups to one unit. However, from this research evaluation perspective, our unit still looks like a loose association of researchers with individual interests and goals. A look at our weighted collaboration graph shows connected components that roughly correspond to some of the research activities mentioned in Sect. 7.2. It is evident from this graph that there is already interaction between some of the different research activities. However, perhaps quite naturally, the weight of the connected components is far larger than that of their connecting links to other clusters of activity within our unit.

Moreover, the weight of the links from the connected components to the external research world is, in general, substantially larger than that of those denoting inter-cluster collaborations. This seems to indicate a natural tendency in our unit to seek research collaboration with kindred spirits in the academic world, rather than necessarily within the unit itself. As elaborated in Sect. 7.8.2 to follow, we have delimited areas for further joint research. These have been chosen in such a way that synergy between different research activities is beneficial for each of the areas. Still, we also believe that “*science is an essentially anarchic enterprise*” (Paul Feyerabend), and that individual researchers should be free to follow their own specific research goals.

7.4 Collaborations

The unit has continued its strong emphasis on collaboration throughout the evaluation period both on a national as well as international level. In particular, the participation in BRICS supplemented with the activities in VVS and ATM, as well as the European project VHS has allowed us to maintain a continuous employment of project researchers, yielding an important contribution to the research environment and activities of the unit. Collaboration within the department and the university has been more sporadic and should be strengthened in the next five year period as we will indicate in Sect. 7.8.

Project Researchers

- *Bertrand Jeannot* October 2000–October 2001. Funded by INRIA (Verification and Validation activity).
- *Uwe Nestmann*, October 1997–September 2000. Funded by BRICS (Calculi for Mobility and Objects activity).
- *Paul Pettersson*, February 1999–June 2000. Funded by BRICS (Verification and Validation activity).
- *Kåre Jelling Kristoffersen*, August 1998–June 2000. Funded by BRICS (Verification and Validation activity).
- *Carsten Weise*, September 1997–March 1999. Funded by BRICS (Verification and Validation activity).
- *Klaus Havelund*, 1996–1997. Funded by BRICS (Verification and Validation activity).

Short Term Visitors

Throughout the evaluation period, and with increasing frequency, a number of short term guests (visits longer than one week) have visited the group. In several cases the visits have resulted in subsequent publications, as witnessed by the publication list of the unit.

- *Augusto Burgueno*, CERT (Centre d'Etudes et de Recherches de Toulouse), France, June–August 1997.
- *Davide Sangiorgi*, INRIA Sophia-Antipolis, November 1997.
- *Massimo Merro*, INRIA Sophia-Antipolis, October–December 1998, June–July 1999, January 2000.
- *Brahim Mammass*, Universite Pierre and Marie Curie, Paris, France, June–August 1998.
- *Huimin Lin*, Institute of Software, Chinese Academy of Sciences, Beijing, China, November 1998.
- *Francois Laroussinie*, Laboratoire Spécification & Vérification, Ecole Normale Supérieure de Cachan, France, November 1998.
- *Matthew Hennessy*, School of Cognitive and Computing Sciences, University of Sussex, UK November 1998.
- *Justin Pearson*, Department of Computer Systems (DoCS) Uppsala University, Sweden, November 1998.
- *Joost-Pieter Katoen*, Faculty of Computer Science, University of Twente, Netherlands, December 1998.
- *Jan Peleska*, Bremen, Germany, March 1999.
- *Ansgar Fehnker*, Computing Science Institute, Nijmegen University, The Netherlands, March 1999.
- *Fernando Cuartero & Valentin Valero*, Department of Computer Science, UCML, Albacete, Spain, April 1999.
- *Martin Steffen*, Institute of Computer Science and Applied Mathematics, Christian-Albrechts-Universität of Kiel, Germany, April 1999.
- *Paritosh Pandya*, TATA, India, May 1999.
- *Mauno Rönkko*, Turku, Finland, June 1999.

- *Alexandre David*, Department of Computer Systems (DoCS) Uppsala University, Sweden, August 1999.
- *Ernst-Rüdiger Olderog*, Oldenburg, Germany, October–December 1999.
- *Franck Cassez*, Ecole Centrale de Nantes, France, October–December 1999.
- *Björn Victor*, Department of Computer Systems (DoCS), Uppsala University, Sweden, November 1999.
- *Anna Philippou*, Department of Computer Science, University of Cyprus, January 2000.
- *António Ravara*, Technical University of Lisbon, April 2000.
- *Wan Fokkink*, CWI, Amsterdam, The Netherlands, April–May 2000.
- *Judi Romjin*, Computing Science Department, University of Nijmegen, The Netherlands, April–May 2000.
- *Christine Röckl*, Technical University of Munich, Germany, August 2000.
- *Zoltan Ésik*, Department of Computer Science, University of Szeged, Hungary, September 2000.

Sabbaticals

Carefully planned sabbaticals have played an important role in maintaining a healthy research group with consolidated and strengthened international contacts.

- *Josva Kleist*, INRIA Sophia Antipolis, September 2000–August 2001.
- *Brian Nielsen*, University of Bremen, November 1999–January 2000.
- *Luca Aceto*, Dipartimento di Sistemi ed Informatica, Università di Firenze, Firenze, Italy. February–August 1998.
- *Kim G. Larsen*, Computer Science Department, University of Twente, The Netherlands, August–December 2000.

Ph.D. Internationalization

Nearly all Ph.D. students have visited other (international) research groups for substantial amounts of time during their Ph.D. period. These stays are based on and strengthen the international network of the unit.

- Gerd Behrmann: Nijmegen University, Nijmegen, The Netherlands, November 1–7, 1999. Twente University, Enschede, The Netherlands, October 7–13, 2000.
- Mikkel Christiansen: University of North Carolina at Chapel Hill, July–November 1999 (plus shorter stays of up to one month since 1999).
- Thomas Heide Clausen: INRIA Rocquencourt, July 1999–February 2002.
- Thomas Hune: Nijmegen University, Nijmegen, The Netherlands, August–December, 1999.
- Henrik E. Jensen: Massachusetts Institute of Technology (M.I.T.), Cambridge, MA, U.S.A., July 1996–July 1998.
- Peter K. Jensen: Carnegie Mellon University, Pittsburgh, U.S.A., August 1996–February 1997.
- Josva Kleist: INRIA, Sophia Antipolis, France, October 1996–January 1997.
- Kaare J. Kristoffersen: Uppsala University, Sweden, February–August 1996.
- Brian Nielsen: University of Illinois at Urbana-Champaign, Illinois, U.S.A., February 1995–January 1996.
- Oliver Möller: University of Pennsylvania, Philadelphia, U.S.A., August–November 1999, and Uppsala University, Sweden, August–December 2000.

Collaboration with industry

As described in details elsewhere, the unit has collaborated with the company BAAN visualSTATE A/S within the CIT project VVS (Verification and Validation of large Systems). However, collaboration with industry has taken place also outside funded projects.

Collaboration with the Danish video/audio company Bang&Olufsen has taken place in the period 1996-97 on the validation and development of real-time protocols. This work was carried out by Klaus Havelund, Arne Skou and Kim G. Larsen.

A five lecture course on UML as a design methodology for embedded systems was given by Anders P. Ravn and Arne Skou for Bosch Telecom in 1999.

Brian Nielsen and Peter K. Jensen gave a series of lectures on real-time operating systems and real-time scheduling for Bang&Olufsen.

During the period of the evaluation, the unit has contributed on four occasions with lectures as part of the Life-long Education Programme at Aalborg University, covering results on Networks, Real-time Systems and Verification and Validation. The participation at the lectures has ranged from 12 to 26 computer scientists and engineers from industry.

7.5 Ph.D. Projects

7.5.1 Overview

During the period 1996–2000 thirteen Ph.D. students have been associated with the DSS unit. Ten of these students have been enrolled with the Ph.D. programme at Aalborg University, and three have been enrolled elsewhere but with strong relations to Aalborg. Four of the thirteen have successfully completed their Ph.D. study during the period. The remaining students are either about to graduate or have they have recently started (1999) their studies.

7.5.2 Project Descriptions

Reasoning about Objects using Process Calculus Techniques

Name: Josva Kleist

Education: Cand. Scient in Computer Science

Duration: 1/8 1995–31/7 1998

Status: Degree awarded 2000

Financial Support: BRICS

Advisor: Kim Guldstrand Larsen and Hans Hüttel

This thesis investigates the applicability of techniques known from the world of process calculi to reason about properties of object-oriented programs.

The investigation is performed upon a small object-oriented language - The object calculus of Abadi and Cardelli. The investigation is twofold:

We investigate translations of object calculi into process calculi, with the idea that one should be able to show properties of object calculus program by showing properties about their translation. We present translations of two object calculi into pi-calculi. A translation of the untyped functional object calculus turns out to be insufficient. Based on our experiences, we present a translation of a typed imperative object calculus, which looks promising. We are able to provide simple proofs of the equivalence of different object calculus objects using this translation. We use a labelled transition system adapted to the object calculus to investigate

the use of process calculi techniques directly on the object calculus. The results obtained are of a fairly theoretical nature. We investigate the connection between the operational and denotational semantics for a typed functional object calculus. The result is that Abadi and Cardelli's denotational model is sound but not complete with respect to the operational semantics. We also construct a modal logic for the typed functional object calculus, provide a translation of types to a sub-logic and prove the translation is sound and complete.

The amount work required to perform these investigations indicate, that although it is perfectly possible to use process calculus techniques on object oriented languages, such techniques will not come to widespread use, but only be limited to reasoning about critical parts of a language or program design.

Compositional Modal Logics

Name: Jørgen H. Andersen

Education: Cand. Scient in Computer Science

Duration: 1/8 1994–1997

Status: Degree awarded 1997

Financial Support: BRICS and Aalborg University.

Advisor: Kim Guldstrand Larsen

My thesis work include a study of general extensions of the classical pure calculi and modal logics to take into account the variety of parameters of extensions such as value-passing, real-time and hybrid systems. The focus will be how to extend existing and/or provide new algorithms for automatic verification in the more general settings. This is, however, quite an ambitious goal which in most cases will be impossible to reach as problems become undecidable with further expressiveness of the modelling languages. When this becomes the case we are confronted with the choices of:

- either give up expressive power of the formalisms to obtain decidability, or
- design verification strategies which can be used for semi-automatic verification.

Tools like EPSILON and TAV are products of the first choice. A tool like ALF is product of the second choice. These are in a sense extremes as the first often are too restrictive and the latter is too general and therefore requires too much interaction. As my focus is on generality of specific kinds of calculi and logics my aim is somewhere in between the two. In parallel with the evolution of theory I will provide prototype tools implementing the theory.

Compositional Verification of Concurrent Systems

Name: Kaare Kristoffersen

Education: Cand. Scient in Computer Science

Duration: 1/8 1995–1997

Status: Degree awarded 1998

Financial Support: The Danish Technical Science Research Council

Advisor: Kim Guldstrand Larsen

This thesis presents a collection of techniques and tools for avoiding the state-explosion problem in verification of state/event systems and real-time systems.

Abstraction-Based Verification of Distributed Systems

Name: Henrik Ejersbo Jensen

Education: Cand. Polyt. (Master of Science in Engineering, Msc. EE.)

Duration: 1/8 1995–1/6 1999

Status: Degree awarded October, 1999

Financial Support: Aalborg University

Advisor: Kim Guldstrand Larsen

This thesis presents abstraction-based proof methods and practical abstraction strategies to support the integration of theorem proving and model checking methods in verification of distributed systems. The thesis presents abstraction frameworks for untimed systems described as I/O automata and for real-time systems described as timed automata. The untimed framework is supported by the Larch Theorem Prover and the SPIN model checker. The timed framework is supported by the UPPAAL model checker.

Abstraction-based proofs are provided for three nontrivial distributed algorithms all parameterized in the number of processes: Burns' Mutual Exclusion algorithm, The Bounded Concurrent Timestamp System (BCTSS) algorithm, and Fischer's Real-Time Mutual Exclusion algorithm.

Bisimulation Congruence in Graph-based Action Calculi

Name: Ole Høgh Jensen

Duration: 1994–

Status: Unfinished

Advisor: Robin Milner

This thesis develops a graph-theoretical description of Milner's Action calculi, a common framework in which to express the dynamics of a range of computational

calculi such as process calculi and lambda calculus. The action graphs serve as a model for the structural congruence that is usually imposed on action terms in the style of the pi-calculus.

A description of several computational calculi and their operational semantics is given both as action terms and action graphs.

The main result shows how to obtain, in a uniform and general way, a natural behavioural congruence for a restricted class of action graphs. That is, starting from any action calculus (of the restricted type) including a number of reduction rules, it is shown how to obtain a suitable set of labels for a labelled semantics in such a way that the set is just large enough to guarantee that the bisimulation equivalence induced by the labelled semantics is a congruence.

Reliable Real-Time Applications. And how to use tests to model and understand

Name: Peter Krogsgaard Jensen

Education: Master of Science in Electrical Engineering

Duration: 1/9 1995–1/3 1999

Status: In progress

Financial Support: Danish Technical Research Council (STVF)

Advisor: Arne Skou

The task of verifying timing properties of critical computing systems during development is important, due to the serious or fatal consequences if such systems fail. The importance is growing because more and more software is involved in controlling the everyday life of humans. As we need to develop and verify more and more complex and larger systems, the software engineering process must be improved to maintain the quality of the end product. One way to improve is by making it more automatic, such that trivial and error prone manual task can be handled by tools.

The development of software for industrial use will involve a testing phase, with the purpose of “verifying” the requirements of the product. Traditionally tests have targeted functional requirements, but the thesis in this work is that tests can also be the foundation of analysis and verification of timing requirements. To support our thesis, we develop tools to analyze timing information from (instrumented) systems under test. An industrial sized control application is used as a case study, and all tools and techniques that we have been developing are used in the case study. The case study works as a proof-of-concept, displaying the pros and cons of the tools and techniques.

Specification and Test of Real-time Systems

Name: Brian Nielsen

Education: Cand. Polyt. (Master of Science in Engineering MSc. EE.)
Duration: July 1994–November 2000
Status: Degree awarded 2000
Financial Support: Danish Technical Research Council
Advisor: Arne Skou

Distributed real-time computer based systems are very complex and intrinsically difficult to specify and implement correctly; in part this is caused by the overwhelming number of possible interactions between system components, but especially by a lack of adequate methods and tools to deal with this complexity.

This thesis proposes a real-time specification and modeling language which facility modular specification and programming of reusable real-time software components. Further, it proposes new techniques for systematic and automated black box conformance testing of real-time systems against densely timed specifications given as a subclass of timed automata. This work emphasizes a thorough treatment of the time dimension of the test generation problem, as well as efficient tool support therefore. Experiments with a prototype tool shows that our technique produces encouragingly small test suites.

The Performance of HTTP Traffic Under Random Early Detection Queue Management

Name: Mikkel Christiansen
Education: Cand.polyt (M.E.) in Computer Science
Duration: August, 1997–February 2001
Status: In progress
Financial Support: Danish Technical Research Council
Advisor: Arne Skou

The active queuing mechanism "Random Early Detection" (RED) is being recommended by the IETF for widespread deployment in internet routers. The idea of RED and other active queuing mechanisms is to replace the traditional FIFO buffer with FCFS scheduling in internet routers with a more complex mechanism called an active queuing mechanism. The active queuing mechanism is usually designed for solving a specific type of problems occurring on the internet, such as avoiding buffer overflows in the router and/or ensuring that no single flow can use more bandwidth than its fair share.

The active queuing mechanism RED is designed for detecting incipient congestion on a router and reacting to this, so that the congestion can be avoided.

None of the empirical evidence, that we are aware of, documents the effects of RED on HTTP response time performance. That is the performance that users

experience when surfing the Web. We find this to be missing research, due to the fact that recent studies of Internet traffic show that 60-80% of the traffic on the internet is HTTP traffic. Therefore, this work has addressed by the question: How is HTTP performance effected by RED and can RED be tuned to optimize it?

This has resulted in an empirical evaluation of the effects of RED on HTTP response time performance. The overall conclusion of the evaluation shows that RED provides no significant advantage over FIFO in a network dominated by HTTP traffic.

This work is in part supervised by Professor Kevin Jeffay and Research Professor F. Donelson Smith, University of North Carolina at Chapel Hill, USA.

Validation and Verification of Embedded Systems

Name: Gerd Behrmann

Education: Cand. Scient. (Master of Science in Computer Science, Msc. CS.)

Duration: 1/2 1999–1/2 2002

Status: In progress

Financial Support: BRICS

Advisor: Kim Guldstrand Larsen

This thesis considers model checking of embedded systems, and especially real-time systems described as timed automata. It focuses on extending current models and providing techniques and data-structures to efficiently model check instances of current and new models. An efficient algorithm for reachability analysis in hierarchical systems; data-structures such as *clock difference diagrams* and *priced zones*; techniques for reachability analysis of stopwatch automata and linearly priced automata; and a distributed algorithm for timed automata are presented.

Analyzing Real-Time Systems using Timed Automata

Name: Thomas Hune

Education: Cand. Scient.

Duration: 1/2 1997–31/1 2001

Status: Ph.D. student

Financial Support: University of Aarhus

Advisor: Mogens Nielsen & Kim G. Larsen

This thesis projects is concerned with analyzing real-time systems based on the formal model of timed automata. Extensions of timed automata has been studied and the algorithm used in Uppaal for analyzing the behavior of timed automata

has been extended to handle a new formalisms. Especially, a notion of cost has been introduced and automatic synthesis of parameters has been made possible. Also a distributed version of the original tool has been implemented and tested on two different kinds of machines. Uppaal has also been applied to an industrial case-study based on the SIDMAR steel production plant in Gent, Belgium.

The decidability of timed automata with cost has been proven and using the categorical framework of open maps a new proof of decidability of timed bisimulation has been presented.

Hierarchical Specification and Verification of Timed Systems

Name: Michael Oliver Möller

Education: Diplom-Informatiker, University of Ulm.

Duration: 1/3 1999–1/3 2002

Status: Ph.D. candidate, expected date of finish: March 2002

Financial Support: BRICS Ph.D. School

Advisor: Kim Guldstrand Larsen

Architectural hierarchies are much valued in system design, as they incorporate powerful mechanisms for data abstraction and high-level simulation in an early stage. The achieved modularity is a popular theme in specification formalisms such as UML or state-charts.

This Ph.D. project aims at providing foundations and methods to make use of hierarchical descriptions in a timed setting, in particular in connection with the real-time verification tool UPPAAL. It is believed, that hierarchies can be exploited in terms of more efficient model-checking algorithms, that make use of locality. Moreover, hierarchical structures provide a natural way to build property-preserving abstractions, which can scale up the applicability of formal verification tools significantly.

Case studies with heuristics for automated construction of abstraction incorporating hierarchies in the model checking tool MOCHA, that were conducted in these lines, suggest that there is potential for improvement in timed settings.

Routing in mobile ad-hoc networks

Name: Thomas Heide Clausen

Education: cand.polyt

Duration: 1/7 1999–1/2 2002

Status: In progress

Financial Support: Aalborg University

Advisor: P. Jacquet (INRIA Rocquencourt), W. A. Pietraszek, A. Skou

This thesis concerns mobile ad-hoc networks (MANETs), specifically routing and other network (OSI layer 3) issues. It focuses on development of efficient routing protocols, including practical experiments, simulation, modelling and formal reasoning on protocol properties. Different scenarios (as defined e.g. by communication and mobility patterns) require different routing paradigms. Thus one area of work includes studying and evaluating different protocols against different such scenarios. Another area of work results from the situation where several routing protocols co-exist. This involves the development of a meta-protocol for MANETs (i.e. a protocol providing a common, physical format for the messages exchanged between the nodes, in which other MANET routing protocols can be expressed and encapsulated) and protocols and principles for injecting routes from one protocol to another (i.e. a kind of "BGP for MANET's").

A Configurable Process for Design of Object-Oriented Software Architectures

Name: Birgitte Lønvg
Education: Cand. Scient.
Duration: 1995–2000
Status: In progress
Financial Support: Aalborg University
Advisor: Arne Skou

7.6 Services

7.6.1 Editorial Services

- Initiator and Steering Committee member of *Nordic Workshop on Programming Theory* series. Started 1989 with an average 50 persons participation. Informal meeting place for Nordic (junior) researcher in the area of programming theory (KGL)
- Editorial and Executive Board member of *Nordic Journal of Computing* (KGL)
- Steering Committee member of *CONCUR* conference series (KGL)
- Initiator and Steering Committee member of *TACAS*, Tools and Algorithms for the Construction and Analysis of Systems, series (KGL)
- Steering Committee of the European Educational Forum, *EEF* (KGL)
- Chairing at SOAP 98: Semantics of Objects as Processes, Aalborg, July 1998 (HH,JK,UN)

- Guest editor of a special issue of Mathematical Structures in Computer Science (MSCS) devoted to "The Difference between Concurrent and Sequential Computation" and to selected papers from EXPRESS'00: 7th International Workshop on Expressiveness in Concurrency. (LA)
- Guest editor of special issue of Information Processing Letters (IPL) on Process Algebra. (LA)
- Chairing at EXPRESS'00: 7th International Workshop on Expressiveness in Concurrency, August 2000. (LA)
- Chairing at HLCL'00: High Level Concurrent Languages 2000. (UN)
- Chairing at FTRTFT'98: Formal Techniques for Real Time and Fault Tolerant Systems, 1998, Lyngby. (APR)
- Chairing at ICALP'98: International Colloquium on Automata, Languages and Programming, Aalborg, July 1998, Aalborg. (KGL)
- Chairing at CONCUR'01: International Conference on Concurrency Theory, 2001, Aalborg. (KGL)
- Chairing at CAV'02: Computer Aided Verification, 2002, Copenhagen. (KGL)
- BRICS Autumn School on Verification, October 28 — November 1, 1996, Aarhus University. (KGL chair)

7.6.2 Conference and Workshop Organization

- CONCUR 01: International Conference on Concurrency Theory, 2001.
- CONCUR 00: International Conference on Concurrency Theory, 2000. (UN)
- ICALP 98: International Colloquium on Automata, Languages and Programming, Aalborg, July 1998. (KGL, HH, OHJ, JK)
- TACAS 95: Workshop on Tools and Algorithms for the construction and Analysis of Systems, Aarhus University, 19–20 May, 1995. (KGL)

7.6.3 Programme Committees

- SOAP 98: Semantics of Objects as Processes, Aalborg, July 1998. (HH, JK, UN)

- SOAP 99: Semantics of Objects as Processes, Lisbon, July 1999. (HH, JK, UN)
- FTRTFT 98: Symposium and School on Formal Techniques in Real-Time and Fault-Tolerant Systems, Lyngby, September 1998. (APR)
- EUROMICRO 98: Tenth Euromicro Real-Time Workshop, Berlin, June 1998. (APR)
- HSCC: Hybrid Systems Computation and Control, 1999. (APR)
- HSCC: Hybrid Systems Computation and Control, 2000. (APR)
- CONCUR 96: Seventh International Conference on Concurrency Theory, Pisa, Italy, August 26-29, 1996. (LA)
- EXPRESS 97: 4th International Workshop on Expressiveness in Concurrency, Santa Margherita Ligure, Italy, 8–12 September 1997. (LA)
- LICS 98: Thirteenth Annual IEEE Symposium on Logic in Computer Science, June 21–24, 1998, Indianapolis, Indiana, United States of America. (LA)
- EXPRESS 98: 5th International Workshop on Expressiveness in Concurrency, Sophia Antipolis, France, September 1998. (LA)
- EXPRESS 99: 6th International Workshop on Expressiveness in Concurrency, Eindhoven, The Netherlands, August 1999. (LA, UN)
- EXPRESS 00: 7th International Workshop on Expressiveness in Concurrency, August 2000. (LA)
- CONCUR 00: International Conference on Concurrency Theory. (UN)
- HLCL 00: High Level Concurrent Languages 2000. (UN)
- HART 97: International Workshop on Hybrid and Real-Time Systems, 1997. (KGL)
- TACAS 95-00: Tools and Algorithms for the Construction and Analysis of Systems. (KGL)
- NWPC 95-00: Nordic Workshop on Program Correctness. (KGL)
- CONCUR 97: Eighth International Conference on Concurrency Theory, 1997. (KGL)
- CAV 96-98 + 00: Computer Aided Verification. (KGL)

- IEEE Real Time Systems Symposium, 1998. (KGL)
- Hybrid Systems Workshop, 1998, 1999. (KGL)
- FTRTFT 98: Formal Techniques for Real Time and Fault Tolerant Systems, 1998. (KGL)
- ICALP98: Programme and Organizing Chairman, 1998. (KGL)
- AMAST workshop on Real Time Systems, 1998. (KGL)
- 11th Euromicro Conference on Real-Time Systems, 1999. (KGL)
- LICS 00: Logic in Computer Science, 2000. (KGL)
- DSVV 00: Distributed System Validation and Verification, 2000. (KGL)
- PAPM 00: 8th International Workshop on Process Algebra and Performance Modelling, 2000. (KGL)
- SPIN'00: The 7th International SPIN workshop on Model Checking of Software, 2000. (KGL)
- ADPM'00: Automation of Mixed Processes: Hybrid Dynamic Systems, 2000 (KGL)

7.6.4 Invited Presentations

- Ph.D. course for students of the Departments of Mathematics and Computer Science of the Universities of Bologna, Pisa and Siena, June 1998. (LA)
- Invited mini-course at the Technical University of Denmark, Copenhagen, 13 November 1997. (LA)
- Annual meeting of the Italian chapter of the European Association for Theoretical Computer Science, Florence, 17 October 1997. (LA)
- EXPRESS Workshop in Scholss Dagstuhl, Germany, September 16, 1996. (LA)
- Freie Universität Berlin. (JK)
- FTRTFT, Formal Techniques in Real-Time and Fault-Tolerant Systems, Uppsala, Sweden, September 1996. (KGL)
- Annual Symposium of Nederlandse Vereniging voor Theoretische Informatica, Utrecht, The Netherlands, February 1997. (KGL)

- Mathematical Foundations of Programming Systems, MFPS97, Pittsburgh, U.S.A., March 1997. (KGL)
- IFIP WG2.2, meeting in Graz, September 1997. (KGL)
- Real Time Systems Summerschool, Lidingö, Sweden, August 1998. (KGL)
- PROBMIV'99, Probabilistic Methods in Verification, Eindhoven, NL, August, 1999. (KGL)
- Verification of LARGE Systems, Chalmers Technical University, Sweden, December, 1999. (KGL)
- MOVEP'00, MOdelisation et VERification des Processus paralleles, Nantes, June, 2000. (KGL)
- 21st International Conference on Application and Theory of Petri Nets, Aarhus, June, 2000. (KGL)
- ADPM2000: The 4th International Conference of Automation of Mixed Processes, Dortmund, September, 2000. (KGL)
- Real-Time Systems, Denmark's Technical University, a series of lectures (see <http://www.cs.auc.dk/~kgl/DTU00/Plan.html>), 2000. (KGL)

7.6.5 Evaluation Committees

- Jose-Luis Vivas, Lic., KTH, Stockholm, March 1999. (HH)
- Dr. Troubitsyna, Doctor, Turku, June 2000. (APR)
- Dr. Goerigk, Doctor, Kiel, March 2000. (APR)
- Flavio Corradini, Ph.D., Università di Roma “La Sapienza”, February 1996. (LA)
- William Ferreira, D.Phil., University of Sussex, April 1996. (LA)
- Jørgen Hedegaard Andersen, Ph.D., Department of Computer Science, Aalborg University, December 1997. (LA)
- Kåre Jelling Kristoffersen, Ph.D., Department of Computer Science, Aalborg University, December 1998. (LA)
- Marcin Jurdziński, Progress report at the BRICS Ph.D. School, Aarhus University, January 1999. (LA)

- Gerd Behrmann, Thesis proposal, Department of Computer Science, Aalborg University, January 1999. (LA)
- Henrik Ejersbo Jensen, Ph.D., Department of Computer Science, Aalborg University, October 1999. (LA)
- Simone Tini, Ph.D., Department of Computer Science, Università di Pisa, January 2000. (LA)
- Thomas Troels Hildebrandt, Ph.D., BRICS Ph.D. School, Aarhus University, February 2000. (LA)
- Øystein Haugen, Ph.D., University of Oslo, 1997. (ASK)
- Keld Ramstedt Bach, Ph.D., Aalborg University, 1999. (ASK)
- Jørgen Sværke Hansen, Ph.D., Copenhagen University, 2000. (ASK)
- Marc Geilen, Ph.D., Technical University of Eindhoven, Eindhoven, The Netherlands, 2000. (KGL)
- Franc Valencia, Aarhus University (Progress), Denmark, 2000. (KGL)
- Lars Christiansen, Ph.D., Aarhus University, Denmark, 2000. (KGL)
- W.O.D. Griffioen, Ph.D., Nijmegen University, Nijmegen, The Netherlands, 2000. (KGL)
- Martin Weichert, Ph.D., Chalmers Technical University, Göteborg, Sweden, 1999. (KGL)
- Brahim Mammass, Ph.D., Universite Pierre & Marie Curie, LIP6, Paris, France, 1999. (KGL)
- Augusto Burgueno Arjona, Ph.D, Onera CERT, Toulouse, France, 1998. (KGL)
- Jesper G. Henriksen, Ph.D. Part A, Aarhus University, 1997. (KGL)
- Thomas T. Hildebrand, Ph.D. Part A, Aarhus University, 1997. (KGL)
- Anders Sandholm, Ph.D. Part A, Aarhus University, 1997. (KGL)
- Pim Kars, Ph.D., Twente University, 1997. (KGL)
- Allan Cheng, Ph.D., Aarhus University, 1996. (KGL)
- William Ferreira, Doctor of Philosophy, University of Sussex at Brighton, 1996. (KGL)

- Marcus Nielson, Lic, Department of Computer Systems, Uppsala University, Sweden, 2000 (KGL)

7.7 The Unit's Own Evaluation

In our opinion the research activities carried out by the new Distributed Systems and Semantics unit are very satisfactory, both in terms of quality and quantity. The unit has strengthened its international position on the semantic foundations of concurrency theory. This is witnessed, e.g., by the three invited chapters contributed by its members to the forthcoming *Handbook of Process Algebra*, and by its ongoing collaborations with some of the leading centres on concurrency theory.

The unit has continued its research activities on semantic theories for mobility and object-orientation. We are especially pleased with the work on the semantic foundations of object-oriented programs. The results concerning a translation of the imperative object calculus into a label-selective pi-calculus served as a starting point for the investigation of properties of the distributed object-oriented language Obliq. Among other things, the results obtained here have served the important purpose of clarifying aspects of the semantics of Obliq. Moreover, the proofs of the results on Obliq have shown the use and usefulness of the pi-calculus in this setting.

The research on verification and validation has been particularly successful. The main contributions offered by the unit in this area of research are the development of algorithmic foundations and their realization in verification tools whose performances are amongst the best within their application domains (real-time and embedded systems). The (academic) success of the verification tool UPPAAL is indisputable, as witnessed, e.g., by the high number of users of the developed tool(s), by the growing number of international research groups taking part in the continued development of UPPAAL, and by the Honorary Doctorate awarded by Uppsala University to Professor Kim G. Larsen in 1999.

Besides pushing the limits of UPPAAL, the conclusion of the research evaluation five years ago called for applications to industrial case studies. Consequently, significant research effort has been devoted to the pursuit of this ambition. Compared with the previous evaluation period, substantially more research has been conducted in direct collaboration with industry. In particular, the VVS project offered an excellent platform for such collaboration. Also, the collaborations with Bang&Olufsen on modelling and analysis of real-time communication protocols were successful. We are happy to see that this line of research has proved to be extremely rewarding both in terms of publications and in valuable feedback to the continuous development of the tool. In the future, we hope to be able to renew these as well as engage in other similar industrial collaborations.

On the experimental side, the work on network protocols and real-time sys-

tems bodes well for the future research in this direction. Based on experiments with the performance of TCP over high speed ATM channels, a new and more efficient transport protocol has been designed and experimentally evaluated; furthermore, the current work on IP router performance analysis has led to new insights on the feasibility of active queue policies; the unit is involved in the experimental evaluation of a forthcoming standard on wireless network routing protocols and, finally, experimental work is currently applying the UPPAAL tool to the automated analysis of observed traces from real-time systems.

The last couple of years have seen collaborative efforts between the formerly independent research groups now comprising the research unit. A recent example of this is the project involving MindPass. Such efforts are, of course, highly important and should be strengthened. Not only do they provide new results on the borderline between theory and applications, but they may also point to new and interesting research directions.

The number of publication by the members of the unit has remained very high. The high quality of the publications is evidenced by their appearance in well-respected journals, as invited contributions to special volumes and in the proceedings of competitive conferences. As was the case for the previous evaluation period, most publications appear in conference proceedings. This follows the general trend in our field. However, we are pleased to note that the proportion of journal publications remains at the high level that was achieved by the end of the previous evaluation period.

The members of the unit have been engaged in national and international collaborations. In particular, the unit is one of the two sites of BRICS (Basic Research in Computer Science), a centre of the Danish National Research Foundation. Moreover, the unit has engaged in two Esprit projects and in a number of projects with Danish funding. Finally, the unit has benefited from collaboration with research institutions and colleagues outside Denmark. In all cases, the basic funding provided by BRICS has been particularly useful.

The members of the unit have served on the programme committees of an impressive number of highly prestigious international conferences. Another main community service offered by the unit has been the hosting of the 25th International Colloquium on Automata, Languages, and Programming (ICALP) in 1998, fulfilling another ambition of the previous evaluation report. In the next five year period we will host CONCUR in 2001 and a member of the unit will be the programme committee chair for CAV in 2002. We view this as a clear indication that our group has achieved international recognition and visibility.

Since our last research evaluation, the number of people contributing to the research of the unit has increased somewhat. The unit has also experienced a significant growth in the number of external researchers visiting for extensive periods of time. More significantly, though, the number of permanent members of staff has increased beyond the critical mass, yielding a unit which is now less sensitive to the assumption of a continuous flow of visiting researchers. However,

if we are to realize the full spectrum of research (and teaching) encompassed by the unit, the number of (permanent) staff must continue to increase.

During the evaluation period the unit has had a large number of Ph.D. students; to a large extent, this was made possible by the ample level of funding. However, towards the end of the period it has become increasingly difficult to attract new Ph.D. students. This is mainly due to the heavy pull from industry. We hope to be able to counter this tendency—e.g., by offering Ph.D. studies which combine the full research spectrum of our unit and by exploiting the new structure of BRICS, in which the Aalborg group will also participate in the activities of the BRICS International Ph.D. school.

There are also causes for concern, however. Because of under-staffing, the teaching load of the permanent members of staff is quite high throughout the Department of Computer Science, and our unit is no exception. If this tendency continues, our research efforts are bound to suffer at some point in the not too distant future.

A further cause for concern stems from the fact that there is a discrepancy between the number of staff members within specific research areas and the number of masters-level students that we are able to attract to the corresponding topics. More specifically, only a small number of students are attracted to the areas related to semantics, where there is a bulk of active research and relatively many tenured researchers. On the other hand, quite a few students are actively interested in the design and implementation of distributed systems, an area which is currently under-staffed.

We are also concerned by an issue related to Ph.D. education within our unit. The expected duration of a Ph.D. project in Denmark is of three years. However, some of our Ph.D. students have taken a substantially longer time to complete their doctoral dissertations, and the delivery of some theses is long overdue. This unsatisfactory state of affairs needs to be discussed within our unit, and appropriate strategies to steer the development of Ph.D. projects must be developed.

7.8 Research Plan 2001–2005

This section describes the research plans of the unit for the next five year period. We begin by presenting the concrete research plans of the five existing research activities. Next we address possible areas of collaborative research involving researchers spanning several existing research activities.

7.8.1 Research Plans for the Existing Activities

Process Calculi The discovery of intriguing connections between process algebra and the classic areas of formal languages, automata theory and max-plus

algebra were, to our mind, amongst the most stimulating results we achieved in the period of this evaluation, and will pave the way to further studies of the equational logic of the mathematical structures that arise in these connections. In particular, in joint work with Zoltan Ésik (University of Szeged), we plan to study the equational logic of several semirings based on the max-plus and min-plus algebra of the natural numbers, the integers, the rationals and the reals. These structures are usually referred to as tropical or exotic semirings, and have applications in, e.g., automata theory, discrete event systems, discrete optimization and non-expansive mappings. Our goals will be to establish results to the effect that these tropical semirings are not finitely based, to study infinite equational axiomatizations for them, and to determine the complexity of deciding their equational theories. As an important step towards achieving these goals, we hope to be in a position to offer a one year visiting research professorship to Zoltan Ésik, with whom this work will be carried out.

We also intend to continue our collaboration with Wan Fokkink (CWI) on the study of the equational logic of processes. In particular, we shall work on establishing non-finite axiomatizability results for the whole family of n -nested simulation (pre)congruences over, e.g., the language of finite synchronization trees. Our joint work with Wan Fokkink will also be devoted to continuing our study of iteration operators in process algebra. For example, we shall aim at establishing non-finite axiomatizability results for several extensions of the language BPA^* with respect to, e.g., variations on bisimulation equivalence.

We shall also aim at studying extensions of Anna Ingólfssdóttir's work on verified implementations of value passing languages to the setting of variations on the pi-calculus.

We shall also work towards extending our theories in the direction of hybrid systems in the context of the European project VHS. The expressiveness results of [53] are particularly encouraging with respect to the future possibility of extending the efficient algorithms of UPPAAL to a large class of hybrid systems.

In collaboration with researchers at Twente University we plan to work on the development of process calculi allowing for both nondeterministic, real-time and probabilistic behaviour to be modelled and specified. This work will aim at providing the theoretical foundation for a planned effort in extending UPPAAL in the direction of performance analysis.

Finally, we expect to continue our work on the meta theory of process algebras by studying the automatic generation of fully abstract domain theoretic models for languages that allow for the description of processes with an infinite behaviour without the use of recursion. This work will build on the experience of [18], but we believe that it will involve intriguing new subtleties.

Over the next few years, and building on our contributions to the *Handbook of Process Algebra*, we also plan to pay more attention than in the past to the dissemination of results in our fields of research. To this end, Luca Aceto, Wan Fokkink and Chris Verhoef have reached an agreement with Cambridge University

Press for the publication of a research monograph based upon the handbook chapter [7] in their Cambridge Tracts in Theoretical Computer Science. Luca Aceto, Wan Fokkink and Anna Ingólfssdóttir will also aim at developing a web resource for the dissemination of results in process algebra.

Calculi for Mobility and Objects An important new research activity within the unit, pursued in collaboration with Björn Victor from Uppsala University, is that of investigating the mathematical properties of the spi-calculus and applying the results to the construction of a software tool for verifying properties of security protocols. A first, important step in this direction is to investigate the limits of equivalence checking in the spi-calculus, providing undecidability results. An early, as yet unpublished result is that even the finite-control fragment of the spi-calculus is Turing-powerful, destroying all hope of obtaining decision procedures for checking any non-trivial notion of behavioural equivalence. Ongoing research aims at defining a symbolic semantics of the spi-calculus which is in agreement with existing results and determining an interesting subset of the calculus where interesting notions of equivalence become decidable.

A longer-term goal is that of applying the spi-calculus and its associated reasoning principles to existing security protocols. A litmus test of the reasoning principles will be whether they can be used to prove already known properties of existing security protocols. This is an important source of possibilities for external collaboration.

The work on Obliq will continue. Most importantly, we would like to establish a result on the operational correspondence between pi-calculus encodings of Øjeblik and the operational semantics. Moreover, we intend to establish a number of equational laws for reasoning about Øjeblik. Finally, a longer-term goal is to use the ideas to understand the safety (or lack thereof) of migration in Java.

Another future activity is the development of a new version of the Mobility Workbench, an important software tool for verifying properties of pi-calculus processes; this work will be pursued in collaboration with Björn Victor from Uppsala University.

Finally, research activities concerning the object calculus will continue. A goal here is to establish a sensible labelled transition semantics, an appropriate notion of bisimulation equivalence and an associated equational theory for the imperative object calculus.

Verification and Validation The bulk of the IT industry is presently focused on office automation, e-commerce, and internet technology in general. Although the broadening of the Internet will continue at an impressive rate, the number of processor nodes can be increased an additional 50-fold by considering the number of embedded devices in the world. One of the distinguishing aspects of these devices is that they are equipped with sensors and/or actuators, and are directly

embedded into a physical environment in order to monitor and control it with little and in many cases no interaction with humans. Consequently, debugging such systems becomes more necessary, but also much harder, than for general purpose software. We see our research on verification and validation as becoming increasingly important for future software engineering practice. Future software engineering practice in the area of embedded systems also calls for research which bridges the gap between control theory and computer science. Again we see our research on timed and hybrid systems as useful contributions in this direction.

A new application domain for UPPAAL is that of scheduling theory. Compared with existing approaches, the use of timed automata provides an easy and flexible way of modelling scheduling problems, making the full range of real-time verification techniques available. Also, automatic synthesis of controllers is possible. Often, schedules satisfying various optimality criteria are sought. As a step towards applying UPPAAL to the synthesis of optimal schedules, we have introduced an extension of timed automata with linear prices [122, 123], and we have shown that the associated minimum-cost reachability question is decidable. This work has been carried out in close collaboration with Frits Vaandrager, Judi Romijn and Ansgar Fehnker from Nijmegen University, and we strongly believe that our future collaborative research will permit extensions of the zone-based techniques of UPPAAL to settle this decision problem efficiently.

For a three year period, Kim G. Larsen will be affiliated with the Department of Computer Science at Twente University as part-time industrial professor. This affiliation formalizes planned collaboration with Ed Brinksma, Joost-Pieter Katoen, Pedro D'Argenio and Holger Hermans from Twente University, on stochastic extensions of UPPAAL. The ambition is to create an efficient tool which allows for a wide range of performance analyses using model-checking techniques. The group at Twente has already achieved remarkable results in this direction for stochastic models based on Continuous Time Markov Chains, and we expect to advance this area even further as a result of the collaboration.

We will continue our work on exploiting structure to obtain more efficient verification techniques. A challenging goal is to transfer the success of the CBR technique for finite-state systems to the timed setting of UPPAAL. An obstacle to this transfer is the global synchronization between timed automata components due to the assumption of global time. This leaves very little room for exploitation of independent behaviour which was a crucial point in the finite-state case. Similar difficulties have been experienced by our colleagues at Uppsala University in applying the so-called partial order reduction techniques to the setting of UPPAAL. Nevertheless we will pursue this quest. Together with our collaborators at Uppsala University, a hierarchical extension of UPPAAL is currently being designed. We hope very much that our success in exploiting hierarchical structures for finite state systems [46] may be reused in the timed setting. In this respect we also expect to $\frac{1}{2}$ benefit from close contacts [118] with the research on MOCHA carried out by Rajeev Alur.

We want to continue the line of research initiated in collaboration with Franck Cassez on extending the efficient verification techniques of UPPAAL to richer models than that of timed automata. The work on stopwatch automata [53] is an important first step in this direction.

A current trend in formal verification tools is to apply the model-checking technique directly to the final source program (written in C or Java) rather than to abstract models of it. Our work on verification of NQC code may be seen as a small contribution in this direction. In our continued research in this direction we intend to bring in techniques known from abstract interpretation and partial interpretation. To boost our expertise on abstract interpretation, a new BRICS postdoc (Bertrand Jeannot) has been employed.

The development of verification and validation tools should of course be carried out in close contact with their applications. Thus, we firmly intend to continue this line of research as explained in more detail in the sections on Real-Time Systems and Networks. Also, application of verification tools needs to be complemented by a suitable design methodology. In particular, we intend to see how our tool efforts can be placed in the context of a design methodology based on the Unified Modelling Language (UML) adapted for embedded and real-time systems.

We also intend to complement our work on verification tools with techniques for validating systems through testing as explained in more detail in the section on Real-Time Systems. Again, here we expect to benefit from the collaboration with colleagues from Twente University, who are renowned for their long-standing contributions within testing theory and practice.

Finally, in light of the significant external interest in UPPAAL, we consider initiating a UPPAAL workshop along the lines of the existing SPIN workshop. Such a workshop would stimulate continued activity in applying and improving UPPAAL, as it would provide a forum for presenting such work. Preliminary contacts indicate the possibility of co-locating a UPPAAL workshop with the existing SPIN workshop.

Real-Time Systems As for real-time systems, the unit identifies two major areas for further research: specification methodology and system test. Also, there will be a particular focus on application of the results on autonomous systems.

The UML notation is gradually being accepted world wide as a specification language for most kinds of systems. However, the notation offers a huge number of facilities, many of which do not seem relevant within the realm of real-time systems. Work needs to be done on identifying parts that are relevant for specific application areas and design methodologies. The unit will pursue these questions during the coming five year period. Industrial case studies and industrial collaboration will be important inputs to this work.

It is generally accepted that system tests must to some extent be derived

automatically from specifications in order to ensure a good coverage. Despite the fact that useful specification formalisms are available, a break-through remains to be seen for automated test generation. This is even more true for real-time systems, where only a few attempts have been seen so far.

The above situation is partly due to the gap between on the one side the available formalisms and on the other hand the variety of system platforms in the different companies. The unit wants to contribute to fill this gap—based on close collaboration with one or more companies.

Experiments with design methodologies within the area of real-time and hybrid systems, in the spirit of the work reported in [98], will continue during the next five year period. A start is the unit's involvement in the design and development of an Autonomous Plant Care System. This is a collection of autonomous vehicles that shall supervise and work farm fields with high value crops. This project is undertaken jointly with the Department of Control Engineering, the Royal Veterinarian and Agricultural University, and the Danish Institute of Agricultural Sciences. The project commences in November 2000 and is currently funded for 2 1/2 years.

Networks As previously stated, it has become clear that the Internet will be the carrier of more and more services. It is also clear that the Internet will become increasingly mobile, e.g., through emerging wireless local area network technologies and the 3rd generation mobile telephone system. Much of our work in the near future will therefore focus on improvements of the Internet protocols, and in particular on mobility related problems.

We plan to further strengthen our research at the router level on efficient handling of web traffic and support for roaming in the Internet. Another challenge is to provide end-to-end quality of service guarantees for real-time traffic in heterogeneous networks consisting of mixed network types (e.g., ATM, best-effort Internet, wireless), or alternatively to devise a communication service which enables applications to adapt to changes in the network quality in a meaningful way.

We plan to collaborate with one of the local mobile telephone companies on software architectures for future generation communication services. The collaboration will include development of future protocols as well as validation of the implementations.

We wish to establish a substantial laboratory for practical experiments, and also to use the Danish research network which in the near future will provide excellent opportunities for experiments over distance between Danish research institutions.

Much of our work involves performance analysis. It would be very interesting to apply formal methods and tools to this application domain. Specifically, we believe that performance models formulated in a stochastic extension of the timed

automata as envisioned in Section 7.2.1 could serve as a powerful basis for such analysis.

7.8.2 A Strategy for the Unit

The unit of Distributed Systems and Semantics sees the creation of a theoretically well-founded development discipline for embedded software systems as an important goal for the coming years. The unit already covers the experimental side and the foundations of embedded software, and, in the period of this evaluation, it has been able to contribute significantly to tool development, which is at the intersection between practice and foundations. However, we must continue to strengthen the involvement in industrially related experiments, because they are essential for generating external funding, and because they are a strong motivation for students that want to enter the field. Furthermore we must carefully plan exploitation of the investment in tool development. Some promising directions are outlined below. We emphasize that the strategy is a supplement to continuation plans for existing research, as detailed for the five research directions in Section 7.8.1. It would be unwise to stifle these plans for individual groups because they produce results which contribute, albeit modestly, to the growth of our Science. It would also remove the possibility of pursuing the unexpected, which is often more fruitful than sticking to a plan. Thus, we expect these activities to continue.

However, in our choice of projects and industrial collaborations, we have an opportunity to act in accordance with the strategy, and we shall strive to do so. An encouraging sign is that, at the end of the current period, several younger members of the staff are engaged in such industrial collaboration projects although they have a background in foundations.

Another area, where we can apply our strategy, is in teaching. There is enormous pressure on the Department to provide high quality education within the area of embedded software both as a specialization and as an ingredient in master's programmes for engineering fields. We should work consistently towards the goal of teaching what we research.

Areas for joint research

In the coming period, we see the following promising directions for exploiting tool and theory developments in connection with practice. The list is not intended to be exhaustive, but we expect at least the first mentioned items to produce results in this period of planning.

Real-time Specification Methodology. The performance of UPPAAL has already allowed rather complex industrial systems to be analyzed. However, to become an integrated part of engineering practice it is important that the tool is

seen as part of a methodology for developing specifications and designs. Thus, we will be working on identifying and tailoring the parts of UML which are relevant for the design of real-time systems, and extending UPPAAL accordingly. In this line of research we will need close industrial collaboration based on case studies, continued development effort on UPPAAL with our collaborators in Uppsala, and ideally collaboration with colleagues from the Information Systems unit.

Testing. The work on verification and validation should be complemented by systematic testing of the actual implementations. In the setting of real-time, automatic test-generation is still rather poorly understood and we aim at contributing to fill this gap. This research will be conducted in close collaboration with one or more industrial partners and will build on our contact with Twente University.

Programming Languages. Several of the issues addressed by our research on distributed systems are not just properties of individual systems, but are intimately connected to the programming languages used for describing these systems. Also, an emerging trend in model-checking is to base the verification not on a model of the program but rather directly on the program itself. Thus, whereas so far the connection to programming languages has largely been implicit, we expect to work much more explicitly with this connection in the future.

Hybrid Systems. This constitutes a promising research area on the borderline between Computer Science and Control Theory. The mixture of continuous and discrete behaviour still poses challenges to both of the research communities. We hope to continue the work of VHS in a follow-up project and initiate collaboration with the Control Theory group within the university.

Modelling and Analysis of Systems and Networks. The planned work on extending UPPAAL to scheduling theory and performance analysis will potentially have several applications in the construction of real-time systems and in the analysis of the performance of different network protocols and policies. On the other hand, the planned extension may benefit from recent advances made within our unit on the understanding of $(max, +)$ -algebras. Also the addition of stochastic information to our models point at potential collaboration with people from the Decision Support Systems unit.

New Application Domains Traditionally, most of our case studies have been addressing issues in (real-time) communication protocols and (distributed) control programs. However, our modelling and verification frameworks extend to other application domains such as multimedia systems and E-Commerce where issues such as soft deadlines and security may be addressed using our techniques.

Thus, we intend to investigate the applicability of our frameworks in such alternative settings.

Conclusion

The strategy which we have outlined, together with the detailed areas mentioned here and in Section 7.2, constitute our plan for the coming years. We must be prepared to change its details, but intend to follow the overall strategy, and in particular to strive to implement it in future staff policy.

7.9 Bibliography

Edited Books

- [1] Luca Aceto and Björn Victor (editors). Preliminary Proceedings of the 7th International Workshop on Expressiveness in Concurrency, EXPRESS'00, (State College, USA, August 21, 2000), volume NS-00-2 of *BRICS Notes Series*, August 2000.
- [2] Allan Cheng, Kim G. Larsen, and Mogens Nielsen. *BRICS Autumn School on Verification*. BRICS, 1996. volume NS-96-1 of BRICS Notes Series.
- [3] Uwe Nestmann and Benjamin C. Pierce (editors). *HLCL '98: High-Level Concurrent Languages (Nice, France, September 12, 1998)*, volume 16.3 of *ENTCS*. Elsevier Science Publishers, 1998.
- [4] K.G. Larsen, S. Skyum, G. Winskel (editors). Automata, Languages and Programming. 25th International Colloquium, ICALP'98, Lecture Notes in Computer Science 1443, 1998.
- [5] Anders P. Ravn and Hans Rischel, (editors). *Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT98)*, volume 1486 of *LNCS*. Springer Verlag, 1998.

Chapters in Books

- [6] Luca Aceto, Wan Fokkink, and Chris Verhoef. Conservative extension in structural operational semantics. In Mogens Nielsen, editor, *The Concurrency Column, Bulletin of the EATCS*, pages 110–132. October 1999. Also available as BRICS report RS-99-24.
- [7] Luca Aceto, Wan Fokkink, and Chris Verhoef. Structural operational semantics. In Bergstra et al., *Handbook of Process Algebra*. Elsevier, 2000. To appear. Pp. 133. Also available as BRICS report RS-99-30.

- [8] Anna Ingólfssdóttir and Huimin Lin. A symbolic approach to value-passing processes. In Bergstra et al., *Handbook of Process Algebra*. Elsevier, 2000. To appear. Pp. 51.
- [9] Bengt Jonsson, Kim G. Larsen, Wang Yi: “Probabilistic Process Algebra”, In Bergstra et al., *Handbook of Process Algebra*. Elsevier, 2000. To appear. Pp. 31.
- [10] Kim G. Larsen and Paul Pettersson. *UPPAAL2k*, volume 70. Bulletin of the EATCS, 1999. Also available in BRICS Newsletter no 10.
- [11] Uwe Nestmann and Björn Victor. Calculi for mobile processes: Bibliography and web pages. *EATCS Bulletin*, 64:139–144, February 1998. Contribution to the Concurrency Column.

Refereed Journal Articles

- [12] Luca Aceto, Zoltan Ésik, and Anna Ingólfssdóttir. The max-plus algebra of the natural numbers has no finite equational basis. *Theoretical Computer Science*, 2000. To appear. Also available as BRICS Report RS-99-33.
- [13] Luca Aceto and Wan Fokkink. An equational axiomatization for multi-exit iteration. *Information and Computation*, 137(2):121–158, 15 September 1997.
- [14] Luca Aceto, Wan Fokkink, and Anna Ingólfssdóttir. A menagerie of non-finitely based process semantics over BPA^* —from ready simulation to completed traces. *Mathematical Structures in Computer Science*, 8(3):193–230, June 1998.
- [15] Luca Aceto, Wan Fokkink, and Anna Ingólfssdóttir. On a question of A. Salomaa: The equational theory of regular expressions over a singleton alphabet is not finitely based. *Theoretical Computer Science*, 209(1–2):163–178, 6 December 1998.
- [16] Luca Aceto, Wan Fokkink, Robert van Glabbeek, and Anna Ingólfssdóttir. Axiomatizing prefix iteration with silent steps. *Information and Computation*, 127(1):26–40, 25 May 1996.
- [17] Luca Aceto and Jan Friso Groote. A complete equational axiomatization for MPA with string iteration. *Theoretical Computer Science*, 211(1–2):339–374, 1999.
- [18] Luca Aceto and Anna Ingólfssdóttir. CPO models for compact GSOS languages. *Information and Computation*, 129(2):107–141, September 1996.

- [19] Luca Aceto and Anna Ingólfssdóttir. A characterization of finitary bisimulation. *Information Processing Letters*, 64(3):127–134, 14 November 1997.
- [20] Luca Aceto and Anna Ingólfssdóttir. Corrigendum: Volume 66 (1991), for the article “A Domain Equation for Bisimulation” by Samson Abramsky, pp. 161–218. *Information and Computation*, 2000. To appear.
- [21] Luca Aceto and David Murphy. Timing and causality in process algebra. *Acta Informatica*, 33(4):317–350, June 1996.
- [22] G. Behrman, K.G. Larsen, H.R. Andersen, H. Hulgaard, and J. Lind-Nielsen. Verification of hierarchical state/event systems using reusability and compositionality. *Formal Methods in System Design*, 2000. (full version) To appear.
- [23] Kirsten M. Hansen, Anders P. Ravn, and Victoria Stavridou. From safety analysis to formal specification. *IEEE Trans. Softw. Eng.*, 24(7):573–584, July 1998.
- [24] Hans Hüttel, Josva Kleist, Uwe Nestmann, and Massimo Merro. Aliasing models for mobile objects. *Information and Computation*, 2000. To appear.
- [25] Anna Ingólfssdóttir and Andrea Schalk. A fully abstract denotational model for observational precongruence. *Theoretical Computer Science*, 2000. To appear.
- [26] K. Kristoffersen, K.G. Larsen, P. Pettersson, and C. Weise: “Experimental Batch Plant — VHS Case Study 1 Using Timed Automata and UPPAAL”, Deliverable of ESPRIT-LTR Project 26270 VHS (Verification of Hybrid Systems). Full version to appear in Special Issue of European Journal of Control on “Verification of Hybrid Systems”.
- [27] K.G. Larsen, W. P. Pettersson, W. Yi: UPPAAL in a Nutshell. *International Journal on Software Tools for Technology Transfer*, volume 1, pages 134–152, 1997.
- [28] K.G. Larsen, B. Steffen, and C. Weise. Continuous modeling of real-time and hybrid systems: from concepts to tools. *International Journal on Software Tools for Technology Transfer*, volume 1, 1997.
- [29] Kim G. Larsen, Carsten Weise, Wang Yi and Justin Pearson. Clock Difference Diagrams (extended version). In *Nordic Journal of Computing* 6, pages 271–298, 1999.
- [30] Kim G. Larsen, Wang Yi. Time Abstracted Bisimulation: Implicit Specifications and Decidability. *Information and Computation*, vol. 134, no. 2, 1997.

- [31] J. Lind-Nielsen, H.R. Andersen, G. Behrman, H. Hulgaard, K. Kristoffersen and K.G. Larsen. Verification of Large State/Event Systems using Compositionality and Dependency Analysis (extended version). Accepted by Formal Methods in System Design.
- [32] Uwe Nestmann and Benjamin C. Pierce. Decoding choice encodings. *Information and Computation*, October/November 2000. To appear. Available as report BRICS-RS-99-42, Universities of Aalborg and Århus, Denmark, 1999.
- [33] Brian Nielsen and Gul Agha. Towards Re-usable Real-Time Objects. *The Annals of Software Engineering*, 7, 1999. Special Issue on Real-Time Software Engineering.
- [34] J. Staunstrup, K.G. Larsen, H.R. Andersen, H. Hulgaard, G. Behrman, K. Kristoffersen, J. Lind-Nielsen, H. Leerberg, A. Skou, and N.B. Theilgaard. Practical Verification of Embedded Software, IEEE Computer, Vol 33, No 5, May 2000.
- [35] T. Hune, K.G. Larsen, and P. Pettersson. Guided synthesis of control programs using UPPAAL. *Nordic Journal of Computer Science*, 2000. (full version) To appear.

Refereed Conference Articles

- [36] Luca Aceto, Patricia Bouyer, Augusto Burgueño, and Kim Guldstrand Larsen. The power of reachability testing for timed automata. In V. Arvind and R. Ramanujam, editors, *Proceedings of the Eighteenth Conference on the Foundations of Software Technology and Theoretical Computer Science*, volume 1530 of *Lecture Notes in Computer Science*, pages 245–256. Springer-Verlag, December 1998.
- [37] Luca Aceto, Augusto Burgueno and Kim G. Larsen: “Model Checking via Reachability Testing for Timed Automata. “, In Proceedings of the 4th International Workshop on Tools and Algorithms for the Construction and Analysis of Systems. Gulbenkian Foundation, Lisbon , Portugal, 31 March–2 April, 1998. LNCS 1384.
- [38] Luca Aceto, Zoltan Ésik, and Anna Ingólfssdóttir. On the two-variable fragment of the equational theory of the max-sum algebra of the natural numbers. In H. Reichel and S. Tison, editors, *Proceedings of the 17th International Symposium on Theoretical Aspects of Computer Science*, volume 1770 of *Lecture Notes in Computer Science*, pages 267–278. Springer-Verlag, February 2000.

- [39] Luca Aceto, Wan Fokkink, and Anna Ingólfssdóttir. A Cook’s tour of equational axiomatizations for prefix iteration. In M. Nivat, editor, *Foundations of Software Science and Computation Structures (Lisbon, 1998)*, Lecture Notes in Computer Science, pages 20–34, Berlin, 1998. Springer-Verlag.
- [40] Luca Aceto, Wan Fokkink, and Anna Ingólfssdóttir. 2-nested simulation is not finitely equationally axiomatizable, 2000. To appear in the *Proceedings of the 18th International Symposium on Theoretical Aspects of Computer Science* (STACS 2001)
- [41] Luca Aceto, Hans Hüttel, Anna Ingólfssdóttir, and Josva Kleist. Relating semantic models for the object calculus. *Electronic Notes in Theoretical Computer Science*, 7, 1997. Proceedings of EXPRESS 97 Workshop.
- [42] Luca Aceto and Anna Ingólfssdóttir. An equational axiomatization of observation congruence for prefix iteration. In Wirsing and Nivat, editors, *Algebraic Methodology and Software Technology: 5th International Conference, AMAST ’96, Munich, Germany*, volume 1101 of *Lecture Notes in Computer Science*. Springer-Verlag, July 1996, pages 195–209. Full version available as BRICS Report RS-95-5, January 1995.
- [43] Luca Aceto and Anna Ingólfssdóttir. Testing Hennessy-Milner logic with recursion. In Wolfgang Thomas, editor, *Foundations of Software Science and Computation Structures (Second International Conference, FOSSACS’99)*, volume 1578 of *Lecture Notes in Computer Science*, pages 41–55. Springer-Verlag, March 1999.
- [44] Luca Aceto and François Laroussinie. Is your model checker on time? On the complexity of model checking for timed modal logics. In *Mathematical Foundations of Computer Science 1999 (Szklarska Poreba)*, volume 1672 of *Lecture Notes in Computer Science*, pages 125–136. Springer-Verlag, 1999.
- [45] Jørgen H Andersen, Carsten H Kristensen, and Arne Skou. Specification and Automated Verification of a Real-Time Behaviour. A Case Study. *Annual Reviews in Control*, 20:55–70, 1996.
- [46] G. Behrman, K.G. Larsen, H.R. Andersen, H. Hulgaard and J. Lind-Nielsen: “Verification of Hierarchical State/Event Systems using Reusability and Compositionality”, In Proceedings of TACAS’99. Lecture Notes in Computer Science 1579.
- [47] G. Behrmann, K.G. Larsen, J. Pearson, C. Weise, W. Yi: “Efficient Timed Reachability Analysis Using Clock Difference Diagrams”, In Proceedings of Computer Aided Verification 1999. Lecture Notes in Computer Science 1633.

- [48] Gerd Behrmann, Thomas Hune and Frits Vaandrager. Distributed Timed Model Checking—How the Search Order Matters. In *Proc. of 12th International Conference on Computer Aided Verification*, Lecture Notes in Computer Science, Springer-Verlag 2000.
- [49] J. Bengtsson, W.O.D. Griffioen, K.J.Kristoffersen, K.G. Larsen, F. Larsson, P. Pettersson, W. Yi “Verification of an Audio Protocol with Bus Collision using UPPAAL”, *Computer Aided Verification*. Lecture Notes in Computer Science 1102, 1996.
- [50] J. Bengtsson, K.G. Larsen, F. Larsson, P. Pettersson, W. Yi “UPPAAL: a tool suite for automatic verification of real-time systems”, *Hybrid Systems III, Verification and Control, DIMACS*. Lecture Notes in Computer Science 1066, 1996.
- [51] Johan Bengtsson, Kim G. Larsen, Fredrik Larsson, Paul Pettersson, Wang Yi and Carsten Weise.: “New Generation of UPPAAL.” In proceedings of the International Workshop on Software Tools for Technology Transfer. Aalborg, Denmark, 12–13 July, 1998.
- [52] J. Bengtsson, K.G. Larsen, P.Pettersson, W. Yi “UPPAAL in 1995”, *Tools and Algorithms for the Construction and Analysis of Systems*. Lecture Notes in Computer Science 1055, 1996.
- [53] F. Cassez, K.G. Larsen: “The Impressive Power of Stopwatches”, In C. Palamidessi, editor, *CONCUR 2000—11th International Conference on Concurrency Theory, University Park, PA. USA*, pages 138–152, Lecture Notes in Computer Science 1877, Springer-Verlag 2000.
- [54] Mikkel Christiansen, Kevin Jeffay, David Ott, and F. Donelson Smith. Tuning red for web traffic. In *Proceedings of SIGCOMM*, August/September 2000. to appear.
- [55] Anders Strandløv Elkjær, Micheal Höhle, Hans Hüttel, and Kasper Overgård. Towards automatic bisimilarity checking in the spi calculus. In C. S. Calude and M. J. Dinneen, editors, *Combinatorics, Computation & Logic*, volume 21 of *Australian Computer Science Communications*, pages 175–189. Springer-Verlag Singapore Pte. Ltd., January 1999. As part of the *Australian Computer Science Week*, January 18–21, 1999.
- [56] Lone Halkjaer, Karen Haervig and Anna Ingólfssdóttir. Verification of the legOS scheduler using Uppaal. In Flavio Corradini and Paola Inverardi, editors, *MTCS 2000, International Workshop on Models for Time-Critical Systems, State College, USA*, volume 39 of *Electronic Notes in Theoretical Computer Science*, pages 80–99. Elsevier, 2000. To appear.

- [57] Klaus Havelund, Kim G. Larsen, and Arne Skou. Formal Verification of a Power Controller Using the Real-Time Model Checker UPPAAL. In Joost-Pieter Katoen, editor, *Proceeding of the 5th International AMAST Workshop, ARTS'99 on Formal Methods for Real-Time and Probabilistic Systems*, volume 1601 of *Lecture Notes in Computer Science*, pages 277–298. Springer, 1999.
- [58] Klaus Havelund, Arne Skou, Kim G. Larsen and Kristian Lund: “Formal Modelling and Analysis of an Audio/Video Protocol: An Industrial Case Study Using UPPAAL.”, In Proceedings of the 18th IEEE Real-Time Systems Symposium, pages 2-13. San Francisco, California, USA, 3-5 December 1997.
- [59] Thomas Hune: “Modelling a real-time language”, In Proceedings, 4th Workshop on Formal Methods for Industrial Critical Systems, FMICS'99.
- [60] Thomas Hune, Kim G. Larsen and Paul Pettersson. Guided Synthesis of Control Program Using Uppaal. In *Proceedings of the IEEE ICDCS International Workshop on Distributed Systems Verification and Validation*, pages E15-E22. IEEE Computer Society Press 2000.
- [61] Hans Hüttel, Josva Kleist, Dan S. Andersen, and Lars H. Pedersen. Object types and modal formulae. In *Proceedings of the Fourth International Workshop on Foundations of Object-Oriented Languages (FOOL4)*, Paris, January 1997.
- [62] Hans Hüttel, Josva Kleist, Uwe Nestmann, and Davide Sangiorgi. Migration = cloning ; aliasing (preliminary version). In *Informal Proceedings of the Sixth International Workshop on Foundations of Object-Oriented Languages (FOOL 6)*, 1999. Sponsored by ACM/SIGPLAN.
- [63] Hans Hüttel, Josva Kleist, Uwe Nestmann, and Davide Sangiorgi. Surrogates in øjeblik: Towards migration in Obliq. In Hans Hüttel and Uwe Nestmann, editors, *SOAP '98: Semantics of Objects as Processes (Aalborg, Denmark, July 18, 1998)*, volume NS-98-5 of *BRICS Notes Series*, pages 43–50, 1998.
- [64] Anna Ingólfssdóttir. Weak semantics based on lighted button pressing experiments: An alternative characterization of the readiness semantics. In Dirk van Dalen, editor, *Proceedings of the 10th Annual International Conference of the European Association for Computer Science Logic, CSL '96, Utrecht*, volume 1258 of *Lecture Notes in Computer Science*, pages 226–243. Springer-Verlag, 1997.
- [65] Anna Ingólfssdóttir and Rosario Pugliese. Towards verified lazy implementation of concurrent value-passing languages. In Ilaria Castellani and

- Björn Victor, editors, *Electronic Notes in Theoretical Computer Science*, volume 27. Elsevier Science Publishers, 2000.
- [66] Anna Ingólfssdóttir and Andrea Schalk. A fully abstract denotational model for observational precongruence. In Büning, editor, *Proceedings CSL 95*, Paderborn, Germany, volume 1092 of *Lecture Notes in Computer Science*. Springer-Verlag, 1996, pages 335–361.
- [67] Torsten K. Iversen, Kåre J. Kristoffersen, Kim G. Larsen, Morten Laursen, Rune G. Madsen, Steffen K. Mortensen, Paul Pettersson and Chris B. Thomasen. Model-Checking Real-Time Control Programs. In *Proceedings of the 12th Euromicro Conference on Real-Time Systems*, pages 147-155. IEEE Computer Society Press 2000.
- [68] Peter K. Jensen. Automated modeling of real-time implementation. *Doctoral Symposium Proceedings of the 13th IEEE Automated Software Engineering*, October 1998.
- [69] Morten Vadskær Jensen and Brian Nielsen. Real-Time Layered Video Compression using SIMD Computation. In *Austrian Center for Parallel Computation (ACPC'99)*, February 1999. Special Track on Image Processing, Video Processing and Multimedia.
- [70] Henrik Ejersbo Jensen, Kim G. Larsen, and Arne Skou. Modelling and Analysis of a Collision Avoidance Protocol Using SPIN and UPPAAL. In J-C. Gregoire, G.J. Holzmann, and D.A. Peled, editors, *Proceedings Second Workshop on the SPIN Verification System*, American Mathematical Society, DIMACS/32, pp. 33-49, 1996.
- [71] Henrik Ejersbo Jensen, Kim Guldstrand Larsen, and Arne Skou. Scaling up Uppaal – Automatic Verification of Real-Time Systems using Compositionality and Abstraction. In *Proceedings Sixth International School and Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems, FTRTFT 2000*, Pune, India, Springer LNCS 1926, pp. 19-30.
- [72] Henrik E. Jensen and Nancy A. Lynch. A Proof of Burns N -Process Mutual Exclusion Algorithm Using Abstraction. In Bernhard Steffen, editor, *TACAS'98, Tools and Algorithms for the Construction and Analysis of Systems*, volume 1384 of *Lecture Notes in Computer Science*, Lisbon, Portugal, March/April 1998, Springer LNCS 1384, pp. 409-423.
- [73] Josva Kleist, Massimo Merro, and Uwe Nestmann. Local π -calculus at work: Mobile objects as mobile processes. In *Proceedings of TCS'2000*, LNCS 1872, pages 390–408, 2000.

- [74] Josva Kleist and Davide Sangiorgi. Imperative objects and mobile processes. In David Gries and Willem-Paul de Roever, editors, *Proceedings of Procomet '98*, pages 285–303. IFIP–The International Federation for Information Processing, Chapman & Hall, 1997.
- [75] K.J. Kristoffersen and F. Laroussinie and K.G. Larsen and P. Pettersson and W. Yi “A Compositional Proof of a Real-Time Mutual Exclusion Protocol”, In Proceedings for FASE 1997. Lecture Notes in Computer Science 1214.
- [76] F. Laroussinie, K.G. Larsen: “CMC: A Tool for Compositional Model-Checking of Real-Time Systems.”, In Proceedings of FORTE/PSTV Conference, Paris, 1998.
- [77] K.G. Larsen: “Verification of Timed and Hybrid Systems” Invited paper for 21st International Conference on Application and Theory of Petri Nets, June 2000.
- [78] K.G. Larsen: “Tool Supported Validation of Timed and Hybrid Systems” Invited paper for ADPM2000: The 4th International Conference of Automation of Mixed Processes, Dortmund, September, 2000.
- [79] Kim G. Larsen, Christel Baier, Ed Brinksma: Validation of Continuous-Time Markov Systems. In Proceedings of PROBMIV’99, Probabilistic Methods in Verification.
- [80] Kim G. Larsen, Fredrik Larsson, Paul Pettersson and Wang Yi: “Efficient Verification of Real-Time Systems: Compact Data Structure and State-Space Reduction”, In Proceedings of the 18th IEEE Real-Time Systems Symposium, pages 14-24. San Francisco, California, USA, 3-5 December 1997.
- [81] Kim G. Larsen, Paul Pettersson and Wang Yi: “UPPAAL: Status and Developments”, In Proceedings of the 9th International Conference on Computer-Aided Verification. Haifa, Israel, 22-25 June 1997.
- [82] K.G. Larsen, B. Steffen, C. Weise “Fischer’s Protocol Revisited: a simple proof using modal constraints”, Hybrid Systems III, Verification and Control, DIMACS. Lecture Notes in Computer Science 1066, 1996.
- [83] Kim G. Larsen, Bernhard Steffen and Carsten Weise “The Methodology of Modal Constraints”, In “Formal Systems Specification: the RPC–memory specification case study”, Editors: M. Broy, S. Merz, K. Speis. Lecture Notes in Computer Science 1169, 1996.
- [84] Kim G. Larsen, Carsten Weise, Wang Yi and Justin Pearson: “Clock Difference Diagrams.”, In Proceedings of Nordic Workshop on Programming Theory, Turku, 1998.

- [85] Fredrik Larsson, Paul Pettersson and Wang Yi. On Memory-Block Traversal Problems in Model-Checking Timed Systems. In *Proceedings of 6th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 127–141. Lecture Notes in Computer Science 1785. Springer–Verlag 2000.
- [86] J. Lind-Nielsen, H.R. Andersen, G. Behrman, H. Hulgaard, K. Kristoffersen and K.G. Larsen: “Verification of Large State/Event Systems using Compositionality and Dependency Analysis“, In Proceedings of the 4th International Workshop on Tools and Algorithms for the Construction and Analysis of Systems. Gulbenkian Foundation, Lisbon , Portugal, 31 March –2 April, 1998. LNCS 1384.
- [87] Zhiming Liu, Anders P. Ravn, and Xiaoshan Li. Verifying duration properties of timed transition systems. In D. Gries and W.-P. de Roever, editors, *Proceedings of the IFIP Working Conference on Programming Concepts and Methods, PROCOMET’98*, pages 327–345. Chapman & Hall, 1998.
- [88] Birgitte Lønvig. How OO methods/processes can benefit from domain analysis and software architectures. In *OOPSLA 97 workshop on Object Technology, Architectures, and Domain Analysis—What’s the Connection?—Is there a Connection?*, 1997. Also available as CMU Technical Report.
- [89] Birgitte Lønvig. Applying Use Cases for Specification of Functional Requirements on a Telecommunication Network Management System. In Max Muhlhauser, editor, *Special Issues in Object-Oriented Programming, Workshop Reader of the 10th ECOOP 96*, pages 244–, Heidelberg, 1997. dPunkt Verlag.
- [90] Birgitte Lønvig. Process Enactment Inconsistency,—Experiences and Viewpoints. In *ICSE 97 workshop on Living with Inconsistency*, Boston, 1997.
- [91] Uwe Nestmann. On the expressive power of joint input. In Palamidessi and Castellani, *EXPRESS ’98: Expressiveness in Concurrency (Nice, France, September 7, 1998)*, volume 16.2 of *ENTCS*. Elsevier Science Publishers, 1998.
- [92] Uwe Nestmann. Mobile objects (a project overview). In Katharina Spies and Bernhard Schätz, editors, *Proceedings of FBT ’99: Formale Beschreibungstechniken für verteilte Systeme (München, Juni 1999)*, Wissenschaft, pages 155–164. Herbert Utz Verlag, 1999.
- [93] Uwe Nestmann, Hans Hüttel, Josva Kleist, and Massimo Merro. Aliasing models for object migration. In Amestoy et al., *Euro-Par ’99 Parallel Processing, 5th International Euro-Par Conference (Toulouse, France)*, volume

1685 of *LNCS*. Springer, August 1999., pages 1353–1368. The full version of this paper has been accepted for publication in *Journal of Information and Computation*.

- [94] Uwe Nestmann and António Ravara. Semantics of objects as processes (soap). In Ana Moreira and Serge Demeyer, editors, *ECOOP '99 Workshop Reader*, volume 1743 of *LNCS*, pages 314–325. Springer, October 1999. An introduction to, and summary of, the 2nd International SOAP-Workshop.
- [95] Brian Nielsen, Shangping Ren, and Gul Agha. Specification of Real-Time Interaction Constraints. In *1st International Symposium on Object-Oriented Real-Time Distributed Computing*, pages 206–214. IEEE, April 1998. Kyoto, Japan.
- [96] Brian Nielsen and Arne Skou. Automated Test Generation from Timed Automata. In *21st IEEE Real-Time Systems Symposium Work-in Progress Paper 2000, RTSS2000*, Walt Disney World, Orlando, Florida, USA, November 2000. IEEE. To appear.
- [97] Ernst-Rüdiger Olderog and Anders P. Ravn. Documenting design refinement. In *Proc. Formal Methods in Software Practice (FMSP'00)*. ACM, 2000.
- [98] Anders P. Ravn, Thomas J. Eriksen, Michael Holdgaard, and Hans Rischel. Engineering of real-time systems with an experiment in hybrid control. In G. Rozenberg and F. W. Vaandrager, editors, *Embedded Systems*, volume 1494 of *LNCS*, pages 316–352. Springer-Verlag, 1998.
- [99] Anders P. Ravn and Hans Rischel. Real-time constraints through the ProCoS layers. In E.-R. Olderog and B. Steffen, editors, *Correct System Design: Recent Insights and Advances*, volume 1710 of *LNCS*, pages 61–78. Springer-Verlag, 1999.
- [100] Mauno Rönkkö and Anders P. Ravn. Action systems with continuous behaviour. In P. J. Antsaklis, W. Kohn, M. Lemmon, A. Nerode, and S. Sastry, editors, *Hybrid Systems V*, volume 1567 of *LNCS*, pages 304–323. Springer-Verlag, 1999. also TUCS–Turku Centre for Computer Science Techn. Rep. 151, Dec. 1997, <http://www.tucs.abo.fi/publications/techreports/TR151.html>.

Dissertations

- [101] Jørgen H. Andersen. *Compositional Modal Logics*. Ph.D. thesis, Department of Computer Science, Aalborg University 1997.

- [102] Henrik Ejersbo Jensen. *Abstraction-Based Verification of Distributed Systems*. Ph.D. thesis, Department of Computer Science, Aalborg University, June 1999.
- [103] Josva Kleist. *Reasoning about Objects using Process Calculus Techniques*. Ph.D. thesis, Department of Computer Science, Aalborg University, 2000.
- [104] Kåre J. Kristoffersen. *Compositional Verification of Concurrent Systems*. Ph.D. thesis, Department of Computer Science, Aalborg University 1999.
- [105] Birgitte Lønvig. *A Configurable Process for Design of Object-Oriented Software Architectures*. Ph.D. thesis, Department of Computer Science, Aalborg University, Denmark, 2000. Submitted.
- [106] Brian Nielsen. *Specification and Test of Real-Time Systems*. Ph.D. thesis, Department of Computer Science, Aalborg University, Denmark, April 2000.

Technical Reports

- [107] J.H. Andersen, K.G. Larsen : “Compositional Safety Logics”, BRICS report RS-97-13.
- [108] Mark Astley and Thomas Heide Clausen. The Actor Foundry: A Modular Platform for Distributed Object Systems. Submission Status: Unknown, 1999.
- [109] Thomas Heide Clausen. *Actor Foundry—a QuickStart*. Aalborg University, Department of Computer Science, Aalborg, Denmark, 1998.
- [110] Mikkel Christiansen, Jesper L. Hagen, Brian Nielsen, Arne Skou, and Kristian Q. Skov. The Design and Implementation of ZCRP Zero Copying Reliable Protocol. Technical Report R-98-5011, Aalborg University, Department of Computer Science, Aalborg, Denmark, 1998.
- [111] Thomas Husfledt, Brian Nielsen, Finn Norman Pedersen, and Dao Van The. Experiments with Video Communications on ATM-networks. Technical Report R-97-5002, Aalborg University, Institute for Electronic Systems, Department of Computer Science, Aalborg, Denmark, June 1997.
- [112] Hans Hüttel and Josva Kleist. Objects as mobile processes. Technical report RR-96-38, BRICS—Basic Research in Computer Science, 1996.
- [113] H. Huttel and S. Shukla. On the complexity of deciding behavioural equivalences and preorders. Technical Report SUNYA-CS-96-03, State University of New York at Albany, December 28, 1996.

- [114] Philippe Jacquet, Paul Muhlethaler, Amir Qayyum, Anis Laouiti, Laurent Viennot, and Thomas Clausen. Optimized Link State Routing Protocol. INTERNET-DRAFT: draft-ietf-manet-olsr-02.txt, July 2000.
- [115] Morten Vadskær Jensen and Brian Nielsen. Design and Implementation of an Efficient, Layered Video Codec for Heterogeneous Networks. Technical Report R-98-5008, Aalborg University, Institute for Electronic Systems, Department of Computer Science, Aalborg, Denmark, November 1998.
- [116] Jesper Krogh and Jan Justesen. A Fault Tolerant Mutual Exclusion Algorithm Handling Dynamic Processes Using Message Passing. Technical Report, Aalborg University, Department of Computer Science, Aalborg, Denmark, June 1999.
- [117] Birgitte Lønvg. Order, Invoice and Payment System—A Rational Objectory Process Example. Technical Report R-98-5003, Aalborg University, Department of Computer Science, Aalborg, Denmark, 1998.
- [118] Oliver M. Möller, Rajeev Alur. Heuristics for Hierarchical Partitioning with Application to Model Checking. BRICS Research Report RS-00-21, Department of Computer Science, University of Aarhus, 2000.
- [119] Mauno Rönkkö and Anders P. Ravn. Differential equations as actions. Technical Report 109, TUCS—Turku Centre for Computer Science, Turku, Finland, April 1997.
- [120] Mauno Rönkkö and Anders P. Ravn. Switches and jumps in hybrid action systems. *Proceedings of the Estonian Academy of Sciences*, 4/2:106–118, 1998. also TUCS—Turku Centre for Computer Science Techn. Rep. 152, Dec. 1997, <http://www.tucs.abo.fi/publications/techreports/TR152.html>.

Others

- [121] G. Behrmann, K.J. Kristoffersen, K.G. Larsen: “Context Dependent Minimization of State/Event Systems”, in Proceedings of Nordic Workshop of Programming Correctness, 1997.
- [122] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim G. Larsen, Paul Pettersson, Judi Romijn, and Frits Vaandrager. Minimum-Cost Reachability for Priced Timed Automata. Submitted for publication 2000.
- [123] Gerd Behrmann, Ansgar Fehnker, Thomas Hune, Kim G. Larsen, Paul Pettersson, and Judi Romijn. Efficient Guiding Towards Cost-Optimality in UPPAAL. Submitted for publication 2000.

- [124] Bodentien, Nicky O. and Vestergaard, Jacob and Friis, Jakob and Kristoffersen, Kåre J. and Larsen, Kim G.: “Verification of State/Event Systems by Quotienting”, In Presented at *Nordic Workshop in Programming Theory*, Uppsala, Sweden, October 6–8, 1999. BRICS report RS-99-41.
- [125] Mikkel Christiansen, Kevin Jeffay, David Ott, and F. Donelson Smith. Tuning red for web traffic. Submitted for ACM Transactions on Networking.
- [126] Mikkel Christiansen. Mscpdll—a language for behavioral abstraction. In *Proceedings of Nordic Workshop on programming theory*, 1997.
- [127] Henrik Ejersbo Jensen. Model Checking Probabilistic Real Time Systems. In Bror Bjerner, Marie Larsson, and Bengt Nordström, editors, *Proceedings of the 7th Nordic Workshop on Programming Theory*, Report 86, Chalmers University of Technology, 1996, pp. 127-140.

Chapter 8

Information Systems

8.1 Profile of the Unit

The Information System Unit conducts research in the application of information technology (IT). The research field is development and use of computerised systems in organisational contexts. The research approach is empirical and experimental. The types of results include:

- Theoretical frameworks as contributions to understanding the field.
- Methodologies and techniques as contributions to guiding practical action.
- Changes as contributions to improving practice within the field.

The research combines an interest in software systems with an interest in the organisational issues of information systems lately supplemented with an interest in human-centred issues.

The research is based on and makes contributions to the intersection between software engineering (as traditionally known within computer science) and information systems (as traditionally known within business administration); lately this has been supplemented with human-computer interaction. Intellectual support is sought in adjacent disciplines: engineering, systems science, organisation theory, sociology, media science, theory of research, and philosophy. The research approaches applied by the unit encompass action research, case research, laboratory experiments, and field experiments. The research is based on an interpretative paradigm.

The target of the research is the professional practitioner working with the development and use of IT in the broadest sense. That includes systems development, systems use, IT management, and use of IT in innovation and change. To contribute to the professionalisation of the IT field it is assumed that the research must address the technical aspects of systems integrated with their human,

organisational, and social contexts. The research thus seeks to improve professionals' ability to engineer systems as well as their ability to plan and manage effective organisational interventions.

8.2 Activities and Results

The activities of the information system research unit combine an established, strong orientation towards software development methodologies and tools with a new and complementary focus on management and use of information technology. The majority of these activities have been conducted in research programs and projects of varying size. These programs and projects are described below.

Throughout the whole period from 1996 to 2000 several activities focussing primarily on software engineering have been conducted within two research programs: the software factory programme and the software modelling fundamentals programme. Both programs have served as a frame for investigating basic issues in software modelling and development. Related research activities in the OOA&D (Object Oriented Analysis & Design) project have been concerned with further development and refinement of an object-oriented software development methodology. Finally, a new position for research and practice in systems development has been formulated as Reflective Systems Development. This position is based on many of the systems development research activities that have been carried out in the research unit over the past decade.

Management issues related to software engineering have played some role throughout the research unit's activities. During this period, however, there has been a strong emphasis on these issues in the Software Process Improvement project. This project has involved many of the members of the unit, researchers from other institutions and tight collaboration with software engineering organisations. The goal has been to explore to what extent and in what ways maturity models and related frameworks for organisational learning can facilitate professionalisation of software organisations.

E-commerce is the main topic of three projects: the PITNIT (Process Integration and Transformation based on new Network Information Technology) project, the LOKNIT (a Danish acronym for Management, Organisation, and Competence Development Based on Network Information Technology) project, and in the EDI (electronic data interchange) Diffusion and e-commerce project. The focus is on the design and management of IT that transcends organisational and functional boundaries. Primarily network economics and emergence theories are applied as theoretical vehicles to explain the evolution and design of these new technologies.

Activities involving a fundamentally different perspective on use of information technology have emerged in the research unit since 1996. Here, the focus is on the individual level where a user is interacting with a computerised system.

These efforts started with a Ph.D. project dealing with functional analysis and development methods for web applications. The Staging project that started in 1998 deals more generally with interactive media, but still with a primary focus on an individual level. In 1999, these activities provided inspiration for a project where the aim is to establish a modern usability laboratory that can be used in empirical evaluations of web applications, mobile devices, and other kinds of interactive media.

The Software Factory Program (SF)

The purpose of this program has been to explore the theoretical and conceptual aspects of a modern software development process, which through the program was called software factory. The process is based on components with high complexity, collaboration between software developers, and industrial production of software. The program was set up as collaboration between the departments' research groups in software engineering, programming, and databases. The groups' contributions were related to object-oriented modelling, object-oriented design methodology, and software development processes. The program lasted from 1995 to 1997 with the whole group as participants. The program was supported on a small scale with 1.2Mkr.

The Software Modelling Fundamentals Program (SMF)

The purpose of this program has been to further explore the theoretical and conceptual aspects of modern software development. The program was set up as collaboration between the departments' research groups in software engineering, programming, and databases. The groups' results fell within concepts and methodologies for modelling in analysis and design, and theory on the role of modelling in software development. The program lasted from 1998 to 2000 with the most of the group as participants. The program was supported on a small scale with 0.8Mkr.

Object-Oriented Analysis & Design (OOA&D)

The OOA&D project has lasted since 1990 with Lars Mathiassen, Andreas Munk-Madsen (Metodica), Peter Axel Nielsen, and Jan Stage as members of the project group. The purpose is to develop a coherent methodology for analysis and design of computerised systems and to study its use in national and international software development educations and organisations. The methodology combines a classical functional approach with recent object-oriented thinking. It combines a Scandinavian approach to system development with experiences from software engineering.

The main effort during 1996 to 2000 falls into four streams of activity. First, the separate elements of the methodology that were published in 1993 and 1995 have been integrated into a coherent overall framework. Second, the design rationale underlying the methodology has been defined explicitly. Third, a considerable number of object-oriented patterns and contemporary architectures have been integrated. Fourth, the methodology has been extended in order to make it useful also for information system development and thereby disseminate it to a wider international audience.

The coherent methodology is reported in an academic book published in Danish, Swedish, and English.

Reflective Systems Development

Reflective Systems Development (RSD) is an approach to systems development research and practice that has been developed by Lars Mathiassen. RSD is published as a Dr. Techn. Thesis and summarised in a journal paper. The thesis was successfully defended September 11th, 1998, at Aalborg University. The underlying perspective of RSD is that systems developers are faced with uncertain, ambiguous, and conflicting situations in which knowledge about the specific situation has to be created. This situational knowledge should then be combined with state-of-the-art techniques to arrive at satisfactory solutions. Systems development research and practice are from this point of view intrinsically related: skillful practitioners are researchers into their own practice, and academic researchers must start out from and interact closely with systems development practice to arrive at relevant knowledge.

Computer-Supported Co-operative Work

The project work has been collaboration between the group and researchers at Risø National Laboratory. The focus has been on expanding and testing a theory of computer-supported co-operative work called Co-ordination Mechanisms. The focus has in particular been on modelling co-operative work and computerised co-ordination mechanisms. Lars Mathiassen, P. A. Nielsen and J. Stage have participated in this together with PhD students M. Divitini, S. Herskind, B. Krogh, and M. Nielsen.

Software Process Improvement (SPI)

The Software Process Improvement (SPI) project started January 1997 and ran for 3 years until December 1999. More than 15 researchers and 4 software organisations took part in the project. The total budget was DKK 30 million, DKK 10 million were sponsored by the Danish Ministries of Research and Industry, and DKK 20 million were sponsored by the participating software organisations and

research institutions. The project was managed from the IS unit and Ivan Aen, Jesper Arent, Jan Damsgaard, Jakob Iversen, Lars Mathiassen, and Peter Axel Nielsen participated actively in the project. In addition, Gro Bjerckness, Karl Heinz Kautz, and Ojelanki Ngwenyama were engaged and served as visiting professors in the IS unit during parts of the project. The project has strengthened the research unit's network of collaborations internationally, to the Danish software industry, and to other Danish academic institutions (Copenhagen Business School, Denmark's Technical University).

Some key research questions were addressed. How can one assess software practices? What are the central needs and opportunities for improvement? How can improvements be implemented to gain sustainable advantages? The project has resulted in documented improvements in the 4 software organisations and knowledge and experiences have been disseminated to the Danish software industry through publications, workshops and conferences. The research results are documented in a number of publications, including 11 journal papers, 30 conference papers, 1 book chapter, and 6 workshop contributions. In addition, two of the participants (Jakob Iversen & Jesper Arendt) have defended a PhD thesis based on results from the project.

The Staging of Virtual 3D-Spaces

This project is a large multidisciplinary effort that started in 1998 and continues until 2001. More than 30 researchers from different Danish research institutions participate in the project. It is supported by the Danish Research Ministry with a total budget of DKK 22 million.

The activities of the project are divided in three teams concerned with analysis, construction, and methodology related to the development of computerised interactive media that stage virtual inhabited 3D-spaces. Peter Bøgh Andersen, Mikael Skov, and Jan Stage represent the information system research unit in the construction and methodology teams.

The construction team has concentrated on developing architectures for agents in 3D-environments. The idea is that agents are the smallest building blocks in interactive narratives, and, in this capacity, should be designed to function as *actors*. It is assumed that interactive narratives are not told, but enacted. The first results of these activities consist of a few simple prototypes, and some theoretical contributions.

The methodology team has focussed on analytical and experimental approaches to software development. These approaches have been characterised and their relevance for design of interactive media has been evaluated through qualitative empirical experiments.

The first results of these activities are documented in L. Qvortrup et al. (ed.): *Interaction in/with Virtual Inhabited 3D Worlds*. Springer Publishers, London (2000). This book includes contributions from all three teams.

Process Integration and Transformation based on new Network Information Technology (PITNIT)

The PITNIT project is a large multidisciplinary research effort that involves researchers from Department of Computer Science, Department of Production, Department of International Business studies at Aalborg University, and Department of Management, Aarhus University. It started in 1999 and continues until 2003. It is supported by the Danish Research Council and has a total budget of DKK 7,5 Million. Participants from information systems are Jan Damsgaard (research director), Jan Karlsbjerg, Rens Scheepers, and Jan Kristensen.

The goal of the PITNIT project is to describe, analyse and offer practical guidelines of the integration and transformation of industrial processes that is enabled by new networked information technology. The key research challenges are

- The merger between a number of process innovation concepts and associated IT.
- The Extended Enterprise that emerges from a multitude of different co-operating organisations and associated IT.
- Network-based interaction with the environment using IT for marketing purposes.

Management, Organisation, and Competence Development Based on Network Information Technology (LOKNIT)

The LOKNIT project focuses collaboration between practitioners and researchers in connection with the PITNIT project. The project was initiated in 1999 and terminates in 2001. The total budget is DKK 4.7 million. Participants from information systems are Jan Damsgaard (research director), Jan Karlsbjerg, Rens Scheepers, and Jan Kristensen. The goal is to disseminate research results mainly from the PITNIT project and to interact with industry representatives. This is done through monthly meetings and annual conferences.

EDI Diffusion and e-commerce

This project examines the diffusion and implementation of EDI and e-commerce systems in Hong Kong, Finland, the Netherlands, and Denmark. One aim is to study the evolution of EDI patterns on an organisational, industry, and environmental level on a longitudinal basis using power dependency analysis and institutional theory. Another aim is to understand the evolution of e-commerce networks as they are subjected to the pressures of usage. One interesting aspect of e-commerce is the standardisation issues studied using emergence theory and

linguistics. Finally the project also studies the evolution of electronic markets and the integration of logistics and distribution channels, especially in the air cargo industry. The project incorporates researchers from University of Jyväskylä, University of Amsterdam, Georgia State University and Hong Kong University. The project is not funded by any external sources.

Interactive Media: Functionality, Aesthetics and Usability

The project was started in 1999 with Lars Bo Eriksen, Mikael Skov, and Jan Stage as project members. From the summer 2000, Lars Bo Eriksen has left the project, and Peter Bøgh Andersen, Jesper Kjeldskov, and Tom Nyvang have joined it. The project is funded by Siemens, Nokia, the Faculty of Technology and Science, and Department of Computer Science.

The first activities in the project have focussed on design and establishment of a modern usability laboratory. The building was finished in the summer 1999, and from the beginning of 2000 a first collection of equipment has been set up and tested on a number of student projects. A second collection of equipment is being installed during the fall of 2000.

The first experiments have dealt with mobile informatics with special focus on mobile telephones. A comparison between web-based and wap-based interfaces has been conducted.

Elastic systems

Elastic Systems is part of Centre for Human Machine Interaction funded by the Danish Basic Research Foundation. Started in 1998 the centre has now moved to the department. The project is funded until 2003. The objective of the centre is to conduct empirical and theoretical research in human-machine interaction in complex work domains. The Elastic Systems subproject is a co-operation between the Danish Maritime Institute, Department of Automation (DTU) and Aalborg University. It aims at designing a new generic type of interface and to apply it to instrumentation of ship bridges. The interface is called “elastic” because it supports smooth changes of the locus of control and of access to system properties. The design is based on extensive field studies aboard a variety of ship types, with special regard to changing needs of information and control.

8.3 Organization and Staff

8.3.1 Staff

The IS research unit has grown in the period from one to two full professors, and from three to four associate professors, thus a moderate growth rate.

Scientific Staff					
	1996	1997	1998	1999	2000
Full Professors	1	1	1	1	2
Associate Professors	3	3	3	3	4
Assistant Professors	1	1	1	1	1.5
Visiting Full Professors		0.25	0.25	1.25	0.3
Visiting Researchers					0.3

The visitors have been Gro Bjerknæs (Avenir), Karl Heinz Kautz (Copenhagen Business School), Kalle Lyytinen (University of Jyväskylä), Ojelanki K. Ngwenyama (Virginia Commonwealth University), Knut Rolland (University of Oslo), Duane P. Truex (Georgia State University) and Detlev Zwick (University of Rhode Island).

Peter Bøgh Andersen

Academic Degrees

- 1991 Dr. Phil. in Information science, University of Aarhus
- 1971 Ph.D. in Danish language, University of Aarhus.

Positions

- 2000- Professor, Department of Computer science, University of Aalborg
- 1993–99 Professor, Department of Information and Media Science, University of Aarhus
- 1981–93 Associate professor, Humanistic-Informatic Curriculum, University of Aarhus
- 1971–80 Associate professor, Department of Scandinavian Languages and Literature, University of Aarhus

Key Publications:

Andersen, P. B. Dynamic Logic. In: Kodikas/Code. *Ars Semeiotica*. An international Journal of Semiotics (18/4), 249-276. 1995. Also in: E. W. B. Hess-Lüttich & B. Schlieben-Lange (Eds): *Signs & Time - Zeit & Zeichen*, Gunter Narr Verlag, Tübingen (1998), 152- 186.

Andersen, P. B. Multimedia phase-spaces. In: *Multimedia Tools and Applications* 1998/6, 207-237.

Andersen, P. B. WWW as a self-organizing system. In: *Cybernetics & Human Knowing*. 1998, 5/2, 5 - 42.

Andersen, P. B. Genres as Self-Organising Systems. In: *Downward Causation. Minds, Bodies and Matter*. Eds, P. Bøgh Andersen, C. Emmeche, N. O. Finne- mann, & P. Voetmann Christiansen. Aarhus University Press: Aarhus, 2000. 214-260.

Jan Damsgaard

Academic Degrees

- 1996 Ph.D. and European Doctorate in Computer Science, Aalborg University
- 1993 Cand. scient. (M.Sc.) in Computer Science (major) and Psychology (minor), Aalborg University

Positions

- 2000– Associate professor, dept. of computer science, Aalborg University
- 2000– Visiting associate professor, Aarhus Business School
- 1999 Associate professor, dept. of management, Aarhus Universitet
- 1999– Visiting associate professor, Dept. of Informatics, Copenhagen Business School
- 1998– Docent in computer science and information systems, University of Jyväskylä, Finland
- 1998 Visiting researcher, dept. of management, School of Business, Hong Kong University
- 1997– Visiting professor, dept. of Informatics, University of Pretoria, South Africa
- 1996–98 Assistant professor, dept of computer science Aalborg University, Denmark

Key Publications:

Damsgaard, Jan and Lyytinen, Kalle (1998). Contours of Electronic Data Interchange in Finland: Overcoming technological barriers and collaborating to make it happen. *The Journal of Strategic Information Systems*. Volume 7, pp. 275-297.

Damsgaard, Jan and Scheepers, Rens (2000). Managing the Crises in Intranet Implementation: a Stage Model. *Information Systems Journal* 10(2)131-150.

Damsgaard, Jan and Duane Truex (2000). The Procrustean Bed of Standards. Binary relations and the limits of EDI standards. *European Journal of Information Systems* 9(3):173-188.

Bansler J. P., J. Damsgaard, E. Havn, J. Thommesen, and R. Scheepers (2000). Corporate Intranet Implementation: Managing emergent technologies and organizational practices. *Journal of AIS*, Vol 1. paper 10, 40 pages.

Lars Bo Eriksen

Academic Degrees

- 2000 Ph.D. in Computer Science, Aalborg University
- 1995 Cand. polyt. (M.Eng.) in Computer Science, Aalborg University

Positions

- 1999–00 Assistant Professor, dept. of computer science, Aalborg University
- 1996 Researcher Viktoria Institute, University of Gothenburg

Lars Mathiassen

Academic Degrees

- 1998 Dr. tech. in Computer Science, Aalborg University
- 1981 Dr. Scient. (Ph.D.) in Computer Science, University of Oslo
- 1975 Cand. scient. (M.Sc.) in Computer Science, Aarhus University

Positions

- 1997–00 Research Director, National Danish Centre for IT-research
- 1997– Visiting Professor, Viktoria Institute, Gothenburg, Sweden.
- 1998 Visiting Professor, Georgia State University, USA
- 1987– Full Professor, Computer Science, Aalborg University

Key Publications:

- L. Mathiassen, A. Munk-Madsen, P. A. Nielsen and J. Stage. *Object-Oriented Analysis & Design*. Marko Publishing House, 2000.
- L. Mathiassen. *Reflective Systems Development*. Dr. Techn. Thesis, Aalborg University, 1998.
- B. Dahlbom and L. Mathiassen. *The Future of Our Profession*. *Communications of the ACM*, Vol. 40, No. 6, 1997.
- K. Lyytinen, L. Mathiassen, and J. Ropponen. *Attention Shaping and Software Risk: A Categorical Analysis of Four Classical Approaches*. *Information Systems Research*, Vol. 9, No. 3, 1998.

Peter Axel Nielsen

Academic Degrees

- 1990 Ph.D. in Systems and Information Management, Lancaster University, UK.
- 1986 Cand. scient. (M.Sc.) in Computer Science, Aarhus University

Positions

- 1992– Associate Professor in Computer Science, Aalborg University

Key Publications:

- Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (2000) *Object-Oriented Analysis & Design*. Aalborg: Marko.

Avison, D., Lau, F., Myers, M., & Nielsen, P.A. (1999) Action Research. *Communications of the ACM*, 42, 1, 94-97.

Iversen, J., Nielsen, P. A., & Nørbjerg, J. (1999) Situated Assessment of Problems in Software Development. *The DATABASE for Advances in Information Systems*, 30, 2, 66-81.

Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1996) "Method Engineering; Who is the Customer?" Proceedings of the International Conference on Method Engineering, Method Engineering. Principles of Method Construction and Tool Support, 232-245.

Palle Nowack

Academic Degrees

- 2000 Ph.D. in Computer Science, University of Southern Denmark
- 1996 Cand. polyt (M.Eng.) in Computer Science, Aalborg University

Positions

- 2000– Assistant professor in Computer Science, Aalborg University
- 2000– Research Assistant professor in Software Engineering, University of Southern Denmark

Key Publications:

P. Nowack: Structures and Interactions - Characterizing Object-Oriented Software Architecture. Ph.D. Thesis, University of Southern Denmark, 2000.

E. E. Jacobsen & P. Nowack: Frameworks & Patterns: Architectural Abstractions. Chapter in Building Application Frameworks: Object Oriented Foundations of Framework Design (Mohamed E. Fayad, Douglas Schmidt, and Ralph Johnson, editors), Wiley & Sons, New York, NY, 1999: 29-54.

E. E. Jacobsen, B. B. Kristensen, P. Nowack, & T. Worm: Software Evolution: Prototypical Deltas. In Proceedings of International Conference on Technology of Object-Oriented Languages and Systems (TOOLS Asia,99), Nanjing, China, 1999. Computer Society Press, 1999: 14-30.

E. E. Jacobsen, B. B. Kristensen, & P. Nowack: Architecture = Abstractions over Software. In Proceedings of International Conference on Technology of Object-Oriented Languages and Systems (TOOLS Pacific,99), Melbourne, Australia, 1999. Computer Society Press, 1999: pp. 89-99.

Jan Stage

Academic Degrees

- 1989 Dr. Scient. (Ph.D.) in Computer Science, University of Oslo
- 1984 Cand. scient. (M.Sc.) in Computer Science, Aalborg University

Positions

- 1991– Associate Professor in Computer Science, Aalborg University

Key Publications:

- Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (2000) Object-Oriented Analysis & Design. Aalborg: Marko.
- Baskerville, R. L., & Stage, J. (1996) Controlling Prototype Development Through Risk Management. Management Information Systems Quarterly, Vol. 20, No. 4, 481-504.
- Eriksen, L. B., & Stage, J. (1998) A Qualitative Empirical Study of CASE Tool Support to Method Learning. Information and Software Technology, 40, 5-6, 339-345.
- Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1996) "Method Engineering; Who is the Customer?" Proceedings of the International Conference on Method Engineering, Method Engineering. Principles of Method Construction and Tool Support, 232-245.

Ivan Aaen

Academic Degrees

- 1990 Lic. scient. (Ph.D.) in Computer Science, Aalborg University
- 1980 Cand. polyt. (M.Eng.) in Systems Construction, Aalborg University

Positions

- 1991– Associate Professor in Computer Science, Aalborg University

8.4 Collaborations

The Information Systems research unit has maintained and developed relations to a broad group of institutions and individuals. This includes collaboration within the group and department as well as with international research communities and Danish IT organisations. Below, we describe our collaboration with partners outside the department.

There has been a long-term collaboration with Binghamton University and Georgia State University, USA. These activities have focussed on key research issues in information systems, and have involved Professors Richard Baskerville,

Ojelanki Ngwenyama, Duane Truex, and Heinz Klein. Lars Mathiassen stayed for one semester at Georgia State University.

The research unit has had a comparable collaboration with University of Gothenburg, Sweden. Lars Bo Eriksen stayed there for one semester and has continued collaboration with individuals. Lars Mathiassen has ongoing research collaboration with industrial researchers at the Viktoria Institute at University of Gothenburg (Pouya Pourkomeliyan, AstraZeneca, Kerstin Nielsson & Ingegerd Andersson, Volvo IT).

Jan Damsgaard has stayed at and collaborated with Hong Kong University, School of Business, University of Jyväskylä, Finland, University of Pretoria, South Africa, Danish Technical University, Aarhus University, Aarhus School of Business, IT University Copenhagen, and Copenhagen Business School.

Mikael Skov stayed for one semester at Swinburne University of Technology, Human-Computer Interaction Laboratory, where he visited Professor Penelope Sanderson. He has continued collaboration with Lorraine Johnston from the same department.

In addition, the following individuals have visited the research unit as part of collaborative efforts: Ojelanki Ngwenyama, Gro Bjercknes, Karl Heinz Kautz, Jeremy Rose, Helana Scheepers, Rens Scheepers, Kalle Lyytinen, and Duane Truex.

The individual members of the research unit have collaborated with a broad circle of IT organisations. From 1996 to 2000 research projects and consulting has been carried out together with Bankdata, Den Danske Bank, Kommunedata, Metodica, Nokia, Teledanmark, Brüel & Kjær, Ericsson, Systematic, Linq Systems, MindPass, Nykredit, and LEGO Group.

The LOKNIT network has over 75 members from a mix of private companies and public organisations. The members attend meetings and seminars on a regular basis. Kalle Lyytinen, Duane Truex, and Ojelanki Ngwenyama are all employed as visiting professors in the LOKNIT project.

8.5 Ph.D. Projects

8.5.1 Overview

During this 5-year period a total of 9 Ph.D.'s have been awarded within the IS research Unit, another 6 Ph.D. studies have been initiated and are in progress.

8.5.2 Project Descriptions

The Diffusion of Electronic Data Interchange: An Institutional and Organizational Analysis of Alternative Diffusion Patterns

Name: Jan Damsgaard

Education: M.Sc. in Computer Science and Psychology, 1993
Duration: August 1993–December 1996
Status: Degree awarded
Advisor: Ivan Aaen

This dissertation develops a framework to study the diffusion of complex networked technologies. An analysis of EDI literature and evaluation of general diffusion theory suggest a multi-layered and multi-theoretical approach to study EDI diffusion. The approach takes into account institutional, industry specific and organisational factors in the study of diffusion processes and thus extends the analysis beyond micro-level organisational accounts.

The thesis delivers a field study of EDI diffusion and a detailed analysis of 16 alternative EDI diffusion patterns. The field study altogether consists of 67 empirical case studies, which are used to delineate typical patterns in the diffusion, and to distill theoretical constructs that explain the evolution of these patterns. The study focuses on diffusion (or non-adoption) patterns in three socio-economic units; Hong Kong (longitudinal), Finland and Denmark. The data is collected in four different sectors; transportation, retail, finance, and paper and pulp.

Intranet Implementation: Influences, Challenges and Role Players

Name: Rens Scheepers
Education: MBA and Honors degree, University of Pretoria, South Africa, 1993
Duration: August 1996–November 1999
Status: Degree awarded
Advisor: Jan Damsgaard

The main focus of the thesis is to examine the organisational influences, challenges and role players associated with the implementation of intranet technology. Intranet technology is the organisational application of Internet technology, with a purely internal focus.

Intranet implementation is examined from a contextual perspective, a process-specific perspective and an individual perspective. From a contextual perspective, the research examines how the intranet implementation process influences and how the process is influenced in return by the organisational context. From a process-specific level, a conceptual model of the intranet implementation process is proposed and it is argued that there are three main existential crises which must be overcome for the intranet to become institutionalised in the organisation. From an individual perspective, key organisational role players in the intranet implementation process are identified and the roles they play during the process are examined.

Coordinating Cooperative Work: A Framework for the Design of Flexible Computer-Based Support

Name: Monica Divitini

Education: M.Sc. in Computer Science, University of Milan, Italy, 1991

Duration: December 1995–December 1999

Status: Degree awarded

Advisor: Jan Stage

The thesis is about the design of computer-based systems supporting in a flexible way the co-ordination of co-operative work. Co-operation among distributed actors requires an additional work for co-ordinating the individual activities. The intrinsic distributed and dynamic nature of co-operative work makes co-ordination complex. In the one hand, the distributed nature of work calls for the definition of explicit mechanisms (co-ordination mechanisms) for reducing the complexity of co-ordinating the co-operative activities. On the other hand, the dynamic nature of work contexts makes the definition and use of these mechanisms problematic.

The project has developed computer-based systems that can be used for augmenting the capabilities of the individuals and helping them in handling the complexity of co-ordination. It has been defined how these systems must be flexible, so to capture the different co-ordination mechanisms needed for supporting co-ordination, and make them adaptable to the changing environment. Requirements for computer-based support are identified and a framework for designing such systems is presented. The tools of the framework have been used in three different real-world cases.

Building Knowledge to Prevent Failure in IT Development and Use

Name: Helana Scheepers

Education: M. Com., University of Pretoria, South Africa, 1993

Duration: August 1996–January 2000

Status: Degree awarded

Advisor: Lars Mathiassen

The main focus of the thesis is on examining how failure in IT development and use can be prevented by building knowledge within some of the involved communities of actors. Two specific cases are explored: learning future systems developers about escalation in systems development, and building knowledge about IT implementation amongst disadvantaged groups of users starting to take IT into use.

Three research questions are explored. How can knowledge be built to proactively prevent failure in IT development and use? How can knowledge about escalation be built during the education of future developers to actively prevent escalation situations from taking place? How can disadvantaged communities build knowledge during IT implementation in order to proactively prevent the dissemination of inappropriate IT? These questions are researched through the two practice-related studies, one addressing the role of IT professionals, the other addressing the role of users. The approach taken to prevent failure is in both cases through knowledge building within the respective communities.

Matching Methods and Problems—Object-oriented Analysis Methods for Coordination Problems

Name: Birgitte Krogh

Education: M.Sc. in Computer Science, Aalborg University, 1993

Duration: August 1994–April 2000

Status: Degree awarded

Advisor: Lars Mathiassen

The main focus of the thesis is on examining how to rationally evaluate the fit between systems development methods and problems. Particular emphasis is given to the match between object-oriented analysis methods and the use of IT-systems to support human co-ordination. A number of techniques for object-oriented analysis are researched in relation to practical co-ordination problems. Specific lessons from these studies are used as a basis for understanding the matching of methods and problems in general.

Two research questions are explored. On a general level: How can the potential match of a system development method to a problem be assessed? On a more specific level: To which extent do object-oriented analysis methods match the requirements posed by co-ordination problems? The topic of matching methods, or methods features, with problems has played an important role within information systems and software engineering research. Key contributions are surveyed in the thesis and the approach taken to the present study is motivated on that basis. ‘Problem frame’ is chosen as the key concept and the study is performed based on laboratory experiments with specific method features (within object oriented analysis) and concrete system development problems (related to mechanisms for co-ordination of human activity).

Strategy and Measurement in Software Process Improvement

Name: Peter S. Bøttcher

Education: M.Sc. in Computer Science and Mathematics, Aalborg

Duration: April 1994–April 2000

Status: Degree awarded

Advisor: Ivan Aaen

The dissertation researches strategy and measurement in software process improvement (SPI) focusing on the following questions: How does the SPI strategy of the SEI Capability Maturity Model (CMM) relate to other improvement strategies? Are there important SPI concepts that are not covered by the CMM? What measurements are needed in SPI? What is required from measurement methods when used in SPI? To answer these questions a number of SPI approaches were investigated: The software factory, Total Quality Management, Software Six Sigma, Measurements based on product attributed, and Bayesian analysis. This research was conducted as an industrial Ph.D. study in collaboration between Aalborg University and Kommunedata. The research approach was primarily theoretical.

Data-Driven Intervention in Software Process Improvement

Name: Jakob Iversen

Education: M.Sc. in Software Engineering, Aalborg University, 1996

Duration: February 1997–June 2000

Status: Degree awarded

Advisor: Lars Mathiassen

Data collection and interpretation play important roles within software process improvement. State-of-the-art techniques are surveyed in this thesis and practices within four software organisations are researched. The thesis provides concepts for understanding data-driven intervention and it discusses these activities in the context of organisational learning. The main focus is on examining the role that data collection and interpretation play in SPI to support effective intervention. Different types of data-driven intervention are identified, approaches to practice data-driven intervention are presented, and problems and opportunities related to data-driven intervention in SPI are discussed.

Three research questions are explored: What types of data-driven intervention can be distinguished in SPI? How should data-driven intervention be performed in SPI? How can data-driven intervention supports SPI? The doctoral work consists of a significant body of empirical work. Jakob Iversen has participated in a national research project involving close collaboration with four systems development organisations and a number of researchers. The primary research approach is action research.

Designing Web News Services – Bridging Software and Domain Orientation

Name: Lars Bo Eriksen

Education: M.Sc. in Computer Science and Mathematics, Aalborg University, 1995

Duration: August 1996–June 2000

Status: Degree awarded

Advisor: Jan Stage

The thesis explores approaches for supporting design of web news services. The design of IT-solutions based on the web platform is a relatively unexplored character. The novelty and opportunities of the web has spawned a plethora of new uses. For designers, and those who intend to support designers, the new platform raises the question of what kind of support is needed in the task of designing a solution based on the web.

Software and domain orientation is introduced as two opposing approaches to support designers. These approaches are described and explored. The empirical foundation of the study consists of public web-sites and a detailed inquiry into a development project. It is argued that existing system design methods are useful for designers, and it is suggested that they should be combined with genre theory as a means to understand services in a domain such as news.

Normative Software Process Improvement

Name: Jesper Arent

Education: M.Sc. in Software Engineering, Aalborg University, 1996

Duration: February 1997–October 2000

Status: Degree awarded

Advisor: Ivan Aaen

A new set of ideas on how to improve quality and productivity within software engineering have developed over the last decade under the notion of software process improvement. One of the key ideas in SPI is to use norms to increase maturity of the software operation. The thesis presents findings of a study into the implementation of normative SPI in four Danish software organisations over a period of three years. In each of the companies, SPI research groups were established to support the companies' improvement efforts. Participation in these research groups gave first-hand insight in major issues in the implementation and management of normative SPI initiatives. The thesis offers a general overview of the characteristics of normative SPI. It provides a survey of state-of-the-art knowledge on SPI and positions SPI in the landscape of strategic thrusts that can be initiated to mature software organisations. In doing so it addresses the

following questions: (1) What are the characteristic features of SPI initiatives? (2) How do SPI initiatives compare to other improvement approaches? (3) What are the key benefits and risks related to SPI initiatives? The doctoral work consists of a significant body of empirical work. Jesper Arent has participated in a national research project involving close collaboration with four systems development organisations and a number of researchers. The primary research approach is action research.

Coordination at work

Name: Morten Nielsen

Education: M.A. in Humanistic Computer Science, Aalborg University, 1995

Duration: January 1996–December 2000 (expected)

Status: In progress

Affiliation: Centre for Human Computer Interaction, University of Aarhus

Advisor: Peter Axel Nielsen

The primary field of interest of the Ph.D. project is: Computer support of the co-ordination involved in distributed co-operative decision making performed under time- and safety critical work conditions. The project is empirical and the work domain being studied is navigation of large container vessels in confined waters.

Modern maritime transportation is characterised by a continuously growing number of still larger and faster vessels. Growing levels of traffic density and diversity coupled with high-speed sailing has created a significantly complex work domain where the time available for decision making is an increasingly scarce resource. Due to hazardous cargo, collisions and grounding carry the potential of environmental disasters. Thus, efficient and effective co-operation and co-ordination among the human agents in the navigational system must be maintained at all times to ensure safe and timely navigation.

The Ph.D. project seeks to contribute to a better understanding of time- and safety related work aspects involved in modern navigation with the objective to accommodate ways and means for a higher degree of safety in maritime operations. More specifically the project investigates how modern information technology can be employed to reduce the work complexity by providing enhanced computer support of co-ordination in co-operative work arrangements.

Notations for Dynamic 3D Multimedia System Development

Name: Mikael Skov

Education: M.Sc. in Computer Science and Mathematics, Aalborg University, 1995

Duration: August 1998–June 2001 (expected)

Status: In progress

Advisor: Jan Stage

This thesis explores the development of interactive narrative (multimedia) systems. Most interactive narratives involve and develop around some sort of a story or a plot. This key characteristic of interactive narratives poses new requirements and challenges to notations, concepts, and techniques used during the design process. In fact, many interactive narratives share many similarities with movies and film, however, the introduction of interactivity makes them more complex and difficult to construct.

The design process of interactive narratives is analysed from two different perspectives. First, the design process is a systems development process and traditional software engineering approaches and systems development concepts and notations are explored and evaluated. Second, the design process is viewed as a narrative production and construction and narrative concepts and notations are examined. The goal of the project is to develop a framework for the design of interactive narratives, which utilises and combines relevant concepts and notations from software engineering, systems development, and narratology.

E-Business Development: Going On-line With the Business

Name: Jan Kristensen

Education: M.Sc. in Information Science, Aarhus School of Business, 1998

Duration: September 2000–December 2001 (expected)

Status: In progress

Advisor: Jan Damsgaard

This project focuses on changes in the information system due to new workflows as result of the new way of doing business. The purpose is to develop a framework for understanding and managing this process. Jan Kristensen is enrolled at Århus Business School and is co-supervised by Jan Damsgaard.

New technology introduces a need to develop a new understanding for companies' routines or procedures, which can incorporate dynamics. When a company establish e-business it is more than just a lift in technological level. It causes a shift from a stepwise evolution to a much more fluent evolution which forces to create a more flexible organisation than hitherto. This Ph.D. project is based on empirical findings done with a case study in two Danish merchants companies. E-business implementation is examined from a contextual, a process and an individual perspective. The study examines how wholesale dealers implement e-business solutions and which consequences this have for organisational design and technological understanding in the company.

Managing Information Technology Infrastructures

Name: Jan Karlsbjerg

Education: M.Sc. in Computer Science, Aalborg University, 1999

Duration: August 1999 – July 2002 (expected)

Status: In progress

Advisor: Jan Damsgaard

This research project aims to contribute to the current scientific knowledge about the use and management of information technology infrastructures. An organisation's information technology infrastructure is the combination of all the interconnected information systems in the organisation, including hardware and software components as well as the rules and standards of use of the systems.

Several medium-sized to large organisations and their respective information technology infrastructures (ITIs) are studied. Through a detailed understanding of the individual cases and through inter-case comparisons the project examines processes that lie behind the development or evolution of ITIs, i.e. how components of the systems are selected, changed, updated, and replaced.

Design and Usability of Embedded Mobile Interfaces

Name: Jesper Kjeldskov

Education: M.A. in Humanistic Computer Science, Aalborg University, 2000

Duration: October 2000–August 2003 (expected)

Status: In progress

Advisor: Jan Stage

The purpose of this Ph.D. project is to inquire into design, construction, and usability of new interfaces for mobile devices. The focus is on possibilities for embedding the interface in the user's context by means of augmented reality (AR) technologies.

The project aims to develop a mobile embedded interface that is commercially relevant and can be implemented in collaboration with one of the mobile telephone producers in the region. The problems characterising today's interfaces on different categories of mobile devices will be described systematically and the potentials of mobile AR interfaces for solving these problems will be explored.

The main activities will involve development, implementation, and usability evaluation of prototypes that realise ideas for new user interfaces. This includes prototypes based on existing technologies as well as visualisation of future technologies by means of advanced virtual reality facilities at VR Center Nord.

Improving MIS in Software Organisations

Name: Helle Damborg Frederiksen

Education: M. Sc. in Computer Science, Aalborg University, 1991

Duration: August 2000—February 2004 (expected)

Status: In progress

Advisor: Lars Mathiassen

This Ph.D. project addresses the management of software organisations based on systematic collection of data, i.e. metrics programs and, in particular, issues of how to improve data collection, interpretation, distribution and usage. The project is carried out as an action research effort in close collaboration with Kommunedata. This organisation has implemented a number of metrics programs. The goal of the study is to analyse strengths and weaknesses of current management practices. Based on that the study will design and enact a number of interventions. These interventions should improve the quality and reliability of the data used, the effectiveness and efficiency of the data collection activities, and, last but not least, the relevance of the data as seen from the point of view in supporting management decisions in the organisation.

8.6 Services

The IS Unit is actively seeking high visibility by engaging all its members in key services within the international research community. It is important for each individual to be engaged, to contribute, and to be recognised by international peers. It is at the same time important that the unit in this way is acknowledged as a key player within IS research.

8.6.1 Editorial Boards

- *European Journal of Information Systems* (L. Mathiassen),
- *Information and organization* (J. Damsgaard),
- *Information, Technology, and People* (L. Mathiassen),
- *Journal of Database Management* (P. A. Nielsen, J. Stage),
- *Scandinavian Journal of Information System* (P. A. Nielsen).

8.6.2 Keynote Speaker

- IRIS 2000: The Twenty-third Information Systems Research Seminar in Scandinavia, Uddevalla, Sweden, 2000 (L. Mathiassen).
- ISD 2000: International Conference on Systems Development, Kristiansand, Norway, 2000 (L. Mathiassen).
- BIT 2000: Conference on Business Information Technology Management, Manchester, UK, 2000 (L. Mathiassen).

8.6.3 Scandinavian Research Community

The research unit has played a major role in further strengthening the Scandinavian Research Community within Information Systems. Peter Axel Nielsen has been editor-in-chief and in charge of production of the Scandinavian Journal of Information Systems from 1993 to 1999, and Ivan Aaen has been in charge of the administration of the journal. The IRIS organisation (Information Systems Research In Scandinavia) was formally established in 1997 and Jan Damsgaard has been a member of the managerial board from 1997 to 2000 and played an active role in shaping the organisation.

8.6.4 IFIP Working Group 8.2

The research unit has a broad long-term involvement in IFIP Working Group 8.2 on Organisational and social perspectives on information systems. This is one of the largest and most active working groups in IFIP with more than 100 members. In June 2000 the research unit hosted a working conference for the group. Jan Stage served as program co-chair (with Richard Baskerville), and Peter Axel Nielsen served as organising chair. The program included invited articles from several of the main characters in the group as well as submitted and refereed articles. More than 100 international colleagues participated in the conference.

8.6.5 Program Committees

- ECIS 96: 4th European Conference on Information Systems, Lisbon, Portugal, 1996 (L. Mathiassen).
- IFIP WG 8.1/8.2 Working Conference on Principles of Method Construction and Tool Support, Atlanta, USA, 1996 (L. Mathiassen).
- International Conference on Object Oriented Information Systems, London, England, 1996 (L. Mathiassen).

- IFIP WG 8.6: Diffusion, Transfer and Implementation of Information Technology, Ambleside, Lake Windermere, UK, 1997 (J. Damsgaard, L. Mathiassen).
- ICIS 97: International Conference on Information Systems, Atlanta, USA, 1997 (L. Mathiassen).
- 21st International Research Seminar in Scandinavia. Sæby, Denmark, 1998 (J. Damsgaard).
- ICIS 98: International Conference on Information Systems, Helsinki, Finland, 1998 (I. Aaen).
- ECIS 99: 7th European Conference on Information Systems, Copenhagen, Denmark, 1999 (P. A. Nielsen).
- The second EGS seminar, Charlotte, North Carolina, 1999 (J. Damsgaard, Co-program chair).
- IFIP WG 8.2: New Information Technologies in Organizational Processes, St. Louis, Missouri, USA, 1999 (I. Aaen).
- IFIP WG 8.2: Organizational and Social Perspectives on Information Technology, Aalborg, Denmark, 2000 (I. Aaen, J. Damsgaard, L. Mathiassen).
- Third EGS seminar, Brisbane, Australia, 2000 (J. Damsgaard, Co-program chair).
- Sixth International Workshop on Requirements engineering: Foundation for Software Quality, Stockholm, 2000 (P. A. Nielsen).
- Fifth CAiSE/IFIP8.1 International Workshop EMMSAD'00 on Evaluation of Modeling Methods in Systems Analysis and Design, 2000 (P. A. Nielsen, J. Stage).

8.6.6 Organising Committees

- IRIS 21: Information Systems Research in Collaboration with Industry, Sæby, Denmark, 1998 (P. A. Nielsen).
- IFIP WG 8.2: Organizational and Social Perspectives on Information Technology, Aalborg, Denmark, 2000 (P. A. Nielsen).

8.6.7 Reviewing

All members of the group serve on a regular basis as reviewers for international conferences and journals including:

- *Communications of AIS*,
- *Communications of the ACM*,
- *European Journal of Information Systems*,
- *Hawaii International Conference on System Sciences*,
- *Human Factors in Computing Systems (CHI)*,
- *Information and Organization*,
- *Information Systems Research*,
- *Information, Technology, and People*,
- *International Conference on Information Systems (ICIS)*,
- *Journal of Database Management*,
- *Journal of Global Information Management*,
- *Journal of Information Systems*,
- *Journal of the Association of Information Systems*,
- *MIS Quarterly*,
- *Scandinavian Journal of Information Systems*,
- *SOFTWARE Practice and Experience*,
- *The DATABASE for Advances in Information Systems*.

All senior members of the group have also served on a number of national and international evaluation committees for Ph.D. degrees and positions as assistant, associate, and full professor. Finally, several members of the unit are invited on a regular basis to present their research at foreign universities and in industrial settings.

The lists of services given above are not all encompassing, but they illustrate the research unit's level of engagement in the international research community.

8.7 Own Evaluation

In the previous research evaluation five years ago the research unit set a number of goals. A main goal was to cover the four topics: (i) IT usage and management, (ii) professionalisation, (iii) strategies for systems development, and (iv) models and design. It was a goal to balance the focus on the discipline of software engineering and the discipline of information systems. We also set the goal to publish better, not necessarily more.

Compared to these goals the strengths comprise to the following:

- We have met the goals both in terms of covering the four topics and in maintaining a dual, balanced focus on both software engineering and information systems. During the five year period there has been an expansion of the group and its volume of published research. We have thus also experienced a broadening of the topics covered by the unit. A new effort in multimedia and human-computer interaction has been launched and is now properly staffed and attracts Ph.D. students.
- We have increased our publishing in international journals and conferences thus meeting higher standards than just five years ago. We find our publication record to have sufficient quality, quantity, and variation. In that sense we have met our goal; but our publication record can improve.
- We have maintained a strong position in the Nordic research community, and we have improved our international position and visibility. We find it a strengths that the unit has a high visibility and that cannot be attributed to a single person only. We have shown a large commitment to research services in the international information systems research community. More members of the unit than five years ago are now visible and have now substantial research services.
- We have maintained a strong empirical approach in all activities. As an integrated part of this we have built substantial research relations to IT-developing and -using companies.
- We have succeeded in getting external funding for a significant part of research activities. Projects have been influential on the day-to-day organisation of research activities.
- Nine Ph.D. students have successfully completed their degrees. Some Ph.D. students have been attracted from the international research community. All senior members of the unit have conducted supervision.

The weaknesses comprise to:

- Only a few of the Ph.D. students completed their degree within the three year limit for a Ph.D. study. None of the Ph.D. students have after their degree wanted to take the available assistant professor positions in the unit. Some have gone to better offers at universities outside Denmark; while some have gone to the local IT industry.
- We have experienced a decrease in the experimental, design-oriented research. Priority has been given to building research relations with the IT industry. Less priority has been given to focus on the new media, new applications, and new methodologies, e.g., mobile devices, internet applications, open source development, extreme programming.
- The strong project organisation has built ties between members of the same projects, but has also weakened ties between members in different projects. The unit has over the last five years become less cohesive. There is a need to create more research organisation across projects.

8.8 Research Plan 2001–2005

It is important that the Information Systems Unit continues to prioritise the variety of possible impacts. Overall, the unit will consolidate and further develop its research within *Information Technology Engineering & Management* (ITEM) over the next five years. This overall paradigm has served well as a guiding framework for the unit's research during the past period. The paradigm expresses a commitment to participate in professionalising the IT field, covering systems development and design, systems maintenance and use, information technology management, innovation, and utilisation in change processes. Intellectual support will be sought in other disciplines such as management, systems, sociology, media science, knowledge production, and technology in general.

The unit will continue its policy to orient itself towards IS (i.e. business school research communities) as well as software engineering (i.e. computer science and engineering research communities). This complementary approach to research IT development and use creates a strong basis for exploring practical problems in powerful ways that are not limited by arbitrary institutional traditions. More specifically, the unit plans to emphasise three particular research areas within this broader paradigm:

- *Business IS & Networked IT*: diffusion and adoption of inter-organisational systems; the evolution of organisational IT infrastructure of network and mobile technologies; design and implementation of *e-services* and *m-services* based on standards.
- *Multimedia, Mobility & Usability*: use of multimedia IT-systems to support process control and co-ordination; multimedia IT-systems based on

distributed and mobile technologies; approaches to usability of contemporary IT-systems.

- *Software Excellence*: software engineering with emphasis on design and architecture; supportive environments for software engineering; knowledge management in software organisations.

Software Excellence will continue and further develop the achievements of the unit over the past decade. *Business IS & Networked IT* will consolidate and further develop research that was initiated during the last five years and at the same time help building a broader profile oriented towards emerging information technologies and applications. *Multimedia, Mobility & Usability* will make it possible to merge the units engagement in the universities multimedia activities with the opportunities provided by the recent inclusion of professor Peter Bøgh Andersen into the unit. The two latter research areas will also help the institute strengthen its knowledge base for further developing its new education in informatics.

The past five-year period has broadened the groups profile and increased its focus on contemporary IT. The three research areas point at different areas in which the unit will seek to optimise its impact internationally over the next five years. At the same time, the unit will as a whole increasingly focus on emerging issues related to network and mobile technologies, hence creating increased opportunities for collaboration and synergy across the three themes. The increased diversity will hence be consolidated over the next five-year period while at the same time strengthening collaboration within the unit and across projects. The unit will continue to strengthen collaboration across projects, initiatives and groups. In particular, 2-day research workshops will be adopted as an annual activity and other means for strengthening internal collaboration will be explored.

The unit will continue to use empirical research as the main approach, and also continue its efforts to collaborate closely with practitioners and industrial partners. At the same time, more emphasis will be put on design of IT-based solutions and on researching useful approaches based on experimental research strategies. The unit will continue its policy to seek impact both within the relevant research communities through publication in international journals and at recognised international conferences, and in industry through publication of academic books for practitioners.

The unit will continue to develop its active international collaboration policy. Many guest will be invited both on short-term and long-term contracts. In addition, the unit will actively support all its members in organising sabbaticals on a regular basis to become members of internationally recognised research groups abroad for one or two semesters. This exchange is considered a key strategy for continued improvement of the units research capabilities. It is, at the same time,

considered a key factor in attracting more young researchers to become members of the unit.

The mixture between basic and ad-hoc research funding over the past period is considered appropriate, and the unit will continue to seek ad-hoc funding for specific projects on a similar level over the coming five year period.

Finally, the unit expects to see more Dr. Techn. or Dr. Scient. theses activities in the five year period to come.

8.9 Bibliography

Academic Books

Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1997) *Objektorienteret analyse og design (in Danish)*. Aalborg: Marko.

Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1998) *Objektorienterad analys och design (in Swedish)*. Lund: Studentlitteratur.

Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (2000) *Object-Oriented Analysis & Design*. Aalborg: Marko.

Articles in Refereed Journals

Avison, D., Lau, F., Myers, M., & Nielsen, P.A. (1999) Action Research. *Communications of the ACM*, 42, 1, 94-97.

Bansler, J. P., Damsgaard, J., Havn, E., Thommesen, J., & Scheepers, R. (2000) Corporate Intranet Implementation: Managing emergent technologies and organizational practices. *Journal of Association of Information Systems (JAIS)*.

Baskerville, R. L. & Stage, J. (1996) Controlling Prototype Development Through Risk Management. *Management Information Systems Quarterly*, Vol. 20, No. 4, 481-504.

Dahlbom, B., & Mathiassen, L. (1997) "The Future of Our Profession." *Communications of the ACM*, 40, No. 6.

Damsgaard, J. (1999) Electronic Markets in Hong Kong's Air Cargo Community. *Electronic Markets*, 8, 3, 46-49.

Damsgaard, J. (1999) Global Logistics System Asia Co., Ltd. *Journal of Information Technology*, 14, 3, 303-314.

Damsgaard, J., & Truex, D. P. (2000) The Procrustean Bed of Standards. Binary relations and the limits of EDI standards. *European Journal of Information Systems* 9(3):173-188.

Damsgaard, J., & Farhoomand, A. F. (1999) Electronic Commerce in Hong Kong

- Special Administrative Region of the People's Republic of China. *Electronic Markets*, **9**, 2, 73-80.
- Damsgaard, J., & Lyytinen, K. (1996) Government Strategies to Promote the diffusion of Electronic Data Interchange: What we know and what we don't know. *Information Infrastructure and Policy*, **5**, 3, 169-190.
- Damsgaard, J., & Lyytinen, K. (1998) Contours of Electronic Data Interchange in Finland: Overcoming technological barriers and collaborating to make it happen. *The Journal of Strategic Information Systems*, **7**, 275-297.
- Damsgaard, J., & Lyytinen, K. (2000) Building Electronic Trading Infrastructure: A private or public responsibility. *Organizational Computing and Electronic Commerce*.
- Damsgaard, J., & Lyytinen, K. (2000) The role of intermediating institutions in diffusion of Electronic Data Interchange (EDI): How Industry associations in the grocery sector intervened in Hong Kong, Finland, and Denmark. *The Information Society*.
- Damsgaard, J., & Scheepers, R. (1999) Power, Influence and Intranet Implementation: A Safari of South African Organizations. *Information, Technology and People*, **12**, 4, 333-358.
- Damsgaard, J., & Scheepers, R. (2000) Managing the Crises in Intranet Implementation: a Stage Model. *Information Systems Journal*, **10**, 2, 131-150.
- Divitini, M., & Simone, C. (2000) Supporting Different Dimensions of Adaptability in Workflow Modeling. *Computer Supported Cooperative Work*, **9**, 3/4, 365-397.
- Eriksen, L. B., & Stage, J. (1998) A Qualitative Empirical Study of CASE Tool Support to Method Learning. *Information and Software Technology*, **40**, 5-6, 339-345.
- Iversen, J., Nielsen, P. A., & Nørbjerg, J. (1999) Situated Assessment of Problems in Software Development. *The DATABASE for Advances in Information Systems*, **30**, 2, 66-81.
- Lyytinen, K., Mathiassen, L., & Ropponen, J. (1996) A Behavioral Model of Software Risk Management. *Scandinavian Journal of Information Systems*, **8**, 2.
- Lyytinen, K., Mathiassen, L., & Ropponen, J. (1998) Attention Shaping and Software Risk: A Categorical Analysis of Four Classical Approaches. *Information Systems Research*, **9**, 3.
- Mathiassen, L., & Sørensen, C. (1996) The Capability Maturity Model and CASE. *Journal of Information Systems*, **6**
- Mathiassen, L., & Stage, J. (1999) Informatics as a Multi-Disciplinary Education. *Scandinavian Journal of Information Systems*, **11**, 1, 13-22.
- Mathiassen, L., & Nielsen, P. A. (2000) Interaction and Transformation in SSM. *Systems Research and Behavioral Science*, **17**, 3.
- Mathiassen, L. (1996) Information Systems Development: Reflections on a Discipline. *Accounting, Management and Information Technology*, **6**, 1/2.
- Mathiassen, L. (1998) Reflective Systems Development. *Scandinavian Journal of*

Information Systems, **10**, 1.

Mathiassen, L., Borum, F., & Pedersen, J. S. (1999) Developing Managerial Skills in IT Organizations—A Case Study based on Action Learning. *Journal of Strategic Information Systems*, **8**.

Nulden, U. & Scheepers, H. (1998) What about Failures? Suggestions for an alternative computing education activity. *Information Technology Management*, **2**, 1, 133-144.

Nulden, U. & Scheepers, H. (1999) Interactive Multimedia and Problem Based Learning: Experiencing project failure. *Journal of Educational Multimedia and Hypermedia*, **8**, 2, 189-215.

Palmer, J., & Eriksen, L. B. (1999) Digital News - Paper, Broadcast and more convergence on the Internet. *International Journal on Media Management*, **1**, 1, 31-34.

Palmer, J., & Eriksen, L. B. (1999) Digital Newspapers explore marketing on the Internet. *Communications of the ACM*, **42**, 9, 32-40.

H. Scheepers & Mathiassen, L. (2000) Out of Scandinavia—Facing social risks in IT development in South Africa. *Journal of Global Information Management*, **8**, 2, 38-49.

Scheepers, H & Nuldén, U. (2000). Experience and interaction in teaching escalation. *Interactive Learning Environments*. **8**, 1, 1-22.

Dissertations

Arent, J. (2000) “Normative Software Process Improvement.” Ph.D. Thesis. Aalborg University: Department of Computer Science.

Damsgaard, J. (1996) “The Diffusion of Electronic Data Interchange: An Institutional and Organizational Analysis of Alternative Diffusion Patterns.” Ph.D. Thesis. Aalborg University: Department of Computer Science.

Divitini, M. (1999) “Coordinating Cooperative Work: A Framework for the Design of Flexible Computer-Based Support.” Ph. D. Thesis. Aalborg University: Department of Computer Science.

Eriksen, L. B. (2000) “Designing Web News Services – Bridging Software and Domain Orientation.” Ph. D. Thesis. Aalborg University: Department of Computer Science.

Iversen, J. (2000) “Data-Driven Intervention in Software process Improvement.” Ph. D. Thesis. Aalborg University: Department of Computer Science.

Krogh, B. (2000) “Matching Methods and Problems—Object-oriented Analysis Problems for Coordination Problems.” Ph. D. Thesis. Aalborg University: Department of Computer Science.

Mathiassen, L. (1998) “Reflective Systems Development.” Dr. Techn. Thesis, Aalborg University: Department of Computer Science.

Scheepers, H. (2000) “Building Knowledge to Prevent failure in IT Development

and Use.” Ph. D. Thesis. Aalborg University: Department of Computer Science.
Scheepers, R. (1999) “Intranet Implementation: Influences, Challenges and Role Players.” Ph. D. Thesis. Aalborg University: Department of Computer Science.

Refereed Conference Articles

- Aaen, I., & Damsgaard, J. (1999) “Software process Improvement: What management tends to forget.” *IFIP 8.2 and 8.6 Joint Working Conference on Information Systems: Current Issues and Future Changes*, Helsinki, Finland, 89-110.
- Aaen, I., Bøttcher, P., & Mathiassen, L. (1996) “Delta Measurement in Software Process Improvement based on Product Attributes.” *Proceedings of the 4th European Conference on Information Systems*, Lisbon, July 2-4, 15-28.
- Aaen, I., Bøttcher, P., & Mathiassen, L. (1998) “The Software Factory: Contributions and Illusions.” *Proceedings of European Conference on Information Systems, Aix en Provence*, France, June 4-6, 736-750.
- Andersen, E. S., Damsgaard, J., Hanseth, O., King, J. L., Markus, M. L., & Monteiro, E. (2000) “Standardization, Network Economics and IT.” *IFIP 8.2. Conference, Organizational and Social Perspectives on Information Technology*, Aalborg, Denmark, 521-525.
- Andersen, P. B., Nielsen, M., & Carstensen, P. (2000). Dimensions of Coordination. *Proc. of the Fifth International Workshop on the Language-Action Perspective on Communication Modelling*, 41-61. Fachgruppe Informatik der RWTH: Aachen.
- Arent, J., Iversen, J., Andersen, C. V., & Bang, S. (2000) “Project Assessments: Supporting Commitment, Participation, and Learning in Software Process Improvement.” *Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS)*, Maui, CD-ROM.
- Bansler, J. P., Damsgaard, J., Havn, E., Thommesen, J., & Scheepers, R. (1999) “Corporate Intranet Implementation: Managing emergent technologies and organizational practices.” *Proceedings of the 7th European Conference in Information Systems*, Copenhagen, Denmark, 750-757.
- Baskerville, R. L., & Stage, J. (2000) “Discourses on the Interaction of Information Systems, Organizations, and Society: Reformation and Transformation.” *IFIP 8.2. Conference, Organizational and Social Perspectives on Information Technology*, Aalborg, Denmark, 1-12.
- Bjerknes, G., & Mathiassen, L. (2000) “Improving the Customer-Supplier Relation in IT Development.” *Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS)*, Maui, CD-ROM.
- Carstensen, P., Krogh, B. & Sørensen, C. (1996) “Object Oriented Modelling of Coordination Mechanisms.” *Second International Conference on the Design of Cooperative Systems*, Juan-les-Pins, France, 12-14 June 1996, 223-240..
- Christiaanse, E., & Damsgaard, J. (2000) “A Tale of Two Airports: A Compar-

- ison of Electronic Infrastructures in the Air Cargo Industry in the Netherlands and Hong Kong SAR.” *Fourth Pacific Asia Conference on Information Systems, Electronic Commerce and Web-based Information Systems*, Hong Kong SAR, 512-523.
- Damsgaard, J. (1998) “Global Logistics System Asia Co.” *Nineteenth International Conference on Information Systems. Crossing Boundaries: Managing Virtual Enterprises*, Helsinki, Finland, 402-404.
- Damsgaard, J. (1997) “The Danish EDI Bandwagon Gaining Momentum.” *Thirtieth Annual Hawaii International Conference on Systems Sciences*, Maui, 31-41.
- Damsgaard, J., & Farhoomand, A. F. (1999) “Building Electronic Commerce Infrastructure: The Hong Kong Retail Sector.” *Proceedings of the 7th European Conference in Information Systems*, Copenhagen, Denmark, 850-865.
- Damsgaard, J., & Lyytinen, K. (2000) “The evolution of factors explaining EDI diffusion in Hong Kong in the late 1990’s.” *Fourth Pacific Asia Conference on Information Systems, Electronic Commerce and Web-based Information Systems*, Hong Kong SAR, 1061-1074.
- Damsgaard, J., & Lyytinen, K. (1999) “International trade at the speed of light: Building an electronic trading infrastructure in Denmark, Finland, and Hong Kong.” *IFIP 8.2 and 8.6 Joint Working Conference on Information Systems: Current Issues and Future Changes*, Helsinki, Finland, 417-438.
- Damsgaard, J., & Lyytinen, K. (1996) “Tracks of the Finnish EDI bandwagon.” *Proceedings of the 4th European Conference on Information Systems*, Lisbon, Portugal, 471-492.
- Damsgaard, J., & Lyytinen, K. (1997) “Hong Kong’s EDI Bandwagon. Derailed or on the right track?” *IFIP 8.6. Conference on Facilitating Technology Transfer Through Partnership: Learning from Practice and Research*, Ambleside, United Kingdom, 39-63.
- Damsgaard, J., & Scheepers, R. (1998) “Intranet Implementation in Large organizations.” *Proceedings of the 21st Information Systems Research Seminar in Scandinavia, Information Systems Research in Collaboration with Industry*, Sæby, Denmark, 165-184.
- Damsgaard, J., & Scheepers, R. (1999) “A stage model of intranet technology implementation and management.” *Proceedings of the 7th European Conference in Information Systems*, Copenhagen, Denmark, 100-116.
- Divitini, M. (1999) “Supporting coordination: a multi-agent approach.” *Proceedings of Norwegian Informatics conference NIK’99*, Norway.
- Divitini, M., & Farshchian, B. A. (1999) “Collaboration and Coordination through Basic Internet Tools: a case study.” *Proceedings of World Multiconference SCI ISAS’99*, vol. II, 550-554.
- Divitini, M., Simone, C., & Schmidt, K. (1996) “ABACO: Coordination Mechanisms in a Multi-agent Perspective.” *Proceedings of the 2nd International Conference on the Design of Cooperative Systems*, 103-112.
- Eriksen, L. B. (1997) “Digital Documents, Work and Technology – Three cases of

- Internet News Publishing.” *Proceedings of 30th Hawaii International Conference on System Science*, IEEE Computer Society Press. CD-ROM.
- Eriksen, L. B. (2000) “Limitations and Opportunities of SDM in WIS development.” *Proceedings of the IFIP TC 8 WG 8.2 Working Conference on Organizational and Social Perspectives on Information Technology*, Aalborg, Denmark, 473-486.
- Eriksen, L. B., Axline, S., & Markus, M. L. (1999) “What Happens After “Going Live” With ERP Systems? Competence Centers Can Support Effective Institutionalization.” *Proceedings of AMCIS 99*, IEEE CD-ROM.
- Eriksen, L. B., & Ihlström C. (1999) “In the path of the pioneers – Longitudinal Study of Web News Genre.” *Proceedings of the 22nd Information Systems Research Seminar in Scandinavia, Information Systems Research in Collaboration with Industry*.
- Eriksen, L. B., & Ihlström, C. (2000) “Evolution of the Web News Genre – The slow move beyond the print metaphor.” *Proceedings of 33rd Hawaii International Conference on Systems Science*. IEEE Computer Society Press. CD-ROM.
- Eriksen, L. B., & Skov, M. (1998) “A Critical Look at OOA&D in Multimedia Systems Development.” *Proceedings of the 21st Information Systems Research Seminar in Scandinavia, Information Systems Research in Collaboration with Industry*, Søby, Denmark, 185-195.
- Eriksen, L. B., & Sørsgaard, P. (1996). “Organisational implementation of WWW in Scandinavian Newspapers: Tradition based approaches dominate.” *Proceedings of the 19th Information Systems Research Seminar in Scandinavia*, 333-349.
- Eriksen, L. B., & Stage, J. (1996) “Applying CASE Tools to Support Method Learning.” *Proceedings of the 19th Information Systems Research Seminar in Scandinavia*, 843-857.
- Eriksen, L. B., & Stage, J. (1997) “Do CASE Tools Support Method Learning?” *Proceedings of the fifth International Conference of the British Computer Society: Training and Education of Methodology Practitioners and Researchers*, 159-169.
- Herskind, S., & Nielsen, P. A. (1998) “Coordination Mechanisms and Ephemeral Organizations.” *Proceedings of the Third International Conference on the Design of Cooperative Systems*, Cannes.
- Iversen, J. H., & Mathiassen, L. (2000) “Lessons from Implementing a Software Metrics Program.” *Proceedings of the 33rd Hawaii International Conference on System Sciences (HICSS)*, Maui, CD-ROM.
- Iversen, J. H., Mathiassen, L., & Nielsen, P. A. (1999) “Managing Risks in Software Process Improvement.” *Proceedings of the 7th European Conference on Information Systems (ECIS 99)*, Copenhagen, Denmark, 451-466.
- Iversen, J., Johansen, J., Nielsen, P. A., & Pries-Heje, J. (1998) “Combining Quantitative and Qualitative Assessment Methods in Software Process Improvement.” *Proceedings of the European Conference on Information Systems (ECIS 98)*, Aix-en-Provence, France, 451-466.
- J. H. Iversen, Nielsen, P. A., & Nørbjerg, J. (1998) “Problem Diagnosis in Soft-

- ware Process Improvement. Information Systems: Current Issues and Future Changes.” *Proceedings of IFIP WG 8.2/8.6 Joint Conference*, Helsinki.
- Johansen, J., & Mathiassen, L. (1998) “Lessons learned in a National SPI Effort.” *Proceedings of EuroSPI '98*, Gothenburg, Sweden, November 16-18, 5.1-17.
- Karlsbjerg, J., Olsen, C., & J. Damsgaard (2000) “Make or buy - A taxonomy of intranet implementation strategies.” *Proceedings of the 23rd Information Systems Research Seminar In Scandinavia*, Uddevalla, Sweden, 1573 - 1588.
- Krogh, B., & Mathiassen, L. (1998) “Coordination Problems and Method Features.” *Proceedings of the 31st Hawaii International Conference on System Sciences (HICSS)*, Maui, Hawaii.
- Krogh, B. (1996) “Object Oriented Analysis of Groupware Applications.” *Proceedings of the fifth International Conference on Information Systems Development*, Gdansk, Poland, 24-26 September 1996, 323-338.
- Krogh, B. (1997) “Modelling the Dynamics of Cooperative Work Arrangements.” *Proceedings of the sixth International Conference on Information Systems Development*, Boise, Idaho, USA, 11-14 August 1997, 61-73.
- Krogh, B., Dutoit, A., Levy, S., & Subrahmanian, E. (1996) “Strictly Class-Based Modelling Considered Harmful.” *Proceedings of the 29th Hawaii International Conference on System Sciences (HICSS)*, Maui, Hawaii.
- Mathiassen, L., & Nielsen, P. A. (1996) “Interaction and Transformation in Soft Systems Methodology.” *Proceedings of European Conference on Information Systems*, Lisbon, Portugal.
- Mathiassen, L. & C. Sørensen (1997): “A Guide to Manage New Software Engineering Tools”. In *2nd IFIP WG8.6 Working Conference: Diffusion, Transfer and Implementation of Information Technology*, 25th - 27th June 1997: Low Wood Hotel, Ambleside, Lake Windermere, United Kingdom, (Eds.) Tom McMaster, Enid Mumford, E. Burton Swanson, Brian Warboys, and David Wastell. 257-272. Chapman & Hall.
- Mathiassen, L. (2000) “Collaborative Practice Research.” *Proceedings of the IFIP TC 8 WG 8.2 Working Conference on Organizational and Social Perspectives on Information Technology*, Aalborg, Denmark.
- Mathiassen, L., Borum, F., & Pedersen, J. S. (1997) “Transforming IT Management through Action Learning.” *Proceedings of the Twentieth Information Systems Research Seminar in Scandinavia*.
- Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1996) “Method Engineering; Who is the Customer?” *Proceedings of the International Conference on Method Engineering, Method Engineering. Principles of Method Construction and Tool Support*, 232-245.
- Mathiassen, L., Munk-Madsen, A., Nielsen, P. A., & Stage, J. (1996) “Method Engineering: Who’s the Customer?” *Proceedings of the 19th Information Systems Research Seminar in Scandinavia*, Sweden, 693-705.
- Nielsen, P. A. (1997) “The Usefulness of Reinventing the Wheel: Implications for Learning Information Systems Modelling.” *Learning Information Systems*

Methodologies.

- Pedersen, J. S., Borum, F & Mathiassen, L. (1997) "Legitimation and Identity: Bridging Between Imitation of Standard Solutions and Creation of Unique, Contextual Solutions." *The 13th EGOS Colloquium*, Budapest, Hungary.
- Scheepers, H. (1999) "The Computer-Ndaba Experience: Introducing IT in rural community in South Africa." *Proceedings of 22nd Information Systems Research Seminar in Scandinavia*, Jyväskylä, Finland, 127-142.
- Scheepers, H. & Mathiassen, L. (1998) "Out of Scandinavia: Facing Social Risks in IT development in South Africa." *IRMA 98*, Boston, Massachusetts, USA, May, 130-142..
- Scheepers, H. & Villiers, C. de (1999) "Teaching of a Computer Literacy Course: A case study using traditional and co-operative learning." *Proceedings of the 7th European Conference on Information Systems (ECIS 99)*, Copenhagen, Denmark, 935-945.
- Scheepers, R., & Damsgaard, J. (1997) "Using Internet technology within the organization: a structural analysis of intranets." *Group '97, the International ACM SIGGROUP Conference on Supporting Group Work "The Integration Challenge"*, Phoenix, Arizona, 9-18.
- Scheepers, R., & Damsgaard, J. (1997) "Using Internet technology within the organisation: a structural analysis of intranets." *Proceedings of the 20th Information Systems Research Seminar In Scandinavia*, Hankö, Norway, 647-664.
- Skov, M., & Stage, J. (1996) "Object-Oriented Design of Process Architecture: An Exploratory Study." *Proceedings of the 19th Information Systems Research Seminar in Scandinavia*, Sweden, 975-1000.

Book Chapters

- Andersen, P. B., & Callesen, J. (2000) "Agents and Actors." In L. Qvortrup, Ed., *Virtual Interaction: Interaction in Virtual Inhabited 3D Worlds*, 132-165. London, Springer Publishers.
- Andersen, P. B., Madsen, C. B., & Granum, E. (2000). "Discussion". In L. Qvortrup, Ed., *Virtual Interaction: Interaction in Virtual Inhabited 3D Worlds*, 209-213. London, Springer Publishers
- Dahlbom, B., & Mathiassen, L. (1996): "Power in Systems Design". In Rob Kling (Ed.), *Computers and Controversy-Value Conflicts and Social Choices*, 903-906, California: Academic Press.
- Damsgaard, J., & Scheepers, R. (1999) "Implementering og styring af intranet (in Danish)." In Kim Viborg Andersen, Jan Damsgaard, and Bjørn Kraft, Eds., *Internethåndbogen*, chapter 10.5. København, Danmark: Børsens ledelseshåndbøger.
- Damsgaard, J., & Lyytinen, K. (1998) "Government intervention in the diffusion of EDI: Goals and Conflicts." In Kim Viborg Andersen, Ed., *EDI and Data Net-*

working in the Public Sector: Government Action, Diffusion and Impacts, 13-41. Boston, MA: Kluwer Publishers.

Eriksen, L. B., Ihlström, C., & Palmer, J. (2000) News: From Product to Service. In Braa, K. *et al.* (Eds.), *Planet Internet*, Chapter 4, 87-112. Lund: Studentlitteratur.

Mathiassen, L., & Stage, J. (1999) Informatics as a Multi-Disciplinary Education. In F. Ljungberg (Ed.), *Informatics in the Next Millennium*, 45-57. Lund: Studentlitteratur.

Skov, M. (2000) Enhancing the Usefulness of Community Chat Rooms: Application of Autonomous Agents. In L. Qvortrup *et al.* (Eds.), Berlin: Springer-Verlag.

Skov, M., Stage J. (2000) Using Software Engineering Approaches to Model Dynamics in Interactive Software Systems. In L. Qvortrup *et al.* (Eds.), Berlin: Springer-Verlag.

Wibroe, M., Nygaard, K. K., & Andersen, P. B. (2000) "Games and stories". In L. Qvortrup, Ed., *Virtual Interaction: Interaction in Virtual Inhabited 3D Worlds*, 166-181. London, Springer Publishers.

Publications Edited

Baskerville, R. L., Stage, J., & DeGroes, J. I., Eds. (2000) *Proceedings of the IFIP TC 8 WG 8.2 Working Conference on Organizational and Social Perspectives on Information Technology*. Boston: Kluwer.

Dahlbom, B., Ljungberg, F., Nuldén, U., Simon, K., Stage, J., & Sørensen, C., Eds. (1996) *Proceedings of the 19th Information Systems Research Seminar in Scandinavia*, Vol. 1 and 2. Gothenburg University: Department of Informatics, Gothenburg Studies in Informatics. Report No. 8.

Buch, N. J., Damsgaard, J., Eriksen, L. B., Iversen, J. H., & Nielsen, P. A., Eds. (1998) *Proceedings of the 21st Information Systems Research Seminar in Scandinavia on Information Systems Research in Collaboration with Industry*, Vol. 1 and 2. Aalborg University: Department of Computer Science.

Kyng, M., & Mathiassen, L., Eds. (1997) *Computers and Design in Context*. MIT Press.

Academic Teaching Cases

Damsgaard, J., & Hørlyck, J. (2000) *Designing www.LEGO.com /shop: Business Issues and Concerns*. Aarhus, Denmark: Department of Management. Teaching Case, 20 pages.

Damsgaard, J. (1998) *Traxon Asia Ltd*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School

Publishing and European Case Clearing House. Teaching Case, 98/01C, 20 pages.

Damsgaard, J. (1998) *CargoNet: Transportation Community Network Limited*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Case, 98/09C, 18 pages.

Damsgaard, J., & Farhoomand, A. F. (1998) *Building Electronic Commerce Infrastructure: Hong Kong Consumer Goods Distribution (A)*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Case, 98/03C, 22 pages.

Damsgaard, J., & Farhoomand, A. F. (1998) *Building Electronic Commerce Infrastructure: Hong Kong Consumer Goods Distribution (B)*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Case, 98/04C, 9 pages.

Farhoomand, A. F., Law, K. K., Ng, P. S. P., & Damsgaard, J. (1998) *Tradelink Electronic Commerce Ltd.: Implementation Strategy*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Case, 98/14C, 20 pages.

Damsgaard, J. & Hørlyck, J. (1999) *Designing www.LEGO.com /shop: Business Issues and Concerns*. Aarhus, Denmark: Department of Management. Teaching Note, 9 pages.

Damsgaard, J. (1998) *CargoNet: Transportation Community Network Limited*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Note, 98/09TN, 6 pages.

Damsgaard, J. (1998) *Traxon Asia Ltd.* : Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Note, 98/01TN, 9 pages.

Damsgaard, J. & Farhoomand, A. F. (1998) *Building Electronic Commerce Infrastructure: Hong Kong Consumer Goods Distribution (A+B)*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Note, 98/04TN, 9 pages.

Farhoomand, A. F., Law, K. K., Ng, P. S. P., & Damsgaard, J. (1998) *Tradelink Electronic Commerce Ltd.: Implementation Strategy*. Available from Centre for Asian Business Cases, School of Business, Hong Kong University and Harvard Business School Publishing and European Case Clearing House. Teaching Note, 98/14TN, 8 pages.

Other Publications

Skov, M., Stage, J. (1996) *Process Architecture Study: Documentation*. Available at the WWW site <http://www.cs.auc.dk/~jans/procarch/>.

Stage, J. (1999) *Object-Oriented Multimedia System Development: Documentation*. Available at the WWW site <http://www.cs.auc.dk/~jans/matchrace/>.

Part III

Summary in Danish

Chapter 9

Resumé

Ifølge retningslinier fra det Teknisk-naturvidenskabelige Fakultet ved Aalborg Universitet skal dets forskningsenheder evalueres hver femte år. Denne rapport dokumenterer i overensstemmelse hermed den tredje forskningsevaluering af forskningsenheden Datalogi (i det følgende benævnt Institutet) og dækker perioden fra 1996 til 2000.

9.1 Formål

Ifølge Fakultetets retningslinier har forskningsevalueringen tre formål, idet den skal

- a) vurdere, om der er tilfredsstillende overensstemmelse mellem de tildelte interne og eksterne forskningsressourcer og det udførte forskningsarbejde,
- b) vurdere om der er rimelig sammenhæng mellem Institutets mål for den forskningsmæssige indsats, herunder som den fremgår af fakultetets langsigtede faglige planlægning, og den gennemførte forskning og
- c) rådgive Institutet om dens fremtidige indsats og forskningsorganisering.

Institutet, der er underinddelt i fire forskningsenheder, tilføjede et ekstra punkt, som tildels understreger visse aspekter af Fakultetets formål:

Inden for forskningsenheden skal evalueringen, processen såvel som den endelige rapport, konstruktivt hjælpe medarbejderne med at evaluere og forbedre deres effektivitet som forskere, gruppeledere og som administratorer af forskning på forskningsenhedens niveau. Specielt er evalueringen en anledning til at vurdere forskningsenhedens forskningspolitik.

Processen, der ledte frem til denne rapport, strakte sig over et år og beskrives kort i det følgende.

9.2 Evalueringsprocessen

Evalueringsprocessen tog sin start i april 2000. En af de vigtige, tidlige aktiviteter var at etablere et evalueringsudvalg. Professor Finn V. Jensen accepterede at være lokal repræsentant i udvalget (lokale medlemmer har ingen stemmeret), og det var med stor tilfredshed i forskningsenheden, at Søren Damgaard (IBM Denmark), professor Matthias Jarke (RWTH Aachen) og professor Moshe Vardi (Rice University) også accepterede at deltage i udvalget. Det var vores fornemmelse, at dette udvalg af anerkendte og erfarne seniorforskere ville være istand til at dække de forskelligartede forskningsområder, som Instituttet rummer, og ville kunne udarbejde en indsigtfuld evaluering.

I løbet af foråret 2000 vedtog vi den overordnede struktur for evalueringsprocessen og for denne rapport. Det var en stor hjælp, at vi allerede havde gennemført to evalueringer.

Rapporten skulle have to hoveddele, nemlig Instituttets beskrivelse af sin forskning og selve udvalgets evaluering af Instituttet. Forskningsbeskrivelsen skulle struktureres, så der var et kapitel for Instituttet generelt og et for hver af de fire forskningsenheder i Instituttet.

Der blev også vedtaget overordnede skabeloner for disse kapitler. Herefter var det op til ledelsen og hver forskningsenhed at udfylde skabelonerne. Formålet med dette design var at tillade de fagligt forskellige enheder en vis grad af frihed til at rapportere deres forskning samtidig med at resultatet stadig blev en tilpas homogen rapport.

I begyndelsen af december 2000 blev en midlertidig udgave af forskningsbeskrivelsen på godt og vel 214 sider (den anden del af denne rapport) sendt til evalueringsudvalget.

Evalueringsprocessen kulminerede i januar 2001 med et to-dags internat, hvor alle forskningsmedarbejdere, herunder Ph.D. studerende og administrativt personale, var inviteret. Den første dag præsenterede hver forskningsenhed sin forskning, og der var lejlighed til at panelet og andre kunne stille spørgsmål til præsentationerne og det skrevne materiale. Den anden halvdel af den første dag arbejdede udvalget separat, mens de øvrige deltog i forskellige aktiviteter omkring temaet forskningsevaluering. Den anden dag fremlagde udvalget sin evaluering, besvarede spørgsmål og modtog kommentarer under den efterfølgende diskussion.

Efter internatet færdiggjorde udvalget sin evaluering. Samtidig producerede ledelsen og forskningsgrupperne reviderede udgaver af deres bidrag til rapporten, idet de baseret på kommentarer fra internatet typisk foretog relativt mindre forbedringer af de foreløbige beskrivelser. Det er disse reviderede udgaver, der danner basis for kapitler i rapportens anden del. Derfor er evalueringen i den første del og beskrivelserne i den anden del ikke fuldstændigt synkroniserede.

9.3 Denne Rapport

Den foreliggende rapport har tre dele. Den første del rummer selve evalueringen af Instituttet. Denne del beskriver først den formelle ramme for evalueringen og præsenterer derefter udvalgets evaluering (konklusionerne herfra følger på dansk i næste kapitel). Som omtalt ovenfor har den anden del et kapitel for Instituttet i sin helhed og et for hver af de fire forskningsenheder i Instituttet. Det første kapitel beskriver generelle forhold ved Instituttet. Det dækker Institutets baggrund og organisatoriske kontekst samt den overordnede organisering af Institutets forskning. Desuden præsenteres forskningsplanen for perioden 1996 til 2000, og de tilgængelige ressourcer for gennemførelsen af planen beskrives. Endelig evalueres den overordnede administration og organisering af forskningen i Instituttet, og en plan for 2001 til 2005 præsenteres. De enkelte forskningsenheders kapitler sammenfatter enheders forskningsaktiviteter og -resultater og beskriver enheders organisering og bemanning. Så følger beskrivelser af forskningssamarbejde, Ph.D. projekter, og service- og forskningsrelaterede aktiviteter. Endelig præsenteres selvevalueringer og planer for den næste periode. Rapportens sidste del er denne opsummering på dansk.

Chapter 10

Evalueringsudvalgets Konklusion

Institut for Datalogi ved Aalborg Universitet er nu godt på vej mod at blive et af Europas førende institutter for datalogi og er allerede et af de stærkeste i Skandinavien. Dette kan ses blandt andet af det faktum, at lederne af de tre enheder som er aktive inden for datalogiske kerneområder alle optræder blandt de 0.5 % mest citerede datalogi-forfattere i ResearchIndex fra the amerikanske projekt *NECI Scientific Literature Digital Library*.

I løbet af rapporteringsperioden er instituttet blevet reorganiseret som fire enheder, der hver ledes af en internationalt velanskreven forsker. Hver af de fire forskningsenheders internationale status, såvel som instituttets som helhed, har forbedret sig signifikant i løbet af de seneste fem år med hensyn til publikationer, samarbejde med industrien og overordnet gennemslagskraft. Meget af denne succes kan tilskrives aktive beslutninger truffet af instituttet i løbet af den seneste fem-årige periode.

Imidlertid skaber den enorme efterspørgsel efter datalogi et strategisk behov for at styre den vækst, man vil kunne forudse, så den er til langsigtet fordel for instituttet med dens væld af nye og hermed forbundne udfordringer. Blandt disse udfordringer er blandt andet dette at definere nye vækstområder, en øget vilje til at acceptere ansvar for styring samt strategier for ansættelse og karriereplanlægning. Den fremragende forberedelse af og de meget klare og åbne diskussioner ved evalueringsmødet har hjulpet evaluatorerne med at udarbejde en række specifikke anbefalinger desangående.

Søren Damgaard Matthias Jarke Moshe Vardi
IBM Denmark RWTH Aachen Rice University