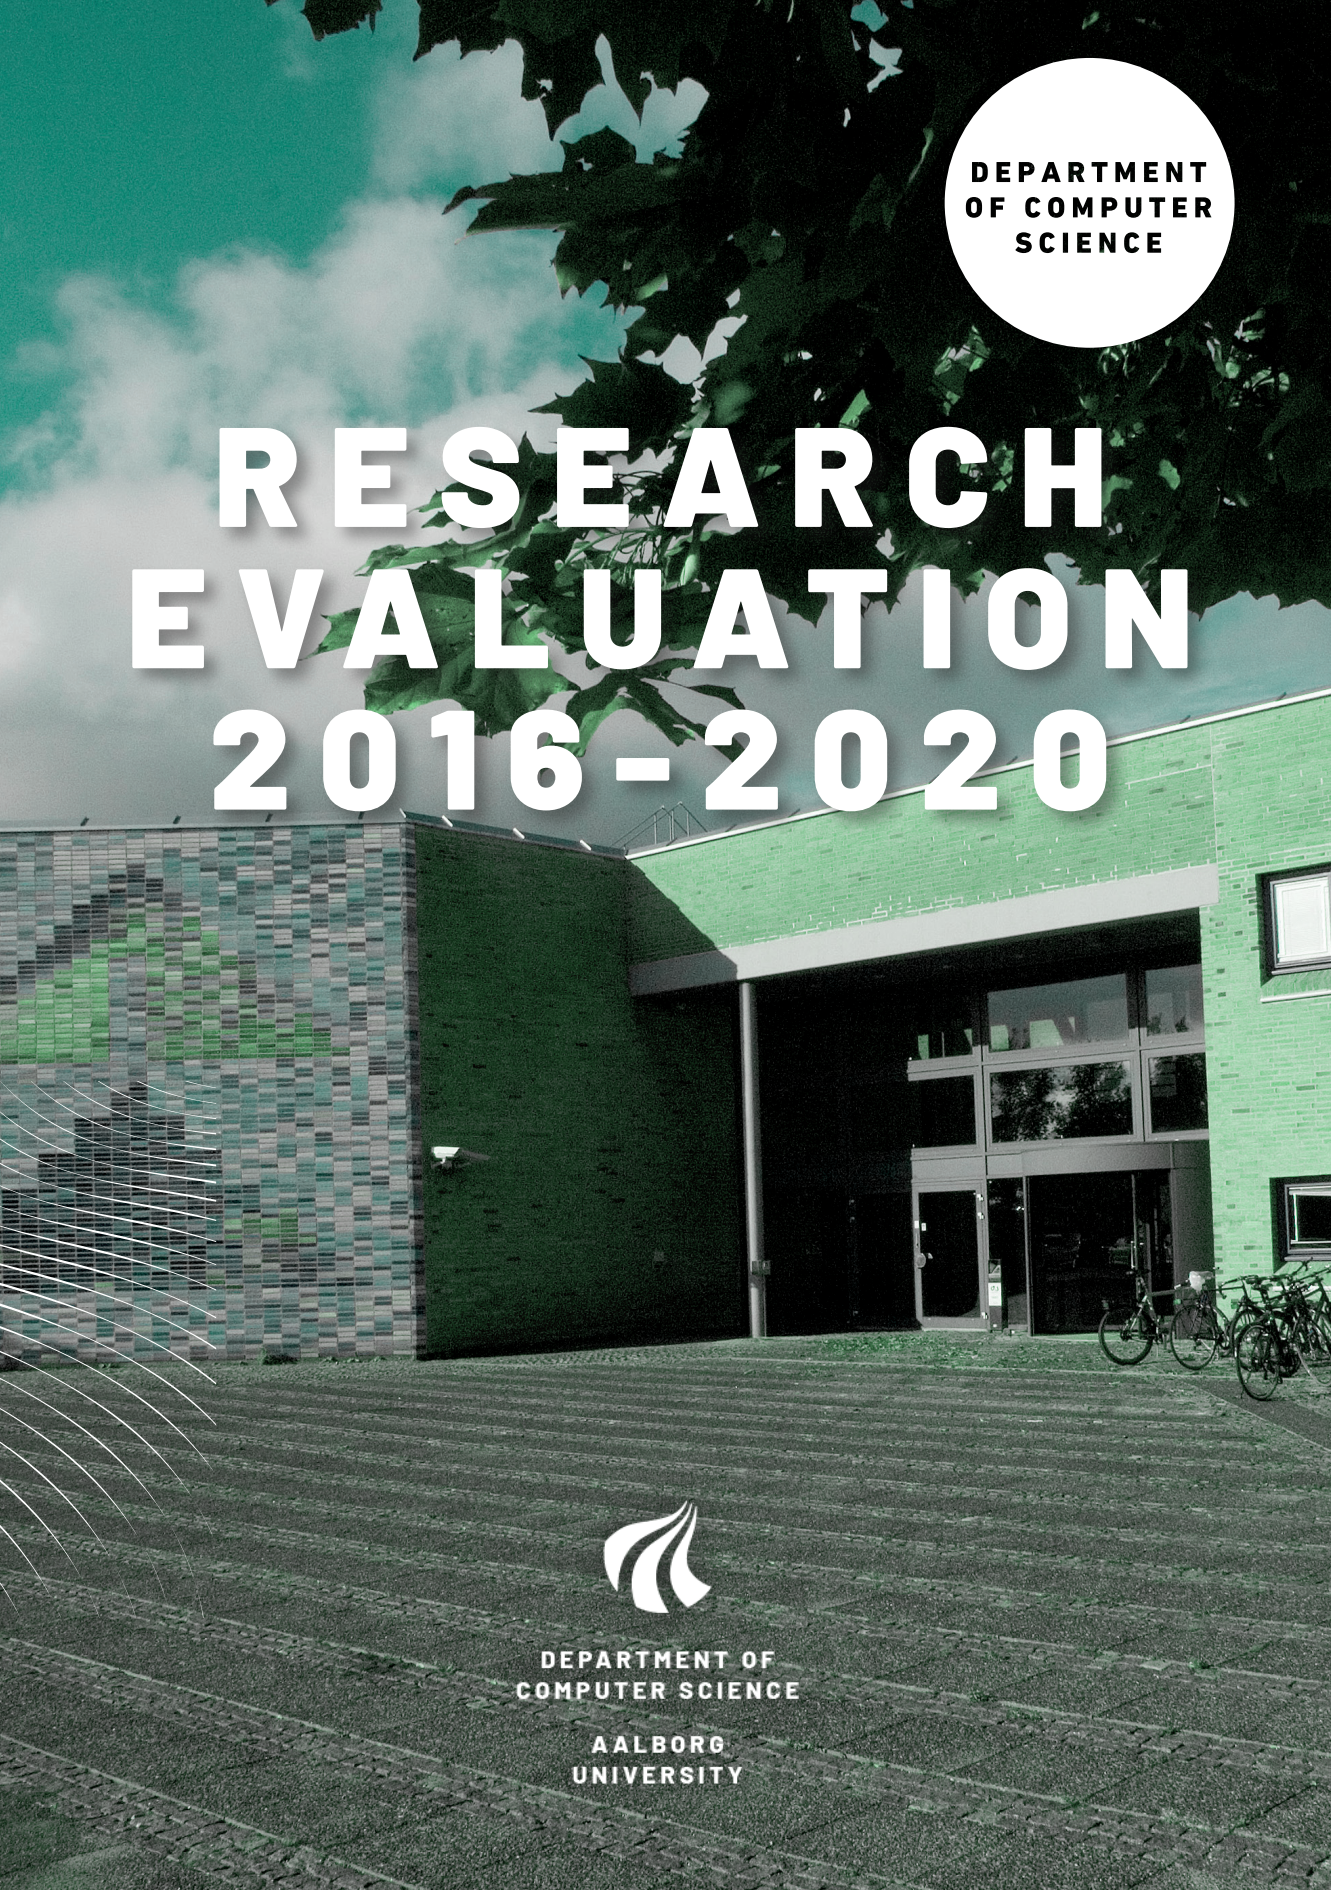

DEPARTMENT OF COMPUTER SCIENCE

RESEARCH EVALUATION 2016-2020

DEPARTMENT OF
COMPUTER SCIENCE

AALBORG
UNIVERSITY

Computer Science, Aalborg University
Research Evaluation 2016-2020

Edited by:
Jesper Kjeldskov and Uffe Kjærulff (sub-editors CS Department)
Bin Yang and Gabriela Montoya (sub-editors DPW)
Kim G. Larsen and Giorgio Bacci (sub-editors DEIS)
Mikael B. Skov and Anders R. Bruun (sub-editors HCC)

All photos by Lars Horn, Baghuset

Printed by Novagraf

ISBN 978-87-991383-3-3

Contents

RESEARCH EVALUATION 2016-2020

Department of Computer Science

Aalborg University 2021

THE EVALUATION COMMITTEE

Due to the Corona epidemic the evaluation process was carried out online (using Zoom) distributed over several dates in April 2021.

From the top-left corner:
Professor Franz Baader,
Professor Marieke Huisman,
Professor Kim Guldstrand Larsen (chairman).

From the bottom-left corner:
Professor Yannis Manolopoulos,
Professor Jan Gulliksen and
Professor Bernhard Mitschang.

1 The Evaluation

1.1 The Evaluation Committee

Franz Baader

Franz Baader is full professor for Automata Theory at TU Dresden (Germany) since 2002 and director of the Institute for Theoretical Computer Science at TU Dresden since 2005. From 2012 to 2015 he was Dean of the Faculty of Computer Science of TU Dresden. He has received his doctoral degree (Dr.-Ing.) in Computer Science from the University of Erlangen-Nürnberg in 1989, and from 1989 to 1993 was senior scientist and project leader at the Germany Research Center for Artificial Intelligence (DFKI) in Kaiserslautern and Saarbrücken. From 1993 to 2002 he was professor for Theoretical Computer Science at RWTH Aachen.

His main research areas are Logic in Computer Science and Artificial Intelligence, in particular Automated Deduction and Knowledge Representation. Notably, he and his research group have worked on Description Logics for 30 years, and have laid the logical and algorithmic foundations for the Description Logics underlying OWL and the OWL 2 profile OWL 2 EL. Franz Baader is co-editor of the highly-cited Description Logic Handbook and co-author of the first textbook on Description Logics, both published by Cambridge University Press. He also wrote, together with Tobias Nipkow, a monograph on Term Rewriting Systems, which is the standard reference for the field. Franz Baader has

published more than 250 research articles in international conferences and journals, is on the editorial board of several journals in logic and AI, and has chaired the PCs of several conferences in these areas. He has received a high number of research grants from the German Research Foundation (DFG) and is speaker of a DFG-funded PhD program on Quantitative Logics and Automata (QuantLA). Franz Baader is a fellow of the European Association for Artificial Intelligence (EurAI) since 2004 and a member of the Academia Europaea since 2011. In 2020 he received the prestigious Herbrand Award for exceptional contributions to the field of Automated Deduction.

Jan Gulliksen

Jan Gulliksen is a professor in Human Computer interaction at KTH Royal Institute of Technology, Stockholm, Sweden, since 2009. He is also the Vice President for Digitalization for the university. Before that he was the Dean of the School of Computer Science and Communication. Before joining KTH he was a professor of Human Computer Interaction at Uppsala University. Jan was between 2012 and 2016 the chairman of the Digital Commission of Sweden's government, Department of the Enterprise and he is now a member of the Digitalization Council serving the Minister for Digitalization with the Swedish Government. He was also appointed Digital champion of Sweden by the European commission in 2012, a role he still keeps. Jan is the Vice President of the International Federation for Information Processing (IFIP) and is and has been an active member of IFIP TC 13 on HCI, he was the founder and chair for the first 20 years for the NordiCHI conference series. Jan has conducted a number of larger Action Research projects with the purpose of increasing the usability of digital work support systems by introducing new user-centered design methodologies particularly at several public authorities. He is doing research on accessibility and the effect of digitalization on people with disabilities. Finally he is working on projects relating to the digitalization of industry, digital work environments and digitalization and leadership. Jan is a fellow of the Academy of Engineering Sciences (IVA) in Sweden. Jan founded the company

Gulliksen Consulting AB in 1999 and is still chairman of the board.

Marieke Huisman

Marieke Huisman is a professor in Software Reliability at the University of Twente, Netherlands. She is well-known for her work on program verification and specification. At the University of Twente, she leads the Formal Methods and Tools group. In 2011, she obtained an ERC Starting Grant for the VerCors project on the verification of concurrent software, where she studied practical verification of concurrent programs, and developed verification techniques for advanced programming and specification constructs. All verifications are supported by the VerCors tool set. In follow-up projects, the results have been expanded to programming languages for other parallel computing paradigms, such as GPUs and distribution (EU project CARP, NWO Top project VerDi). She currently is working on adding support to reason at a more abstract level, and to increase the level of automation of the verification process, as part of her NWO VICI project Mercedes (2018 - 2022). In 2013, she received the Netherlands Prize for ICT research 2013.

Marieke Huisman plays an active role in the wider CS community. She is the chair of the ETAPS association that is responsible for a series of top conferences in her area, for several years she chaired the association of Dutch software researchers VERSEN, she is a member of The Informatics Platform Netherlands (IPN) and she also chairs the informatics board of the Lorentz center in Leiden. Moreover, she has co-organised workshops at dedicated workshop venues such as Dagstuhl and the Lorentz center, as well as multiple verification competitions. Finally, at the moment she also actively contributes to setting up a satellite location of the University of Twente in Apeldoorn.

In addition to her scientific work, she also developed many activities in the area of diversity & inclusion. In addition to high school visits and other outreach activities, she chairs the D&I committee at the University of Twente (ambassadors' network), and the national IPN working group on Equity, Diversity, and Inclusion. Moreover, in January 2020, she was one of the key people behind the Alice & Eve event and exhibition, a celebration of women in computing.

Yannis Manolopoulos

Yannis Manolopoulos holds a 5-years Diploma degree in Electrical Engineering (1981) and a Ph.D. degree in Computer Engineering (1986), both from the Aristotle University of Thessaloniki. He is Professor and Vice-Rector of the Open University of Cyprus as well as Professor Emeritus of the Aristotle University of Thessaloniki. He has been with the University of Toronto, the University of Maryland at College Park, the University of Cyprus and the Hellenic Open University. He has also served as President of the Board of the University of Western Macedonia in Greece and Vice-President of the Greek Computer Society. His research interest focuses in Data Management. He has co-authored 6 monographs and 10 textbooks in Greek, as well as >350 journal and conference papers. He has received >15700 citations from >2300 distinct academic institutions from >100 countries (h-index=57). He has also received 5 best paper awards from SIGMOD, ECML/PKDD, MEDES (2) and ISSPIT conferences. He delivered keynote talks at 24 conferences at: Albania, Algeria, Austria, Bulgaria, Cyprus, Czechia, France, Greece, Italy, Kosovo, Lebanon, Luxembourg, Montenegro, Morocco, Poland, Romania and Russia. He served as external member of 20 doctoral examination committees at: Brazil, Denmark, France, Italy, Poland and Spain. Evaluator for national funding agencies: Austria, Canada, Cyprus, Czech Republic, EU, Estonia, Georgia, Greece, Hong-Kong, Israel, Italy, Russia, Switzerland. Currently, he serves in the Editorial Boards of the following journals (among others): Digital (EiC), Information Systems, World Wide Web, Computer Journal, Data Science and Analytics, as well as in the Board of the Research and Innovation Foundation of Cyprus.

Bernhard Mitschang

Since 1998, Bernhard Mitschang holds the chair for Database and Information Systems and further is full professor as well as head of the department 'Applications of Parallel and Distributed Systems' that is part of the Institute of Parallel and Distributed Systems (IPVS) at the Faculty of Computer Science, Electrical Engineering, and Information Technology at the Universität Stuttgart, Stuttgart, Germany.

From 1994 to 1998 he held the position of a professor at the Technische Universität München. In 1988 he received the Ph.D. degree (Dr.-Ing.) in Computer Science from the University of Kaiserslautern, and in 1994 he got the venia legendi (Dr. habil.) for practical Computer Science from the University of Kaiserslautern. From 1989 to 1990 he was on leave to IBM Almaden Research Center, San Jose, CA as a visiting scientist, and during Summer semester 2003, 2008, 2013, and 2019 he was on sabbatical leave to IBM Research and Development Lab in Böblingen, Germany. Since 2014, Bernhard Mitschang is head of the Technology Partnership Lab at the university and head of the board of managing directors of the Graduate School of Excellence on advanced Manufacturing Engineering (GSaME) as well as member of GSaME.

Both research and teaching spectra of his department cover on one hand data-intensive applications ranging from business applications to engineering systems and on the other hand fundamental data management techniques, data analytics as well as scalable data processing architectures.

Kim Guldstrand Larsen (chairman)

Kim Guldstrand Larsen is a Professor of Computer Science at Aalborg University (AAU), since 1993. He received his Ph.D. in Computer Science from Edinburgh University, Scotland, in 1986 under the supervision of Professor Robin Milner. He received Honorary Doctorate from Uppsala University (1999), ENS Cachan (2007), International Chair at INRIA (2016) and Distinguished Professor at North-Eastern University, Shenyang, China (2018). His research interests cover modeling, verification, performance analysis of real-time and embedded systems with applications to concurrency theory, model checking and machine learning.

He is the prime investigator of the verification tool UPPAAL for which he received the CAV Award in 2013. Other prized received include Danish Citation Laureates Award, Thomson Scientific Award as the most cited Danish Computer Scientist in the period 1990-2004 (2005), Grundfos Prize (2016), Ridder af Dannebrog (2007). He is member of the Royal Danish Academy of Sciences and Letters, The Danish Academy of Technical Science, where he is Digital wiseman. Also, he is member of the Academia Europaea. In 2015 he received the prestigious ERC Advanced Granat (LASSO), and in 2021 he won Villum Investigator Grant (S4OS). He is on the editorial board of the journal Formal Methods in System Design, and serves on the steering committee (and has been co-founder) of several conferences series, e.g. ETAPS., CONCUR, TACAS, FORMATS, SETTA. He has been PI and director of several large centers and initiatives including CISS (Center for Embedded Software systems, 2002-2008), MT-LAB (Villum-Kahn Rasmussen Center of Excellence, 2009-2013), IDEA4CPS (Danish-Chinese Research Center, 2011-2017), INFINIT (National ICT Innovation Network, 2009-2020), DiCyPS (Innovation Fund Center, 2015-2021). Finally, he is co-founder of the companies UP4ALL (2000), ATS (2017) and VeriAal (2020).

1.2 Evaluation Guidelines

Below are the guidelines for the evaluation committee members for the research evaluation of The Department of Computer Science at Aalborg University that took place during the spring of 2021:

- The members should:
  - Apply the best of their abilities, professional skills, knowledge and ethics, in the evaluation.
  - Work together with the department representative on the evaluation panel.
  - Take part in the preparation meeting on April 6, 2021, and the online evaluation meetings on April 12, April 16 and April 23.
  - Contribute to the written report as specified below, with the report being ready at the end of the meeting or very soon thereafter.
  - Contribute to the presentations of the evaluation results, as specified below.

- Four documents were sent to the committee members on February 21, 2021, consisting of the first drafts of the self-evaluation of the three research groups (DPW, DEIS and HCC) as well as a self-evaluation of at the department level.

On April 6, 2021, the final and collected version of the Self-evaluation was sent to the committee members including detailing Appendices from the three research groups.

• Each evaluator has main responsibility for one research unit. The responsibilities are as follows: Database and Programming Technologies: Yannis Manolopoulos and Bernhard Mitschang; Distributed, Embedded and Intelligent Systems: Marieke Huisman and Franz Baader; Human Centered Computing: Jan Gulliksen.

• All members should read the main report. The main responsible members should read the relevant group chapter in the main report with special attention.

• The main responsible members should read the appendix for the given research group and should provide comments on all aspects of the unit chapter. The other members are encouraged to comment on selected aspects. The main responsible(s) should integrate the comments into the evaluation report.

• All committee members should provide comments on the department chapter.

• Comments should be on the actual research topics, results, and performance, but also on strategic issues like funding, internal organization and synergies, possible new directions, collaboration with industry, internationalization, SDG, positioning IT as a key enabler in society, etc.

• Comments should be given at the department, research group, and, if relevant, subgroup level.

• The main findings should be presented in dedicated feedback sessions, at the research group and department level, as specified by the evaluation program. The primary responsible(s) will present the main findings at the research group level. The external evaluators will jointly present the main findings at the department level.

1.3 Evaluation Process

Due to the Corona epidemic the evaluation process was carried out online (using Zoom) distributed over several dates in April 2021.

On April 6, 2021, the evaluation had a preparation meeting focusing on alignment of expectations, discussion on the plan for the 3 evaluation days as well as planning the practical execution of the research evaluation.

On the three dates April 4, April 16 and April 23, 2021, the evaluation took place. Given the online setting, the entire evaluation activity (including internal committee work) was for each of the three dates confined to time period 9.00-14.00.

- On the first day (April 4), overview talks on the department and each of the three research groups were given. Also, the committee meet and interviewed selected Assistant Professor/Postdocs and selected PhD students to get a more complete view of the research environment.
- The second day (April 16) started with two parallel sessions where each research group presented further details and had question and answer sessions with their main responsible committee member(s). Following the committee reflection, the day ended with a meeting with the committee, Head of Department and the Professors of the Department.
- On the evening of April 21, the committee arranged themselves an additional (private) online dinner to reflect over the findings of the two first evaluation days in a less time-pressed setting.
- The last day (April 23) started with committee discussions, followed by a meeting discussing strategic issues with Torben Larsen, Vice-Dean of Research at The Technical Faculty of IT and Design, and Jesper Kjeldskov, Head of Department of Computer Science. The days ended with plenary feedback sessions with assessment and recommendations for each research unit, the department, as well as the faculty/university level. These final feedback sessions were participated by Henrik Pedersen, Dean of The Technical Faculty of IT and Design.
- Following the seminar, the committee prepared written final versions of their assessments and recommendations. They are included in this report at the end of the chapters for the different organizational and research units.

2 The Department of Computer Science

2.1 Introduction

Department of Computer Science consists of 210 employees and 1300 students mostly located in, or nearby, the "Cassiopeia" building in Aalborg University's Campus East. Until 2020, the department had all of its activities centred at the Aalborg Campus, but in September 2020, a new group was established at the Copenhagen Campus.

Department of Computer Science is part of Aalborg University's relatively new Technical Faculty of IT and Design (TECH) that comprises four departments. The faculty was formed in 2016 when the Faculty of Engineering and Science was split into two smaller faculties.

The department is formally managed by a Head of Department (Professor Jesper Kjeldskov) responsible for the overall operation, budget and staff, a Deputy Head of Department (Associate Professor Uffe Kjærulff) responsible for education, and a Head of Administration (Helle Westmark) responsible for support and support staff.

In early 2020, a Department Leadership Group was established to oversee the operations of the department. This group includes the formal management as well as two of the department's prominent professors, Christian S. Jensen and Peter Axel Nielsen.

In late 2020, a new role of Deputy Head of Studies was established to support managing the department's educations (Associate Professor Ulrik Nyman).

The department's scientific staff is organised in a matrix structure with three core research groups, and two "teams" going across these groups (Figure 2.1).

Database and Web Technologies (DW) emphasizes querying, data mining, and machine learning in two broad areas: data-intensive systems, including spatio-temporal and multidimensional data management and analytics; and web science

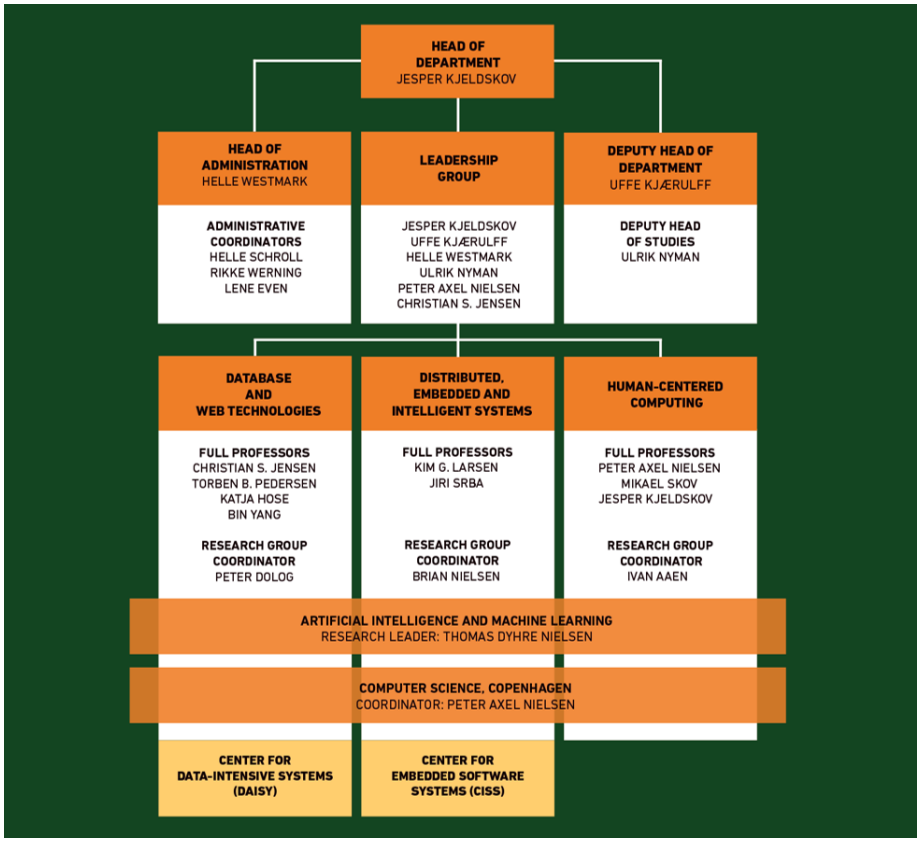

Figure 2.1: The organisational structure of Department of Computer Science.

and knowledge engineering, including graph data management, knowledge graphs, recommender systems. The research group is led by Professor Christian S. Jensen, Professor Torben B. Pedersen, Professor Katja Hose, and Professor Bin Yang, alongside Associate Professor Peter Dolog who represents the group in the department's coordination group.

Distributed, Embedded and Intelligent Systems (DEIS) emphasizes foundational and logical theories, algorithms and tools for verification and validation, methodologies for real-time and embedded systems, networks and operating systems, probabilistic graphical models, machine learning, and programming technologies. The research group is led by Professor Kim G. Larsen and Professor Jiri Srba, alongside Associate Professor Brian Nielsen who represents the group in the department's coordination group.

Human-Centered Computing (HCC) emphasizes the design, development, and use of interactive computer systems from a human-centered perspective, including software development, digitalization, human-computer interaction, interaction design, user experience, and Innovation. The research group is led by Professor Peter Axel Nielsen, Professor Mikael B. Skov, and Professor Jesper Kjeldskov, alongside Associate Professor Ivan Aaen who represents the group in the department's coordination group.

Artificial Intelligence and Machine Learning (AI-ML) gathers researchers from all research groups at the Department of Computer Science, working with different aspects of Artificial Intelligence and Machine Learning, under the research leadership of Professor MSO Thomas Dyhre Nielsen. The team is set up to be the department's bridgehead in the area and responding to the current significant interest and opportunities, including funding and collaboration with industry and other disciplines. The team is also designated to roll out the new Data Science education (expected to be renamed Data Science & Artificial Intelligence in 2022).

Computer Science, Copenhagen (CS-CPH) is responsible for implementing and running the department's new Software education on Aalborg University's Copenhagen Campus. The team comprises of members from all three research groups and is coordinated by Professor Peter Axel Nielsen. The team is set to gradually expand as the Software education is rolled out over the next 4 years. Recruiting is done with a focus on gradually acquiring the competences needed for delivering the teaching required on the Software education. In the short term, we are recruiting assistant and associate professors.

2.2 Students and Teaching

The three research groups deliver teaching at the department's educations. Most educations, such as Computer Science, Software, Interaction Design, and Information Technology, involves teaching from all three research groups. But newer educations, such as Data Science and IDA, are more narrowly targeted, and some run in collaboration with other departments, and therefore only involve teaching from some of the department's groups The department currently has 6 bachelor's and 7 master's educations (with the master's in Data Science starting in 2022, and the master's in Software, Copenhagen starting in 2023).

Bachelor's educations

- Computer Science
- Software
- Software, Copenhagen
- Information Technology
- Interaction Design
- Data Science

Master's educations

- Computer Science
- Software
- Software, Copenhagen (from 2023)
- Interaction Design
- Data Science (from 2022)
- IT Design and Application Development (IDA)
- Computer Science (IT) (CS-IT)

Common for all educations is a core of computer science subjects and projects, in some cases combined with content from other disciplines, such as design, business,

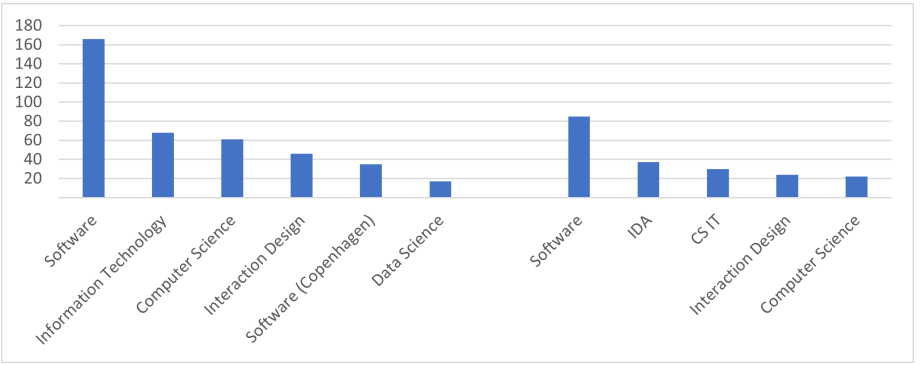

Figure 2.2: Student intake on bachelor (left) and master (right) educations.

the humanities, or social science. As of 2020, the department's student intake is distributed as shown in Figure 2.2.

As expected in the last research evaluation (2015), the department's educations have grown substantially in number of students, as shown in Figure 2.3. In 2015, the total intake was 259. By 2020 this number had increased to 393 (52%). However, unlike the previous 5 years period, the increase in student intake has not been a steady linear growth within existing educations, but can mainly be attributed to the introduction of the new Software education in Copenhagen in 2020, and notable jumps in intake on Interaction Design and Information Technology in 2020. The new education in Data Science is also contributing to the growth. While in most of the department's educations have experienced a growth in intake over the last five years, the Computer Science education has experienced a drop of 23% from 80 to 61 students enrolling on the 1. semester.

In the same period, the intake on the department's master's educations has grown steadily from 119 in 2015 to 198 in 2020 (66%) as shown in Figure 2.4.

With the increased intake, the department's total student population has grown as shown in Figure 2.5. In 2015, the total number was 941. By 2020 this number had increased to 1263 (34%). As expected in the last evaluation (2015) this 5-year growth is smaller than the previous 5 years (where the population increased by 84%). This slowdown in population growth is obviously caused by rather stable intake in 2016 to 2019 (see Figure 2.3) followed by a sudden notable increase in 2020 due to a new popular education. As this education is fully implemented over the next 5 years, the total student population is expected to grow substantially, followed by a flattening of the curve at around 1500 students in 2025 (unless further educations are initiated or some educations experience a decrease in intake).

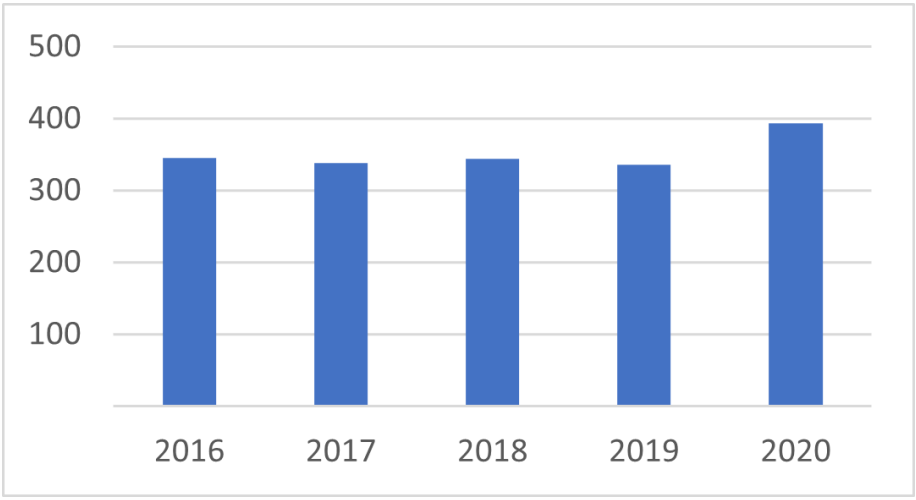

Figure 2.3: Student intake (Bachelor).

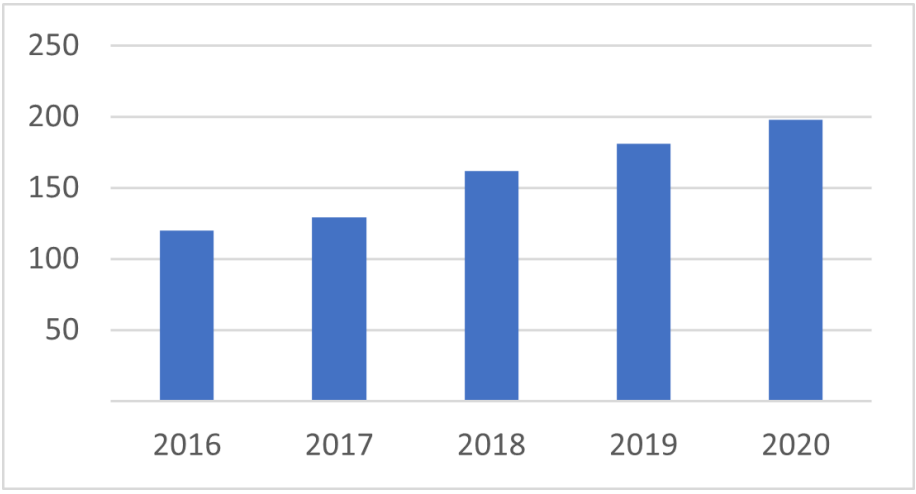

Figure 2.4: Student intake (Master).

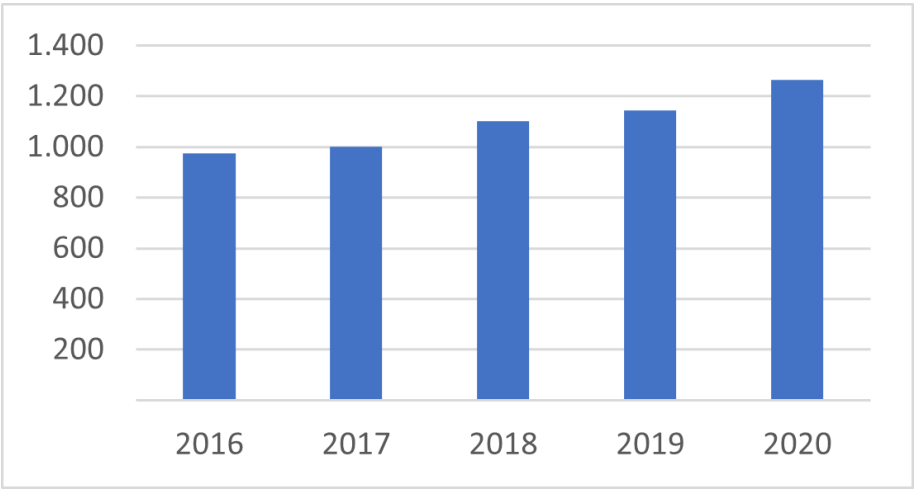

Figure 2.5: Student population.

In the same period, the teaching hours required to run our educations has increased from 41.829 hours/year to 64.027 hours/year (53%). This growth exceeds the increase in student population (35%) and show that our educations are run slightly less efficient. This is due to two factors 1) a deliberate choice to allow slightly smaller classes than before (which the positive current economic situation permits), and 2) the education in Data Science are not yet having high enough student numbers to be in balance economically.

Out of the 64.027 hours/year worth of teaching in 2020, 11% are outsourced to others, including Department of Mathematical Sciences (1900), Department of Electronic Systems (1650), and Department of Architecture, Design & Media Technology (650).

2.3 Funding

The department is mainly funded from two sources: 1) internally allocated funding from the Faculty, and 2) external funding granted by national and international research councils, agencies, and foundations. The internal funding is mainly allocated based on delivered teaching (66%), but also on publications and turnover from external grants (33%) and other factors (1%). From 2018-20, the university implemented a new principle for the allocation of internal funding, which passes a greater amount of funding per student on to the departments. Since the Depart-

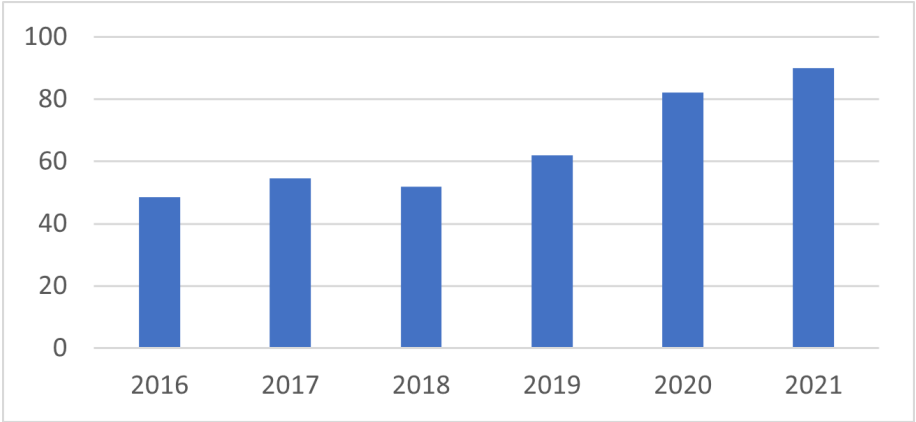

Figure 2.6: Internal income in million DKK from the Faculty (2021 included).

ment of Computer Science's educations were already highly efficient economically, this new model has resulted in an increasing surplus. Consequently, the department has had economical opportunity to recruit massively, invest in new strategic initiatives, and to invest in the physical workspace for staff and students. This opportunity will persist in the near future.

Over the last decade, an increasing proportion of national research funding for universities has been allocated through competitive applications to national research councils. This has put a pressure on the department to increase its activities targeted towards external funding, not only from national government, but also from the EU and private foundations. This has led to a consistent level around 25-30 million DKK of new grants per year.

2.3.1 Internal income

After a long period of relatively stable income from the Faculty, at around 45-50 million DKR. per year, the department's internally allocated income began increasing rapidly from 2018, surpassing 80 million DKK in 2020, and reaching almost 90 million in 2021 (Figure 2.6). This means that the department's internal income has grown by 64% over the last 5-year period and has for the last 3 years been growing steadily at a pace of around 12.5% per year.

This development is the direct result of two factors: 1) the department running its portfolio of educations in an extremely efficient manner economically, and 2) the university rolling out a new economy model for the internal distribution of funds

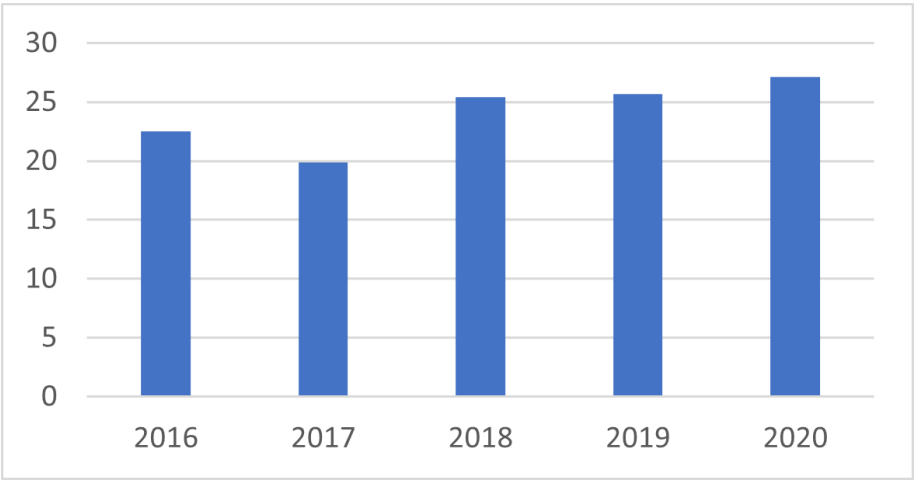

Figure 2.7: External funding (spent) in million DKK.

based on the departments' production, and hence awarding departments that run educations efficiently.

Although Figure 2.6 show a linear growth in internal income from 2018-21, it is important to note that the annual increase is most likely to flatten out after 2021 as the university's new economy model has now been fully rolled out.

2.3.2 External income

The department has experienced a 15% increase in its spending of external income from 22.5 million DKK in 2016 to 26 million in 2020, as can be seen in Figure 2.7. The current level of externally funded expenditure is similar to the last two years, and at a high point since 2013 (where it surpassed 30 million DKK).

It is expected that the spending of external funding will increase in the years to come as a direct effect of an elevated level of funding applications. In 2020, the amount of awarded external funding grew to 33,2 million DKK. Based on analysis at the TECH Faculty, and among the four Heads of Departments, targets are that the department's external funding will grow to 46 million DKK in 2021 and 57 million DKK in 2022.

Working to meet these expectations, the department has hired a full-time fundraiser, and taken several steps to increase the number of applications submitted. Figures 2.8 and 2.9 show the development in applications for external funding. The num-

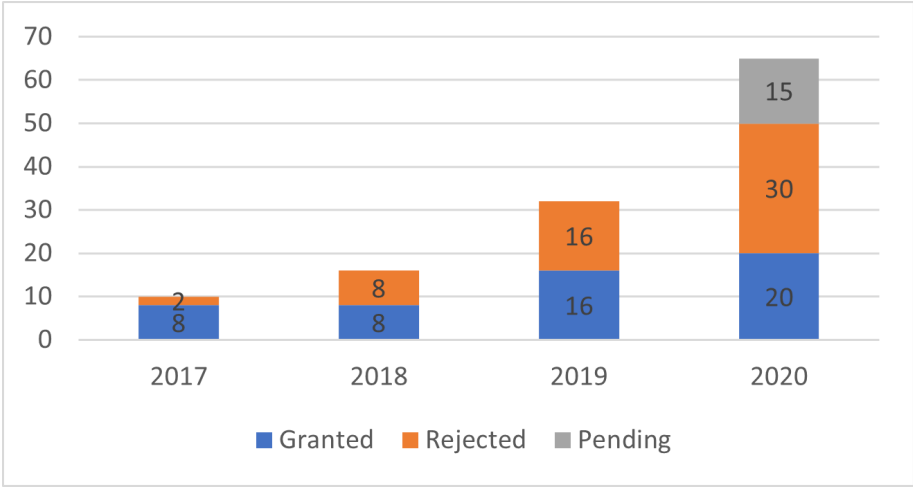

Figure 2.8: Funding applications submitted.

bers start from 2017 when systematic registration of funding applications was first introduced, and from 2019 the numbers reliably show both granted, rejected and pending applications.

Applications are registered in the calendar year they are submitted, approved, or rejected, meaning that some rejections or approvals in one calendar year may be the result of a submission in the previous year.

As can be seen from figures 2.8 and 2.9, the department has more than doubled its activity on external funding applications. From 2019 to 2020, there was an increase of 101% in number of applications submitted (from 32 to 65), and 108% in the amount applied for. The elevated level of applications is already reflected in the total external funding attracted. In 2020 the department had 25% more projects granted than in 2019, amounting to an increase in external funding of 57,9 million DKK. However, the whole effect of the recent doubling in applications is still to be seen. Out of the 65 applications submitted in 2020, 15 are still pending decision, representing a total of 57,9 million DKK.

Figures 2.10 and 2.11 show where funding was applied from, and granted, in 2019 and 2020.

Figure 2.10 show that the increase in funding applications from 2019 to 2020 is caused by a doubling of the amount applied for from Danish research councils, an increase in amount from the EU, and a quadrupling of the amount from private foundations. Figure 11 show that the increase in granted funding from 2019 to

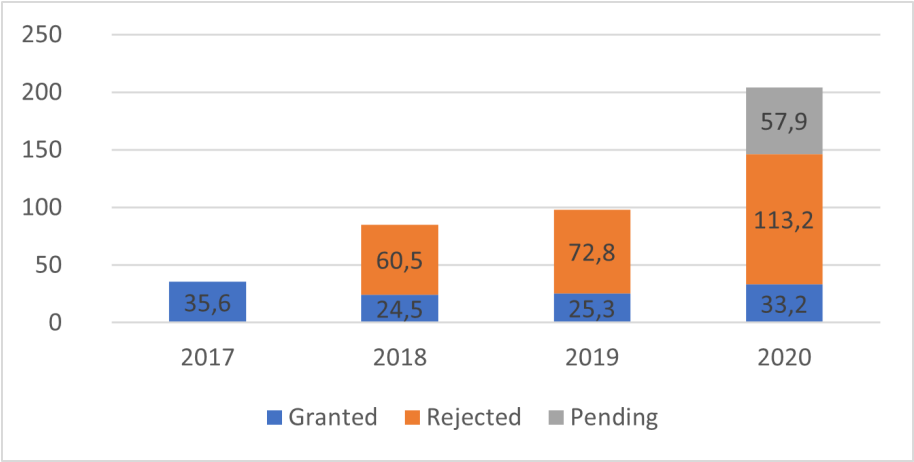

Figure 2.9: Amount of funding applied for (million DKK).

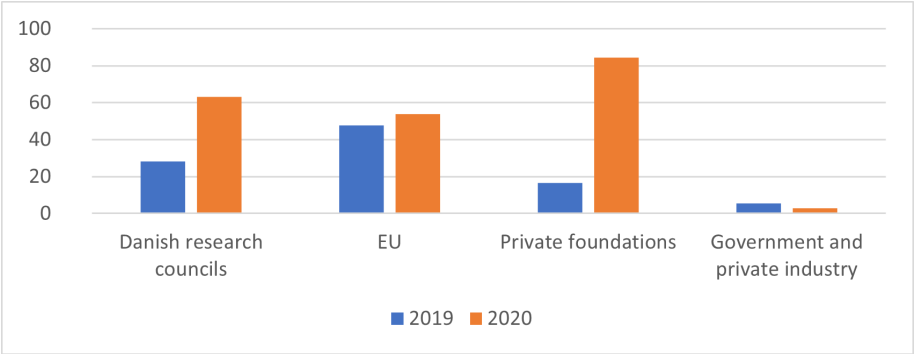

Figure 2.10: Funding source applied from (million DKK).

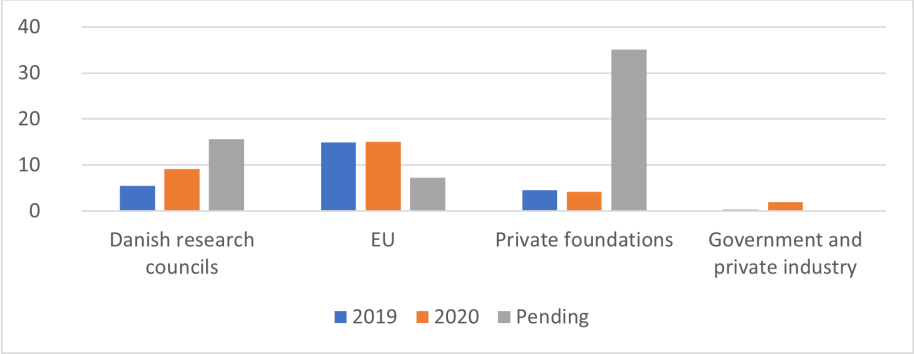

Figure 2.11: Funding source granted or pending (million DKK).

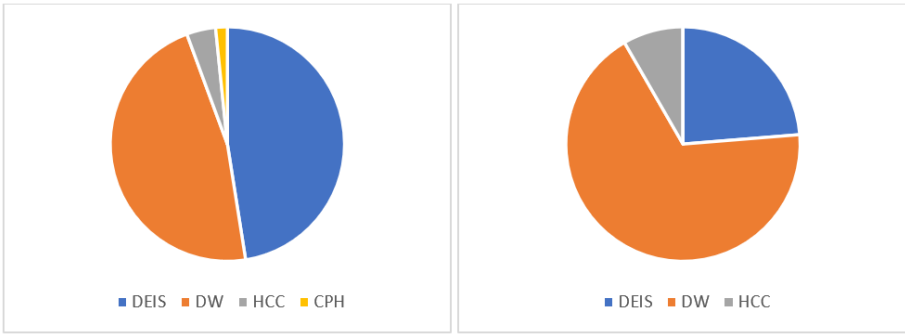

Figure 2.12: Distribution of funding applied for (left) and granted (right) in 2020.

2020 is caused by a better success rate at the Danish research councils, and that a considerable amount of funding applied for is still pending decision, especially from private foundations.

Figure 2.12 shows the distribution of amount of funding applied for (left) and granted (right) on the research groups in 2020. This shows that DEIS and DW has been equally active in applying, with HCC behind, and the team in Copenhagen already beginning to be active. In 2020, DW has had the most success in securing external funding, with HCC having a larger share of granted funding than its share of funding applied for. The Copenhagen team is yet to successfully secure external funding, with all applications from 2020 still pending decision.

2.3.3 Total turnover

As a result of the significantly increased internal income (64%), and a growing amount of external income (15%), the department's total annual turnover has grown by 57% from a relatively steady level around 70 million DKK. in 2016 to 110 million in 2020 (Figure 2.13). As both internal and external funding is expected to grow further in the next 5 years term, it is to be expected that the total turnover will reach 140 by 2023 and have exceeded 150 million by the next research evaluation in 2025 (36% growth over the next 5 years).

2.4 Staff

From 2018 onwards, the department's increased student population and associated growth in internal income has both necessitated and allowed for an active recruit-

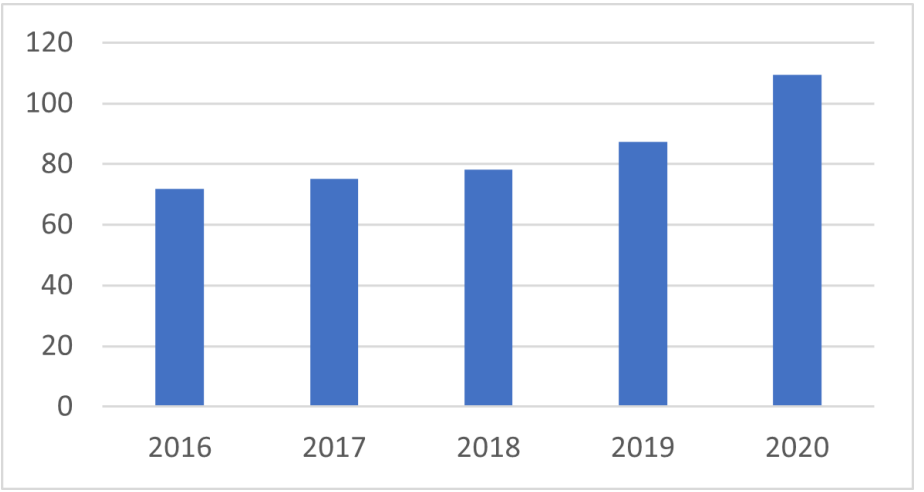

Figure 2.13: Total turnover (internal and external funding spent) in million DKK.

ment strategy. This has led to a 75% growth in staff in only two years, from 84 people in 2018 to 147 people in 2020 (Figure 2.14). The number of scientific staff has grown from 70 in 2016 to 124 in 2020 (77%) and the number of support staff from 15 in 2016 to 23 in 2020 (53%).

The development of numbers within the different staff categories is shown in Figure 2.15. This shows that especially the number of Assistant Professors has increased drastically, from 5 to 21 (420%). This is a direct result of a decision from 2019 onwards to recruit young talents at this level. Also notable is that the total number of Full Professors / Professor MSO has doubled (with Full Professors up by 57%

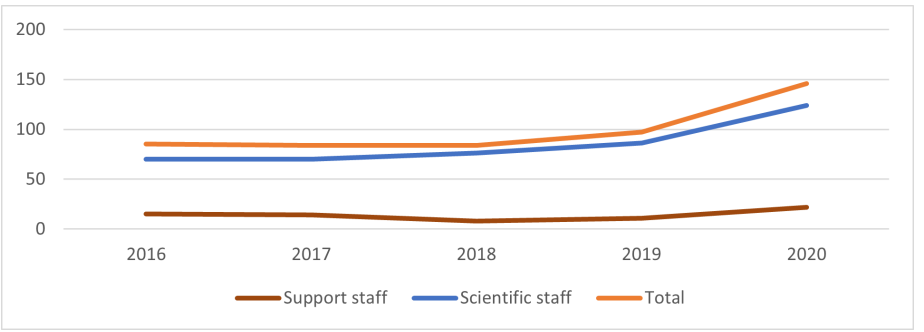

Figure 2.14: Number of scientific and support staff.

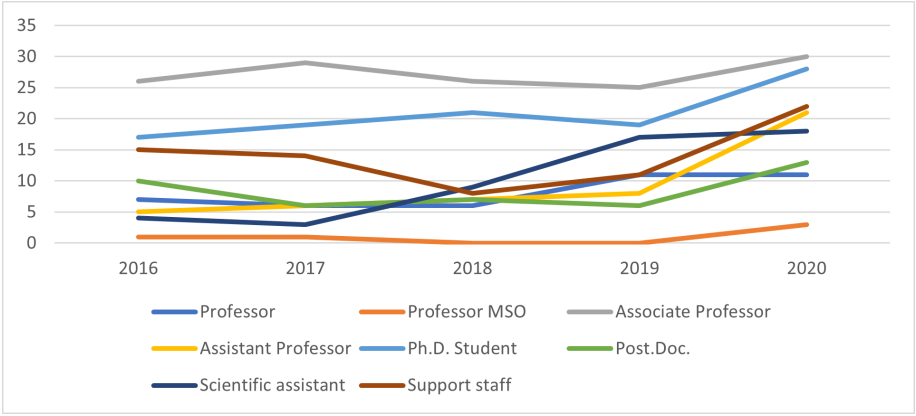

Figure 2.15: Number of staff, by different staff categories.

and the number of Professor MSO up by 150%). This is a direct result of a decision made in 2018-19 to strengthen the departments research leadership capacity, and subsequently promoting a number of highly successful Associate Professors. The number of Associate Professors remains relatively stable (15% increase). The number of PhD students has increased by 58%.

The increase in support staff from 15 to 23 is due to the department's increased level of activity, as well as creation of new departmental support functions in response to new tasks and challenges facing us. The upscaling of support staff includes areas of communication, marketing, diversity, fundraising, finances, research lab support, and local IT support.

The current division of staff on groups is depicted in Figure 2.16. This illustrates the relative sizes of the department's three research groups, support staff, and management. DEIS and DW are of equal size, with HCC roughly half their size. Looking at tenured scientific staff, DEIS accounts for roughly half the department, being 17, with DW being 10, and HCC 8.

Figure 2.17 breaks down the current scientific staff in each research group by type of position. This shows that the size of DEIS is influenced by having many Associate Professors, while the size of DW is influenced by having many Ph.D. students and postdocs. Relatively to its size, HCC has many Assistant Professors and Ph.D. students, reflecting a recent strategy to recruit mainly at junior levels.

Figure 2.18 show the gender balance. This illustrates one of the challenges facing the department, namely the overrepresentation of males among scientific staff. Among the tenured staff, only 11% are female. Among PhD students, the ratio of

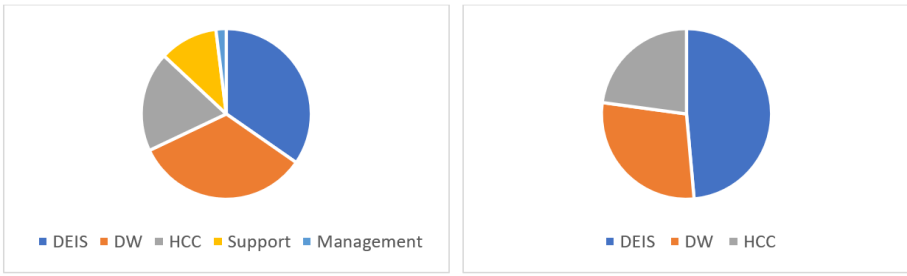

Figure 2.16: Division of all staff (left) and tenured scientific staff (right) (December 2020).

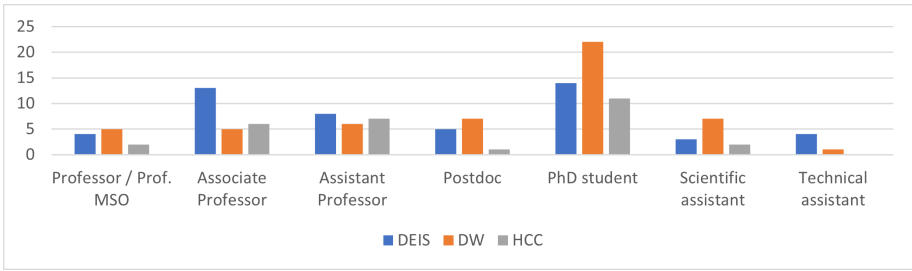

Figure 2.17: Distribution on staff categories in the research groups (December 2020).

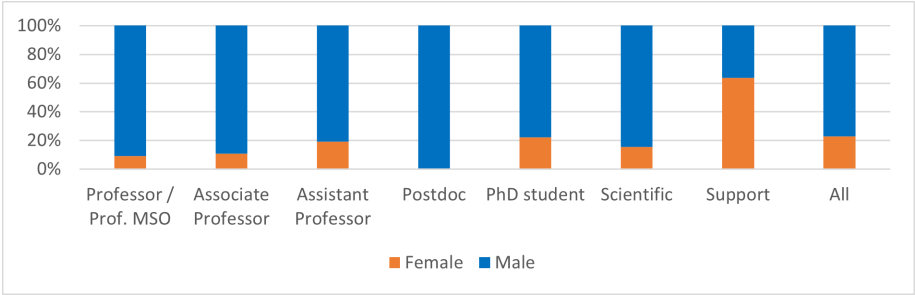

Figure 2.18: Gender balance in different staff categories (December 2020).

females is doubled (22%). This indicate that a proportion of females embarking on an academic career don't make tenure. It is worth noting that the department currently has no female postdocs.

2.5 Publications

Parts of the internal funding is based on publications, as measured by the Danish bibliometric research indicator (BFI). This indicator assigns points to publications in recognized outlets, ranked by reputation. Conference papers, journal articles, and book chapters are given 1-5 points. For each point, the department is granted an amount (currently 21.000 DKK.). Figure 2.19 show the BFI points obtained in the current 5-year research evaluation period, along with numbers for the previous 5-year period for putting recent numbers into context. The numbers for 2020 are forecast, and not yet finalized.

As can be seen from Figure 2.19, the amount of BFI points varies between 160 and 210, with 4 out of 5 years in the current evaluation period (2016-20) coming out lower than the end level of the last evaluation period (2011-15). This is an unfortunate trend and indicates that the scientific staff have less time to do research (and publish) than previously, which would correlate with years of higher teaching load. The forecast numbers for 2020 indicate that the downwards trend is now flattening out. In the next years we expect the number of BFI points to be on an upwards trajectory. This is important, also economically, as reaching the level of productivity of 2014, 15, 18 would release 1.3 million DKK / year of internal funding, which would be enough to finance, for example, two Assistant Professors.

Two factors are expected to increase the number of BFI points obtained in the future. Firstly, the growth in scientific staff should in itself lead to more publications

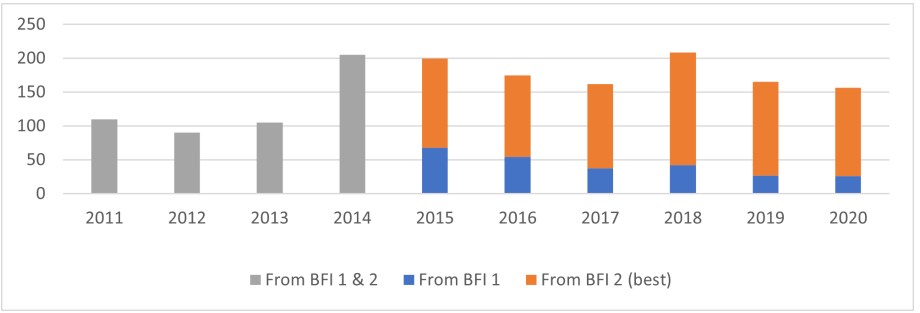

Figure 2.19: BFI points obtained. BFI Level 2 covers the highest ranked outlets. Prior to 2017 no record was kept that separate the source of points from level 1 or 2.

in total, and the focus on attracting quality highly qualified candidates should sustain the trend of Level 2 publications over level 1. Secondly, in 2020 the teaching load for Assistant, Associate and Full Professors, was lowered by 5% This will allow more time for research. However, as there is a considerable lag between the time of planning and doing the research, publishing the results, and the publication releasing BFI points, the full effect of these two factors may first be seen from 2022 onwards. It is also likely, that the time required for applying for more external funding will have a limiting effect on scientific production, restraining the level of BFI points obtained in the short term, until the publication effect from increased research project activity is achieved.

Looking at the years from 2015 onward where the registration of BFI points separated between level 1 and level 2, it is clear that publications are predominantly, and increasingly, level 2 (best). This is a good trend, showing a priority of quality over quantity. The number of publications in BFI level 1 and level 2 outlets can be seen in Figure 2.20. This shows that 44-63% of the department's publications are in BFI level 2 outlets. For comparison, the equivalent numbers for Aalborg University as a whole are 25-36%.

Figure 2.21 show the number of BFI points per full-time employee (FTE) at Department of Computer Science, with the average level at Aalborg University as reference. This shows that the department's historical performance well above average is receding.

The appropriateness of the BFI system for (partially) distributing internal research funds has been debated, as it does not take into account the variances in publication tradition between different areas of research. For the area of Computer Science, it

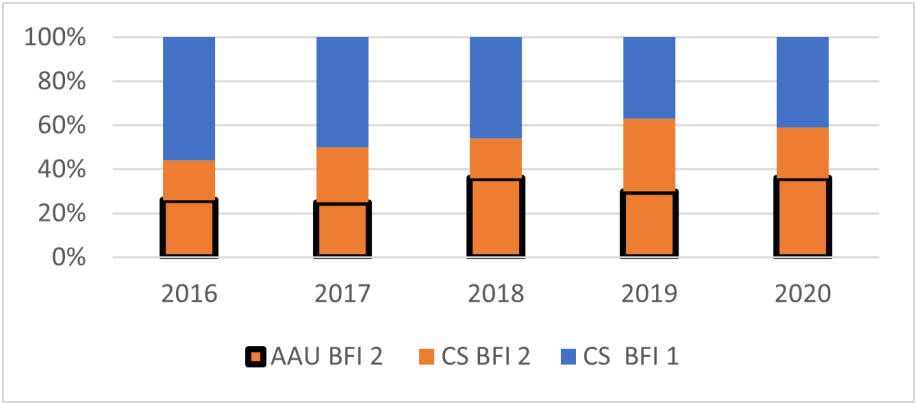

Figure 2.20: Percentage of publications in BFI level 1 and 2 (with AAU level as reference).

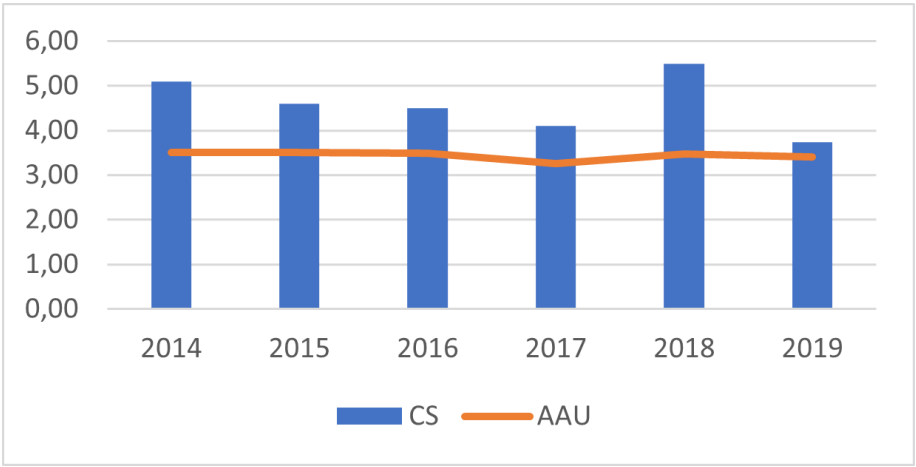

Figure 2.21: BFI points per FTE at Computer Science (with AAU level as reference).

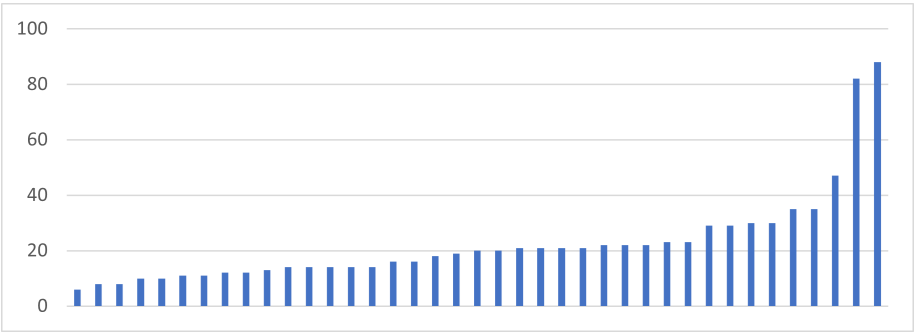

Figure 2.22: H-indices for tenured staff (Google Scholar, December 2020).

has been essential to have full papers at conferences recognized at the same level of journal articles. It is also essential that the ranking of relevant conferences is kept up to date in the BFI system, for example when conferences grow in reputation over time.

2.6 H-Indices for Tenured Staff

The h-index provides an indicator of research impact based on number of citations and publications. All scientific staff at the Department of Computer Science maintain a Google Scholar profile which, among others, calculates this index. Figure 2.22 show the h-index profile for the department in terms of tenured staff as of December 2020.

The pattern of h-index profile for 2020 is very similar to 2015, but with numbers at a higher level, now peaking very close to 90. This of course partly happens because the index is accumulative, but it is also influenced by new researchers joining the department and, of course, by sustained productivity of new research publications that are being cited.

In 2015, we had 3 researchers with an h-index above 30. This has now grown to 5 researchers, closely followed by another 4 researchers at 29 or 30. Where in 2015 we had 12 researchers with an h-index above 20. This has now grown to 18 researchers.

Most notable, of course, is the h-index of Professor Christian S. Jensen (88) and Professor Kim G. Larsen (82), who remain the highest-ranking Computer Science researchers in Denmark. Still towering over the rest is Professor Torben Bach Pedersen (47).

2.7 External Recognition

The department's scientific staff continues to be active on national and international boards, and are regularly recognised with prestigious awards, appointments, and honours. Below are listed some highlights from these.

2.7.1 Awards (by year)

- Professor Torben Bach Pedersen was awarded an Honorary Doctorate (dr. h. c.) from TU Dresden, Germany in 2020.
- Professor Katja Hose was named Top-1 female European researcher in Knowledge Engineering according to AI 2000 top-50 most influential scholars in 2020.
- Associate Professor Hans Hüttel received the CONCUR Test of Time Award in 2020.
- Professor Christian S. Jensen received the IEEE TCDE Impact Award in 2019.
- Professor Bin Yang was awarded Sapere Aude Research Leader by Independent Research Fund Denmark in 2018.
- Professor Katja Hose was awarded Sapere Aude Research Leader by Independent Research Fund Denmark in 2018.
- Professor Christian S. Jensen was awarded the Order of The Dannebrog (Ridder af Dannebrog) by Her Majesty Queen Margrethe II of Denmark in 2016.
- Professor Kim G. Larsen was awarded the Grundfos Prize in 2016.

2.7.2 Appointments (by year)

- Professor Christian S. Jensen was appointed as member of the Scientific Advisory Board, Max Planck Institute for Informatics, Germany in 2019.
- Assistant Professor Daniel Russo was appointed to the ACM Future of Computing Academy in 2019.

- Professor Jesper Kjeldskov was appointed Member of the Danish expert panel on Digital Europe in 2019.
- Professor Jesper Kjeldskov was appointed Chairman of the board for Vesthimmerlands High School in 2019.
- Professor Christian S. Jensen was appointed member of the Board of Directors, VILLUM FONDEN in 2018.
- Professor Kim G. Larsen was appointed Foreign Expert of China in 2018.
- Professor Kim G. Larsen was appointed Distinguished Professor at Northeastern University, China in 2018.
- Professor Kim G. Larsen was appointed INRIA International Chair in 2016.

2.7.3 Honours (by year)

- Associate Professor Peter Dolog became an ACM Senior Member in 2020.
- Associate Professor Lone Leth Thomsen became an IEEE Senior Member in 2019.

2.8 Ambition and goals for Department of Computer Science

In 2020 the department went through a process of defining its ambition, goals and strategy for the next 5 years. This was facilitated by an external consultancy and involved two department workshop days for scientific and support staff in August and November. After both workshop days, the department leadership group merged and summarized the input provided from the workshops. The result (Figure 2.23) describes the department's ambition in relation to research, education, engagement with society, and the workplace. This is responded to with goals and actions for the medium term (2-3 years). The next step of this process is to define goals and actions for the immediate future (1 year), following the research evaluation.

2.8.1 Research

We are recognised internationally and nationally for research in selected areas of computer science at the highest level of excellence according to community standards, prioritizing quality over quantity. Our research is recognised for balancing

Ambition
We are a highly regarded Computer Science Department in Denmark among researchers, students, and external partners. These are attracted to us because of our leading **research**, **education**, and **engagement with society**, and our positive **workplace**. These impact the research communities we are part of, the students we educate, the society we engage with, and the way we work and collaborate in the department. We are internationally leading in the research areas we focus on, successful in acquiring external funding, and our results make a difference to society. We are recognised by students and employers as a high-quality institution with relevant and highly recognised studies, and our graduates do great work.

Figure 2.23: The department's ambition in relation to research, education, engagement with society, and the workplace.

applications and theory and for its societal and scientific impact. We apply research to solve real-world problems, while also conducting foundational research with potential long-term impact. We are known for creating and building on synergies between education and research and for collaborating closely with external partners.

Within the next 2-3 years, we will:

1. Prepare for mission driven research by creating new transformative knowledge and technology that support a global society in balance, with humans at the center.
2. Develop an organization that can identify new strategic synergies and collaborations, and act agile and opportunistically in response to funding opportunities.
3. Demonstrate excellence in large strategic application oriented as well as foundational projects, contributing to selected sustainable development goals.
4. Demonstrate, by concrete examples, how our competences and technologies are core for solving societal challenges.
5. Submit, and have granted, a higher number of project applications, carefully targeting our competences and strongholds toward suitable calls and funding bodies.
6. Have established a well-known and respected cross-group research team in the area of artificial intelligence and machine learning (AIML).
7. Have established two more cross-group research teams, in areas of strategic importance, and two new groups in Web Data, and in Security.

2.8.2 Education

We are recognised for high-quality, problem-based educations in the broader area of computer science. We are recognised for producing highly skilled graduates with very high employability, who are aware of their competences, experienced with project work on solving real-world problems in diverse teams, and able to think critically about the technologies they develop. This makes us an attractive institution of education for students, and an attractive source of graduates for employers. We are known for continuously matching our portfolio of educations to the needs of employers and society as well as students' interests.

Within the next 2-3 years, we will:

1. Educate entrepreneurial graduates who can facilitate increased competitiveness and innovation in solving society's challenges.
2. Develop an increased awareness of evolving societal and technical needs, to make fast adjustments to our curricula and improving the employability of our graduates.
3. Make our master's programs more attractive, with more research-based teaching, electives, extracurricular activities, and projects with external partners to solve real-world problems.
4. Attract more of our master's students to continue in the Ph.D. program, involving them in research, and increasing the number of industrial Ph.D. scholarships.
5. Improve the gender balance on our most technical educations, emphasizing the use of deep IT skills to solve real world problems, in both student projects and in our PR.
6. Improve the student intake on the Data Science education to around 40 per year, focusing on what Data Science can be used for, and what job opportunities it gives.
7. Develop a portfolio of continuing education offers supported by digital platforms, making them visible, accessible, and flexible for our alumni and others.

2.8.3 Engagement with society

We are acknowledged and respected for close collaboration with industry in applying research to create impact and value for society, in particular within the area of sustainable development. We are recognized as an institution capable of engaging in agile and interdisciplinary projects, creating and building on synergies between outstanding applied and foundational research and outstanding problem-based ed-

ucation. This makes us a desirable computer science collaborator in Denmark and internationally, and a go-to partner for working on important societal challenges (e.g., SDGs).

Within the next 2-3 years, we will:

1. Establish strategic partnerships with external collaborators, for transferring our knowledge and technology into real world solutions to societal problems (e.g., SDGs).
2. Establish a strong awareness of the department's applied and foundational research in relation to solving societal problems (e.g., through our Industry Ambassador).
3. Enable all scientific staff members to engage and collaborate with external partners in their research, focusing on solving societal problems.
4. Define student project themes on 2-3 semesters of all educations to focus on solving societal problems with external partners, producing proof-of-concept prototypes.
5. Form a cross-disciplinary team to work with consortia partners on a very large grant application (national or EU) in the area of SDG (e.g., Sustainable Mobility).

2.8.4 The workplace

We strive for diversity in the workplace, for maintaining an inspiring physical work environment, and for ensuring a work-life balance that contributes to ensuring healthy and content employees. Our staff enjoy going to work, interacting with colleagues, and solving tasks in collaboration. Work is done as a team effort, is inclusive, and we always communicate in a constructive and respectful tone. Our ability to achieve excellence is bound by staff and resources. Our staff work hard without overworking themselves. We strive for diversity in terms of nationality, culture, and gender. We recruit internationally and are active in pursuing a gender balance, on all levels.

Within the next 2-3 years, we will:

1. Recruit and retain the right scientific staff to equally teach and do research, with an enrollment process that transfer culture and skills to new employees.
2. Grow the Copenhagen CS group to a sustainable size, ensuring enough people to deliver the required teaching, and critical mass to deliver research results.

3. Improve the gender balance among scientific staff, systematically attracting more women to apply for positions by using insights gained from the diversity project.
4. Nurture talent and support career opportunities at all levels of employment, through start packages, mentoring programs, seed funding etc.
5. Develop a workplace culture characterized by a sense of community, shared understanding and mutual respect among all colleagues.
6. Include more junior staff to participate actively in developing the department's research, education and administration, with a 10-year perspective.
7. Support a healthy work-life balance that prevents stress and overworking, including recognition of working "smart", not just hard, and planning time for deep work.

2.9 Committee Evaluation: The University, Faculty and Department Levels

The report "Research Evaluation 2016-2020" as well as its Appendices authored by the Department of Computer Science are quite extensive and informative. The on-line presentation on "Research Evaluation of the Department of Computer Science" and the discussions on Monday, 12th of April 2021, Friday, 16th of April 2021, as well as the final discussions on Friday, 23rd of April 2021 were very informative and helped to clarify issues. Based on these sources our observations and recommendations are as follows.

Observations As a result of the split of the Faculty of Engineering and Science in 2016 into two smaller ones, the new Technical Faculty of IT and Design was created in 2016, i.e., right at the start of the evaluation period. The Department of Computer Science is now part of this new Technical Faculty of IT and Design. The Department of Computer Science is well managed by the Head of Department (Professor Jesper Kjeldskov), who is supported by the Deputy Head (Associate Professor Uffe Kjærulff and Associate Professor Ulrik Nyman in 2020) and the Head of Administration (Helle Westmark). In addition, in 2020, i.e., at the end of this evaluation period, a Department Leadership Group was established to oversee the operations of the department. The department runs several study programs at B.Sc. and M.Sc. level as well as a Ph.D. program. The department is organized in a matrix structure with three core research groups, and two cross-teams going across these groups performing collaboration activities. These groups and cross-teams are separately discussed in the evaluation paragraphs below.

Our most important observation is that the Department of Computer Science performs excellently in Denmark, Europe and world-wide in terms of publications, visibility, industrial collaboration, start-ups, and societal impact with respect to the selected SDGs. Compared to the previous evaluation period, all key performance indicators, like numbers of students, staff, funding, publications, citations, awards, and honors (just to name a few), evolved in a positive direction. In addition, the department's financial situation improved significantly due to increased funding success and increased number of students, as well as due to a new economy model at the university. This positive situation will probably stay because the department has installed a strategy towards increased external funding due to new specific staff positions for fund-raiser and industry-ambassador. A very effective instrument for flexibility in interdisciplinary collaboration seems to be the matrix organization of the department groups and cross-teams (AI/ML and Copenhagen campus). A further proof of excellent interdepartmental collaboration between all its groups can be exemplified by the key projects DiCyPS and DIREC.

Current and active investments of the department in research-supporting laboratories with permanent staff and facilities are needed for further evolvement with respect to all criteria mentioned above. Here the department has actively taken the first important step that should be further sustained by faculty or university investments. A very good idea of the departmental strategy is to have SDG-focused activities (e.g., energy or transportation issues) becoming cross-cutting themes that can easily be supported and rolled out by the department's matrix organization. This should naturally lead to high impact of the research results.

In the short run, the Department of Computer Science is on a very positive and successful track. However, in the medium to long run the department must counteract several challenges. The most important challenge is to keep the high level of research and research outcome. Currently, this critically depends on a few key persons. The integration of constantly incoming new staff, as well as the substitution of constant leaving (key) staff is key, here. In this respect the integration of assistant professors/tenure trackers due to the large increase in numbers is the current issue and task to solve. Another important challenge is to ensure that Computer Science competences are truly included in interdisciplinary activities with other departments at AAU. Diversity with respect to, e.g., gender, nationality, culture, and inclusion is another challenge that will absorb resources, but finally should pay off. In solving these tasks and counteracting the challenges, the Department of Computer Science must be supported by the Technical Faculty of IT and Design and by Aalborg University with the necessary resources.

The list of recommendations given below should help in getting there.

University and Faculty Level Recommendations Digitalization of the society is important and should also be considered at the university level.

Digitalization is an important future for research and, furthermore, the digital transformation of the university should be research-based. There is a huge potential in using data and learning analytics to improve your education and pedagogics. Some activities are happening on digital teaching but there is clearly room to use what AAU is good at and work on the digital transformation of that, such as digitally assisted PBL. Open science and open data are an important goal and require conscious digitalization processes. The university needs a strategy for digital transformation with a focus on both education and research, and the CS department can and should play an important role in this transformation.

Department Level Recommendations Digitalization is at the core of a Department of Computer Science, and the department's groups and cross-teams are all more or less involved in digital transformation either in teaching or in research projects or even in both. With this involvement and the potential of the topic, the department must build a strong strategy to support digital transformation in various ways.

As many new staff members have joined the department during the last couple of years, time and effort is necessary for their smooth integration into the department's life and making the high aims of the department (research- and teaching-wise) their issue. For example, assistant professors and tenure trackers should be involved in strategic discussions and be coached for future positions (also with assistant professors that you do not expect to stay in your department, make sure they are well-equipped for applications elsewhere); for assistant professors the expectation management should be improved.

Another related issue is the difference in group sizes, which creates the risks that some groups overshadow the other ones. Currently, the performance at department level and at group level is top, and therefore we see no need to address this point urgently. However, the department will face challenges after a few years as senior staff will retire and new leaders will appear, most likely, with different visions and priorities. Therefore, the department should well in advance start working towards building the future leaders from inside as well as to attract mature scientists to replace senior staff after they leave.

In an overall setting, all the groups of the Department of Computer Science as well as the department itself should develop a more coherent long-term vision on research and focus on nurturing future leadership.

Both cross-teams, AI/ML and Copenhagen campus, need some attention. For the cross-team AI/ML the department should develop a stronger vision and strategy on how you would like this team to evolve and improve communication (on the Web and in the reports) about results in this topic. An explicit evaluation in 5 years is recommended. For the cross-team Copenhagen campus the department should improve communication about this and develop a vision and strategy on how you would like to strengthen this branch.

The department has installed a strategy towards increased external funding due to new specific staff positions for fund-raiser and industry-ambassador. Both provide necessary administrative support, but this should be complemented with measures to improve the scientific quality of grant applications, by systematic coaching during application writing, and considering external support for the writing itself.

Diversity and Inclusion are topics of prime interest at both department and group level. Some diversity-related recommendations are given here: when hiring, all members of staff should actively be involved in scouting for diverse applicants; make vacancy texts broad, do not hire only for a specific expertise, but for an expert in a larger area. With respect to inclusion, all members of staff should receive a training on implicit biases. Gender balance is still a concern (e.g., only male post-docs). Hence, consider how to attract female role-models as faculty and, in addition, with respect to tenure track, think in advance how to handle exceptional circumstances during the tenure track period (pregnancy, paternal leave, long sick leave) and how to address this when evaluating for tenure.

Impact of COVID-related effects should be monitored and considered, e.g., for tenure trackers (discuss the impact of the COVID situation including personal circumstances, such as home schooling and possible consequences for their tenure evaluation), but also for all staff members (make sure to discuss impact of COVID home schooling, impossibility to visit family abroad, transferring to online education and how this impacts personal research results).

Quite easy to achieve is an enhancement of the department and the groups' Web appearances and thus increase visibility internally as well as externally. For instance, the groups may be further described by a list of focus topics that reflect the competences and current work areas more comprehensively.

The success of the research work of the department triggered many open-sourced prototypes and noticeable start-up activities. These activities should be supported by the department as well as at faculty and university levels. Support may comprise infrastructure (labs and staff, compute resources, etc.) and other incentives that ease that kind of technology transfer. A general strategy on start-ups and

technology transfer seems worthwhile.

Alumni and alumni networks may become a game changer with respect to industry collaborations and exchange of personnel. We recommend to 'monitor' your graduates (industry, academia, Denmark, abroad) and to build up an active alumni network that may help when applying for funding as well as to achieve broader visibility and new and sustained industrial collaboration relationships or to provide life-long learning for.

3 Database, Programming, and Web Technologies

3.1 Executive Summary

The Database, Programming, and Web Technologies (DPW) group performs research in data management, data mining, machine learning, programming languages, and Web science. The group hosts Center for Data-intensive Systems, Daisy, that embraces all activities concerning data-intensive systems and data-driven value creation. Daisy aims to be an internationally leading research center within its areas of specialization.

The group performed very well during the period according to key performance indicators, including publications, external funding, Ph.D. education, scientific service, and industrial collaboration and dissemination.

The group has published 338 refereed journal and conference papers in the period, of which 136 are in the very top outlets within data management and an additional 84 are in top specialized outlets relevant to Daisy, which ranks the group highly in the database area in Europe.

The group has participated in 46 externally funded projects and has received some DKK 107 million in funding from a wide range of sources. The project portfolio encompasses substantial industrial participation, and many projects involve other groups in the department or the university. The recent Digital Research Centre Denmark, DIREC, will enable the group and the entire department to engage in new collaborations with Danish CS departments.

Two spinout companies have been established. Several software packages and tools have been released as Open Source. The group has trained new researchers, with 19 Ph.D. degrees having been awarded and 18 Ph.D. studies in progress. DPW researchers have received substantial peer recognition and have performed extensive service to the scientific community. The group was successful in attracting top

scholars for positions in data management and machine learning but less successful in attracting and retaining new staff in programming technology. In September 2020, the sub-group on programming technology ceased to be part of the group. To increase cross fertilization and collaboration, the majority of its members joined the DEIS group. In conclusion, the DPW group believes it has demonstrated excellent research performance in the reporting period. In particular, the DPW group believes that its publication performance in top outlets and level of industrial collaboration compare favorably with those of similar groups across Europe.

3.2 Profile

The research in the DPW group covers a wide range of areas within data-intensive systems and programming technology. The group's research approach is primarily constructive: theoretically well-founded, purposeful artifacts such as frameworks, data structures, indexes, algorithms, languages, tools, and systems are prototyped and subjected to empirical study. Thus, the research often involves labor-intensive software development and experimental studies.

Further, the research is mostly driven by real-world applications, the objective being to provide solutions that enable new or better applications. The primary application areas include intelligent transport systems, logistics, energy grids, healthcare, and web querying.

In *data-intensive systems*, substantial research concerns aspects of what is now often referred to as big data. Central research topics belong to three areas:

- **Temporal, spatial, spatio-temporal, spatio-textual, and metric data management and analytics** This has been an active area of research within the group since its creation, and a substantial fraction of the group's research continues to fall into this area. While some research on temporal data management continues, the focus has been on spatio-temporal and spatio-textual data management. Here, predominant spatio-temporal topics include trajectory analytics, traffic analytics, vehicle routing, and outlier detection. Much of the research in spatio-textual analytics concerns the enabling of new, efficient query functionality related to data with textual and spatial attributes, including microblog posts, reviews and rating, and other web content with geo locations. Finally, the group has embraced the indexing and efficient querying of metric data. During the review period, machine learning has become an increasingly important aspect of this research.

- **Deep and extreme-scale analytics** The group has a long history within data analytics, starting with data warehousing, OLAP/multidimensional databases and business intelligence in the late 90s, later adding ETL, specific data mining topics, cloud-based Big Data analytics, and semantic web analytics. Most recently, we are also investigating deep (predictive and prescriptive) analytics and extreme-scale analytics on massive IoT time series. Our main, deep application domain is within digital energy where we focus on energy flexibility and massive energy time series.
- **Semantic data management, graph databases, knowledge engineering, text and natural language processing, recommender systems, and social networks** The group has a continued focus on semantic data management, recommender systems, social network analysis, and expanded the scope to graph databases, knowledge engineering, and natural language processing. Research in these areas covers fundamental problems of query processing, knowledge and data exploration, semantic data warehousing, data integration, natural language processing as well as understanding user preferences in spatial contexts and in recommender systems in deep neural networks with attention mechanism, history, and adversarial approaches.

3.3 Staff

3.3.1 Current Staff

Professors: Christian S. Jensen, Torben Bach Pedersen, Katja Hose, Bin Yang, Hua Lu (MSO, 20%), Kristian Torp (MSO)

Associate Professors: Chenjuan Guo, Christian Thomsen, Peter Dolog, Simonas Šaltenis (20%), Gabriela Montoya

Assistant Professors: Dalin Zhang (tenure track), Daniele Dell'Aglio (tenure track), Huan Li, Ilkcan Keles (on leave), Jilin Hu (tenure track), Johannes Bjerva (tenure track), Kaixuan Chen (tenure track), Nguyen Thi Thao Ho, Yan Zhao

Post-docs: Matteo Lissandrini, Hamdi Ben Hamadou, Bijay Neupane (20%), Tomer Sagi (20%), Daniel Hernandez, Amir Laadhar

Ph.D. Students: Bhuvan Gummidi, Christian Aebeloe, Emil Riis Hansen, Jonas Brusokas, Kashif Rabbani, Mohsin Iqbal, Olivier Pelgrin, Razvan-Gabriel Cirstea, Sean Bin Yang, Simon Aagaard Pedersen, Suela Isaj, Theis Erik Jendal, Tiantian Liu, Tianyi Li, Tung Kieu, Van Ho Long, Ye Yuan, Xinle Wu

Research Assistants: Carlos E. Muniz Cuza, Kasper Fromm Pedersen, Magnus

Nørhave Hansen, Miroslav Pakanec, Roshni Chakraborty, Søren Kejser Jensen, Fabio Lilliu, Tobias Skovgaard Jepsen

Academic Technical Staff (AC-TAP): Laurynas Siksnys (40%)

3.3.2 Staff Development

Over the period, the staff evolved substantially around a stable core. In 2018, Katja Hose and Bin Yang were promoted to Professor. In 2019, Hua Lu, Bent Thomsen, and Kristian Torp were promoted to Professor MSO, and Thomas Bøgholm, Chenjuan Guo, and Gabriela Montoya were promoted to Associate Professor.

A number of professors, associate professors, assistant professors, and postdocs joined the group or left it during the period. Towards the end of the reporting period, the Programming Technology sub-group – comprising Professor MSO Bent Thomsen, Associate Professor Lone Leth Thomsen, Associate Professor Kurt Nørmark, Associate Professor Thomas Bøgholm, and Associate Professor Hans Hüttel – ceased to be part of the group. The former four joined the DEIS group. Alexander Asp Bock joined the group as a postdoc in 2020 for six weeks. Lu Chen joined as Assistant Professor in 2017, was promoted to Associate Professor in 2019, and left in 2020 for Zhejiang University. Muhammad Aftab joined the group as a postdoc in 2018 and left in 2019 for MHI Vestas Offshore Wind. Magnus Madsen joined the group as Assistant Professor in 2017 and left in 2018 for Aarhus University. Bolong Zheng joined the group in 2017 as postdoc and left in 2018 for Huazhong University of Science and Technology. Gönenç Ercan joined the group as a postdoc in 2018 for four months while being on sabbatical from Hacettepe University. Panagiotis Karras joined as Associate Professor in 2016 and left in 2018 for Aarhus University. Vinay Jayarama Setty joined the group as Assistant Professor in 2016 and left in 2017 for University of Stavanger. Luis Antonio Galarraga Del Prado joined as a postdoc in 2016 and left in 2017 for INRIA/IRISA. Xike Xie joined the group as Assistant Professor in 2012 and left in 2016 for University of Science and Technology of China. Benjamin Krogh joined the group as postdoc in 2015 and left in 2016 for Blip Systems. At the beginning of the reporting period, two members of the former Machine Intelligence group – Associate Professor Peter Dolog and Assistant Professor Nattiya Kanhabua – joined the group. The latter left in 2017 for NTent Barcelona. In addition, Eleftherios Zervakis, Thomas Frisk Olsen, Guangyue Xu, and Dalia Kaulakiene were employed as research assistants.

A large number of Ph.D. students and research assistants pursuing Ph.D. degrees were employed during the period, but have already graduated or left the group: Alvis Logins, Davide Frazzetto, Emmanouil Valsomatzis, Faisal Moeen Orakzai,

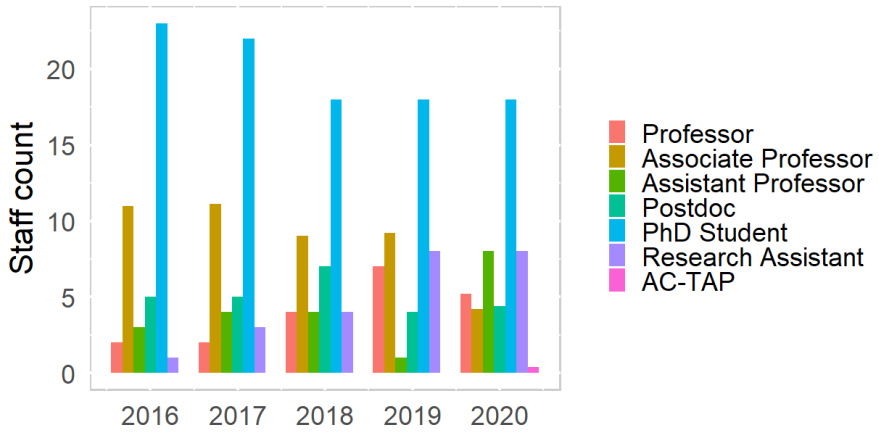

Figure 3.1: Staff count per year, weighted by percentage of employment

Felipe Soares da Costa, Jacobo Rouces Gonzalez, Johannes Lindhart Borresen, Jovan Varga, Kim Ahlstrøm Meyn Mathiassen, Muhammad Aamir Saleem, Lawan Thamsuhang Subba, Nurefsan Gür, Olga Rybnytska, Ove Andersen, Robert Waury, Rudra Pratap Deb Nath, Shumet Tadesse Nigatu, Jahan Lisa Nusrat, and Ying Wang. We also note that we hosted a number of visitors during the period, including Ph.D. students that were supervised by members of the group during their visits.

Figure 3.1 shows the staff count per position category at the end of each year in the reporting period. The numbers of Professors, Assistant Professors, and Research Assistants have increased considerably during the reporting period. The considerable decrease in the number of Associate Professors is due mainly to the Programming Technology sub-group leaving the group towards the end of the reporting period.

Figure 3.2 shows the gender distribution per position category at the end of the reporting period. The group has higher female faculty members in most categories compared to the department average. For example, the group has the only female professor in the department, 2 female associate professors, and 3 female assistant professors.

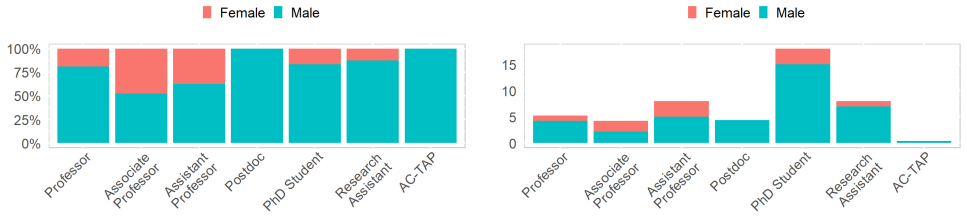

(a) Gender percentage

(b) Gender count

Figure 3.2: Gender distribution per category, weighted by percentage of employment

3.4 Goals 2016–2020

The detailed plan for the evaluation period may be found in the report on the previous evaluation period. The group conducted both fundamental and applied research and viewed these types of research as equally valuable. During the reporting period, the group focused on excellence and impact, realizing synergies, obtaining laboratory support, and increasing it capacity for seizing new opportunities.

The group covered two large areas: data-intensive systems and programming technology, where data-intensive systems was further clustered into three main topics: (i) temporal, spatial, spatio-temporal, spatio-textual, and metric data management and analytics, (ii) advanced analytics and data warehousing, and (iii) web and social media data management.

Excellence and Impact

The group aimed to focus on excellence and impact in both fundamental and applied research. In terms of fundamental research, performance indicators include publication in top scientific outlets (excellence) and citations (scientific impact), as well as awards, invited talks, and other forms of peer recognition. In terms of application-oriented research, a more diverse and less established range of indicators exists: demonstrations, industry and user funding, industrial uses of results and software, spinouts, software downloads, and system and service users.

Synergies

The group aimed at realizing latent potentials for synergies by (i) identifying and focusing some of its activities around selected, cross-cutting themes and (ii) integrating fundamental and applied research.

Cross-cutting themes existed where group members studied the same types of techniques in different settings or application domains. By bringing together such group members, we expected to leverage each other's insights and to improve the innovative nature of our research.

No single group member can focus on both fundamental research and applied research at the same time. Cross-fertilization between fundamental and applied research was established by means of collaboration among group members who often position themselves differently along the fundamental-to-applied dimension. Those focusing on more fundamental aspects got inspiration for new and relevant research problems and may see their results being implemented in systems that create value for collaborators and society in an accelerated fashion. Those focusing on challenges closer to applications may benefit from the availability of fundamental insights and advanced techniques.

Laboratory Support

The group develops, maintains, supports, and uses increasing amounts of software and an increasing number of systems. This is important for long term research and industrial collaboration and uptake, and it contributes to the visibility of the group. These activities are increasingly labor-intensive and are difficult to sustain via external funding, in part because externally funded projects are relatively short-lived and because they often address new challenges. Thus, the research activities in the group would benefit substantially from longer-term laboratory support. It was an objective to achieve such support.

Ability to Seize Opportunities

The group aimed to leverage its size to maintain and build its capacity for seizing opportunities for achieving further research activities and results in our area. More specifically, the group focused on the ability to compete for funding in areas, where we expect future funding to be available, and the ability to collaborate with industrial partners.

3.4.1 Temporal, Spatial, Spatio-Temporal, and Spatio-Textual Data

At the onset of the evaluation period, we planned to focus on enabling real-time mobility and transportation analytics, which calls for improved real-time data ingestion and online querying capabilities. We therefore expected main-memory technologies to play an important role in our research. Next, while we had so far focused on GPS data from vehicles, we expected to consider also additional data sources and the integrated use of multiple data sources in analytics. Expected outcomes of this research included (i) a state-of-the-art platform for network-based mobility and transportation analytics that supports real-time, on-line analytics, (ii) a new paradigm for trajectory-based routing, (iii) and enablement of a range of new applications developed in collaboration with external collaborators.

Next, we aimed to leverage advances made by group members in the previous reporting period in the area of keyword-based querying of spatio-textual data. We planned to continue activities in this area. One goal was to enable better evaluation of query result quality, e.g., by means of crowdsourcing. Success in this area would enable new studies of advanced ranking functions.

Finally, members of the group were among those scientists who were early to identify and study the unique challenges of enabling indoor location-based services. We aimed to continue research in this area on providing efficient support for functionality needed by indoor services that utilize cloud computing and are delivered to smartphone users. We also planned to study analytics on indoor mobility data, to the extent we were able to gain access to such data.

3.4.2 Advanced Analytics and Data Warehousing

One planned focus was data warehousing and analytics for sensor data such as large-scale fine-grained time series, e.g., from wind turbine sensors or cyber-physical systems. A key characteristic would be storing and querying data in the form of higher-level models, e.g., based on mathematical functions rather than as individual low-level data items. Data ingestion and extract-transform-load was to be investigated, both for historical and streaming data, as well as indexing and query optimization. The developed techniques should be integrated in a scalable cloud-based platform.

Another topic was concerning tighter integration of analytics with decision making and optimization. One aspect of this was scaling multi-criteria decision support techniques such as skyline queries to Big Data, including novel algorithms for cloud-based computing paradigms and novel recommendation approaches apply-

ing dominance analytics to big data. Another aspect was the tight integration of (predictive) analytics with optimization problem solving to prescribe the best course of action, so-called prescriptive analytics. Here, we wanted to investigate how to integrate solvers for different problem classes tightly with data management functionality using a single data model and query and optimization language, and how to optimize such analytical workflows and scale them to big data. Finally, we wanted to continue the successful work on energy data management. The FlexOffer concept should be extended to cover even more advanced cases, e.g., to include (running) state information. New aggregation and analytics techniques for FlexOffers should be developed, along with new techniques for predicting available flexibility based on user behavior and suggesting FlexOffers for it. FlexOffer concepts and techniques should also be integrated and aligned with existing smart grid technologies and (existing and emerging) energy markets.

3.4.3 Web and Social Media Data Management

A focus in this area was to integrate activities on Linked Open Data and multidimensional (OLAP) support for RDF – including modeling and management of metadata, distributed and federated query processing, semantic extract-transform-load, and physical level optimizations – with the overall goal of enabling easy and efficient querying of the Web of Data, covering both "self-service" business intelligence and more general user queries. We planned to consider several novel aspects, such as provenance, quality, and context in the (Linked) Open Data setting as well as techniques for discovering new data sources and determining their relevance and interestingness. We also planned to work on techniques for guidance and recommendation, explaining query results as well as annotating and integrating data sources with additional multidimensional semantics.

Regarding social network platforms, we planned to continue our research on value and relevance of search results and – in collaboration with researchers from humanities and social sciences – to study how to detect interesting human-involved patterns from social data to obtain insights into particular types of transmissible societal problems, such as food poisoning.

3.4.4 Intelligent Web and Information Systems

We planned to focus on three directions of research in this area: Web long tail, computational methods for web user interfaces, and understanding and unlocking the Web. We planned to focus on applying and extending machine learning meth-

ods to improve recommendations by recommending also items from the Web long tail. We planned to study latent factors based methods, graph based methods, reinforcement learning based methods and their combination for this task. We planned to develop novel computational methods for web user interfaces facilitating exploration in news and social media, but also explorative user interfaces for Web querying and Web Search. We planned to also study and develop information filtering methods which go beyond simply observing the Web supporting advanced analytics, time- and entity- aware retrieval as well as context-aware summarization.

3.4.5 Programming Technology

The sub-group on programming technology planned to pursue research on the following topics: interactive and programmatic programming in live programming environments; declarative and interactive, parallel and distributed technologies for Big Data Analysis; and high-level programming of low-level components for embedded systems and the Internet of Things. The sub-group also planned to pursue research on teaching of programming as well as technologies and techniques for supporting teaching.

3.5 Activities and Results

Excellence and Impact

DPW has assessed its performance with respect to several indicators. Most of the publications of the group are in top scientific outlets. Thus, 74.53% of the journal publications are in outlets at level 2 in the Danish national bibliometric ranking. For conference papers, the percentage at level 2 is 66.8%. Compared with the last evaluation period, the group has nearly doubled its level 2 publications, from 123 to 238.

According to csrankings.org that monitors publication of full papers in the top-3 database systems conferences (SIGMOD, PVLDB, and ICDE), which align with the main focus of the group, the group ranks *number one in Europe* in the period, ahead of TU Munich, ETH, and EPFL, as shown in Figure 3.3.

The group's publications are very often cited, for instance, 35 publications from the reporting period have 20+ citations, and 10 publications have 50+ citations, according to Google Scholar.

Figure 3.4 shows the scientific impact of the DPW group, as of the end of the reporting period, it is measured as the Google Scholar h-index.

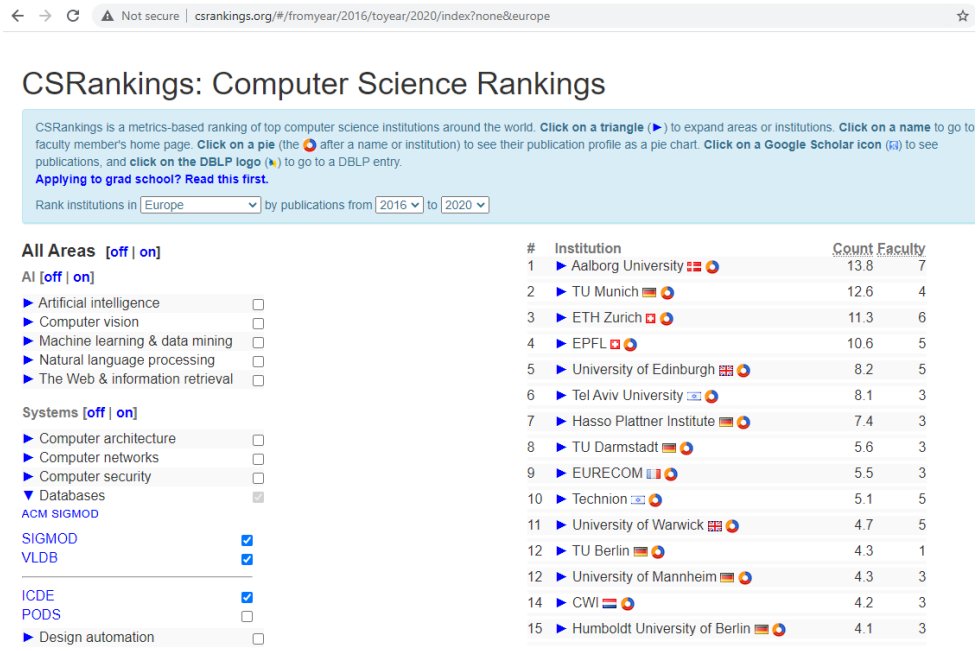

Figure 3.3: Ranking of universities in Europe, from csranking.org, based on full papers in SIGMOD, VLDB, and ICDE during 2016—2020

DPW's members have given 30 keynote talks and 30 other invited talks or lectures. Moreover, DPW's members have been appointed to roles that encompass leadership in the organization of 35 conferences and workshops. Members of the group have also received prestigious awards and recognition in the reporting period, such as the IEEE TCDE Impact Award (Christian S. Jensen), Sapere Aude Research Leader by Independent Research Fund Denmark (Bin Yang, Katja Hose), the Poul Due Jensen Professorate in Computer Science (Katja Hose), Honorary Doctorate (dr. h. c.) from TU Dresden, Germany (Torben Bach Pedersen), AAU Talent Programme (Katja Hose), IEEE MDM Early Career Distinguished Lecturer (Bin Yang), Distinguished Scholar (Bin Yang), Top-1 female European researcher in the field of Knowledge Engineering according to AI 2000 top-50 most influential scholars (Katja Hose), member of the Board of Directors of Villum Fonden (Christian S. Jensen), member of the Scientific Advisory Board of the Max Planck Institute for Informatics, Saarbrücken, Germany (Christian S. Jensen), and Knight of the Order of Dannebrog (Ridder af Dannebrog), appointed by Her Majesty Queen Margrethe II of Denmark (Christian S. Jensen). Additionally, the publications of the

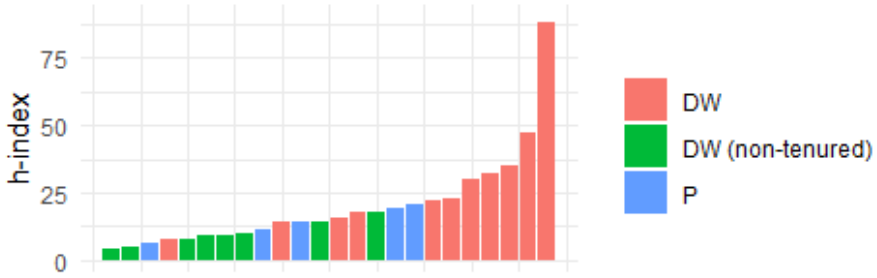

Figure 3.4: Google Scholar h-indices of DPW staff, assistant professor level and above

group have been recognized among the best in venues such as ACM Transactions on Database Systems, ADBIS, AIME, BigData, e-Energy, ESWC, FLAIRS, ICDE, Journal of Computer Languages, PVLDB, WebIST, WWW.

The application-oriented activities of the group are also recognized as valuable. For example the GOFLEX flexibility platform has been demonstrated over multiple years and locations (Wunsiedel, Germany, Sion, Switzerland, Cyprus) reaching more than 410 prosumers. The technologies developed have given place to two spinout companies. FlexShape Aps is based on the AAU-Daisy FlexOffer flexible energy technology, and Modelardata IVS is based on the ModelarDB technology. Additionally, the open source tools developed by the group are used by large number of users from different locations. For example, pygrametl, an ETL tool, is used by over 7400 user from over 14 different countries.

Additionally, the members of group invest a fair amount of time in teaching, and this way, the group achieves a high impact through teaching. Members of the group teach across regular educations offered by the Department, e.g., Software, Computer Science, as well as continuous educations, e.g., IT Vest. Courses taught by the group range from core courses with large numbers of students and specialized courses with a high connection with the research topics investigated by the group. Moreover, members of the group have led the development of the Data Science education.

Synergies

- Within the transportation domain the group has significant experiences with handling GPS and fuel consumption data. This has led to synergies across

projects. As examples, the data management done in the DiCyPS, ELVITEN, and OptiTruck projects are based on the same underlying software platforms, i.e, pygrametl and SimpleETL, both frameworks developed by the group members.

- With the DiCyPS project collaboration with a larger Danish insurance company made it possible to study how to influence drivers while driving. The influence is particular on driving within the speed limit. Speeding is both a major killer in traffic and a large source to extended CO2 emissions as the fuel consumption increases dramatically with speed.

- With the OptiTruck project the focus was directly on reducing the fuel consumption from 40+ tons trucks with 20%. The GPS/CANbus data from truck is new to the group. The project was done in collaboration with Ford Trucks Europe (Ford Otosan), IAV (truck parts), Elidadis (trucking company), and Codognotto (trucking company)

- With the ELVITEN project the focus was on using electric light vehicles (EL-Vs) in cities. These vehicles are mostly two and three wheelers. This represented a number of technical challenges, e.g., with map-matching because two-wheelers behave very differently from cars in traffic.

- Within the energy domain, the group has worked a lot on energy flexibility, including proposing the FlexOffer format and protocols and developing an associated platform for flexibility acquisition, aggregation and management. These contributions were re-used and further developed in more than a handful of funded projects over the period and were the foundation of the spin-out company FlexShape.

- Within data analytics, in particular for energy, the group has significant experience in handling extremely large and fast time series, e.g., from wind turbine sensors. Here, the group developed a powerful platform for model-based data storage and querying which was further developed in several funded projects and the basis of the spin-out company ModelarData IVS.

Laboratory Support

Towards the end of the reporting period, the department invested substantially on laboratory hardware. In particular, a data cluster with eight nodes and two GPU workstations are particularly useful for the group's research activities. Due to the

lack of long-term technicians, the group instead hired many student programmers to maintain software stacks and systems.

Ability to Seize Opportunities

- The Villum synergy programme offers funding opportunities to strengthen the interdisciplinary field of data driven science in Denmark. The programme started in 2019 and the group has two projects that are funded by the programme. One project aims at light-weight AI for renewable energy, more specifically, enabling cognitive power electronics. The other project aims at improving the recovery of microbial genomes and evidence-based analysis by combining DNA sequencing, machine learning, and graph-based analysis.

- The group has worked on digital energy, in particular energy data management, since 2010 and made a range of contributions in the area, including the award-winning FlexOffer concept. Due to the increased focus on the Green Energy Transition within the reporting period, the group was able to leverage this foundation to acquire a handful of large funded projects from EU and Danish agencies and significantly engage with the energy industry.

- For green transportation the group extended to this area by working with new vehicle types. This is the +40 tons trucks in the optiTruck project and working with the many types of light electric vehicles in the ELVITENT project.

3.5.1 Temporal, Spatial, Spatio-Temporal, Spatio-Textual, and Metric Data Management and Analytics

Temporal data management

Some two decades ago, we defined a semantic foundation, based on so-called statement modifiers, for the design of temporal query languages. Teradata later used statement modifiers to offer limited temporal support in their product. In 2016, we (finally!) reported an open-source extension of PostgreSQL in TODS that enables efficient, wholesale support for query languages based on statement modifiers. The group participates in a recently funded project based at Free University of Bozen-Bolzano, Italy that also involves University of Zurich. That project aims to further advance the support for temporal data management and analytics.

Spatial and spatio-temporal data

Some of the subsequent subsections cover contributions related to specific topics within spatial and spatio-temporal data. Here, we cover contributions that go beyond these topics.

We developed indexing techniques for both Euclidean and road-network settings, including learned spatial indexing (CoRR, PVLDB) supporting, point, range, kNN, and range aggregate queries and indexing supporting reachability queries in public transport networks (ADBIS, best paper). Next, we surveyed safe-region based techniques for continuous spatial query processing (Comput. Surv.), we proposed means of spatio-temporal blockchain query processing (Future Gener. Comput. Syst.), we enabled efficient geographic service sharing using imprecise Voronoi cells, and we enabled many-core GPU processing of repeated range queries over streaming moving-object data (Concurr. Comput. Pract. Exp.).

We also proposed a range of machine learning based solutions to problems related to road networks. Notably, we studied network embeddings for machine learning on road networks and proposed so-called relational fusion networks that extend graph convolutional networks to road networks (TITS, SIGSPATIAL, BigData), we proposed spatial network embeddings by extending classic graph embeddings to incorporate spatial properties with multitask learning and employed them to enable context-aware path ranking (TKDE, ICDE), we enabled stochastic-weight completion in road networks by means of graph convolutional networks (ICDE), and we enabled the forecasting of stochastic origin-destination matrices for transportation using recurrent neural networks with graph convolution (ICDE).

Data-Intensive vehicle routing

Increasingly available vehicle trajectory data provides us a solid data foundation to enable data-intensive routing that holds strong potential to improve routing quality. Here, we studied a range of solutions to stochastic routing, where travel costs (e.g., travel time, fuel consumption) of paths are modeled as distributions rather than deterministic values. To capture travel cost distributions accurately, the travel cost distribution dependencies among different edges must be accounted for. To this end, we studied different means, including a path-centric paradigm (PVLDB, VLDBJ, MDM) and a machine learning paradigm (PVLDB, ICDE). We proposed time-dependent uncertain contraction hierarchies to support efficient stochastic routing (VLDBJ) and proposed means to select optimal paths according to different risk profiles while considering the path cost distributions (VLDBJ). In addition, we

study routing preferences learning from trajectory data to enable preference-based routing (ICDE, VLDBJ).

Trajectory analytics

Here, we contributed infrastructure for trajectory analytics as well as support for specific analytics functionality.

Next, we proposed indexing techniques for network-constrained trajectories to support a variety of important queries, including path queries and queries that retrieve travel-time histograms for paths (EDBT, MDM, SIGSPATIAL). We also proposed a path-based API for the analysis of trajectories (SSTD). Further, we provided means of using deep representation learning for trajectory similarity computation (ICDE), and we enabled compression of uncertain trajectories by means of representative trajectories (PVLDB).

Considering next solutions that enable specific analytics functionality, we studied different kinds of trajectory joins, including sequential and parallel similarity joins of trajectories and semantic trajectories (PVLDB, VLDBJ, ICDE) and parallel trajectory-to-location joins. We enabled trajectory search by regions of interest (TKDE). We also enabled non-dominated path retrieval, real-time co-movement mining and pattern detection for streaming trajectories, and convoy pattern mining (MDM, SIGSPATIAL, GeoInformatica, SIGMOD, PVLDB). Further, we contributed means of distinguishing trajectories from different drivers in sets of incompletely labeled trajectories (CIKM), we developed means of using trajectories for analyzing the performance of road intersections (SIGSPATIAL), and we offered means of trajectory segmentation and driving preference mining (SIGSPATIAL workshop).

Spatio-textual data

Here, a substantial fraction of the work encompasses proposals for new kinds of queries, including their efficient implementation that often involves new algorithms and indexing and pruning techniques. Thus, we proposed new query functionality and accompanying efficient implementations for, e.g., why-not spatial keyword queries, semantic region retrieval, co-located community search, similar-region retrieval, event-partner retrieval, location-aware top-k term publish/subscribe, clue-based spatio-textual querying, direction-aware why-not spatial keyword querying, reverse keyword-based location search, reverse geo-social keyword querying, and spatial-textual cluster retrieval. This work is documented in about a dozen papers

in outlets such as CIKM, DASFAA, ICDE, PVLDB, and TKDE. We also initiated work on so-called spatio-textual-temporal cubes that extend OLAP-style analytics and pre-aggregation to complex spatio-textual data.

We pursued the objective of enabling evaluation of the rankings of results produced by different spatio-textual query functionality, dedicating a Ph.D. study to the topic. Specifically, we proposed means of identifying visited points of interest (PoIs) from GPS trajectories and provided means of extracting rankings of PoIs from GPS data. Next, we provided a solution that leverages crowdsourcing to evaluate ranking functions. This work was published in 5 papers in pertinent specialized outlets.

Spatial crowdsourcing

In addition to the work on crowdsourced rankings already mentioned in the coverage of spatio-textual data, we reported a comprehensive literature survey and worked on task assignment for time-continuous tasks and public transport commuters and enrichment of OpenStreetMap tasks. Next, we proposed a solution for price and time aware ride sharing and collective travel planning. We also proposed solutions to demand forecasting and delivery scheduling as well as to pickup and delivery scheduling, resulting in 9 publications in, e.g., ICDE, PVLDB, SIGSPATIAL, SSDBM, TKDE, and TODS.

Indoor location-based services

In the indoor setting, we worked on indexing, keyword-aware routing, time-varying shortest path search, similarity search for trajectories, location privacy, and indoor-outdoor unified spatial queries, resulting in half a dozen publications in PVLDB, ICDE, SIGSPATIAL, DASFAA, and ISA. In the setting of uncertain indoor positioning and tracking data, we worked on data cleansing, on mining dense locations and popular POIs in online and offline settings, and on predicting risks of late arriving objects, resulting in about a dozen publications in outlets such as SIGMOD, ICDE, EDBT, TKDE, GeoInformatica, and MDM. We also built software tools that generate indoor mobility data for real buildings and that enable visual mobility analytics on raw indoor positioning data, resulting two demo publications at PVLDB. Together with an international collaborator, we also gave a tutorial on indoor data management at ICDE.

Time series outlier detection and forecasting

We studied various deep neural network-based approaches to time series analytics, including outlier detection and forecasting. For time series outlier detection, we invented different solutions based on autoencoders, including recurrent neural network based autoencoders and recurrent autoencoder ensembles, capable of detecting point-wise outliers in time series. For time series forecasting, we studied how to employ multitask learning and dynamic filter generation to enhance forecasting accuracy. This work is documented in 5 papers in outlets such as IJCAI, ICDE, CIKM, and MDM.

Skyline and related queries

A new skyline query processing approach was devised for multi-dimensional data indexed by an R-tree. The MBRs of the R-tree are compared to each other in an I/O efficient manner that enables aggressive pruning of subtrees and thus data, thereby accelerating query processing. The work has been published at ICDE. Multi-criteria optimization has also been integrated with spatial constraints in a unified framework that enables multiple types of queries on spatial objects with respect to their so-called skyline dominators. Queries are processed by means of unified algorithmic templates of different complexity. This work was published at Information Systems.

Metric data

Any data that has an associated metric distance function is metric data. Here, the focus was on enabling efficient and scalable evaluation of similarity queries. Key ingredients include advanced indexing and algorithms that exploit the indexing for pruning. Contributions include a proposal for indexing that enables efficient similar-object retrieval and similarity joins, a proposal for aggregate kNN queries in metric data, a proposal for why-not querying of metric probabilistic range queries, a proposal for pivot-based indexing of metric data, a proposal for the indexing of metric uncertain data, proposals for distributed similarity querying and density-based clustering, and a proposal for similarity search on quasi-metric graphs. The results are published in some 10 papers in outlets such as ICDE, PVLDB, TKDE, and VLDBJ.

3.5.2 Advanced Analytics and Data Warehousing

Model-based time series management

Industrial systems, such wind turbines, generate so much sensor data that it is infeasible to store and analyze it all with current technology. We have therefore started work on using models, e.g., mathematical functions, to represent time series and created the open-source tool ModelarDB. The models can represent values within a user-defined error bound (possibly zero) and by only storing parameters for the models, instead of the raw data, huge compression factors can be achieved. In addition, query performance also benefits from the fact that many aggregations can be calculated directly from the models without recreating the raw data points first. This new work has resulted in three publications in TKDE, PVLDB, and SIGMOD (and one to appear in ICDE). Further, this work is the basis for our work in a newly started H2020 project and has led to the spinoff company ModelarData IVS.

Programmatic ETL

Our efforts within Programmatic ETL have continued. We continued to extend and maintain our open-source ETL tool *pygrametl*. The home page for the tool on average has 20 daily visitors from around the world and the tool is being used in different industrial domains. In collaboration with a company, we have also worked on the tool *SimpleETL* which builds on pygrametl. Finally, work has been done on automatic repair of ETL flows when changes are made to the sources and/or destinations. The work has resulted in two papers and a book chapter.

Time series correlations based on mutual information

A new research line within this topic investigated efficient and scalable mining of correlations in time series based on the information-theoretical concept of mutual information. Algorithms for both top-down and bottom-up mining were developed and refined with approximate algorithms that delivered high quality results with orders of magnitude performance gains over state-of-the-art techniques. This new line of work resulted in 3 publications at ICDE, EDBT and IEEE Transactions on Big Data.

Integrating analytics with decision-making and optimization

Within the wider topic of *advanced database technology for data analytics*, we worked on several subtopics. Within *prescriptive analytics*, we developed the SolveDB engine based on PostgreSQL which integrates optimization problem solving and integrated prediction in a unified SQL-based framework. We also made the first comprehensive literature survey on prescriptive analytics in the database community and the first encyclopedia article. In another subtopic, we developed the pgFMU engine based on PostgreSQL which integrates data analytics with physical system models based on Functional Mockup Units. In a third subtopic, we developed INVERDA, the first database framework and DBMS for seamless multi-version schema management. Finally, we worked on advanced support for probabilistic data in multidimensional OLAP. This work resulted in 8 publications in VLDBJ, TKDE, ICDE, EDBT, SSDBM, and EDBS.

Digital energy and energy flexibility

The group continued it work on digital energy and energy data management, with a focus on energy flexibility. The Flexoffer model was extended to so-called Dependency FlexOffers that effectively capture state information in an approximate way that can scale far beyond the previous state-of-the-art. A number of predictive analytics techniques for FlexOffer generation based on device-level forecasts were developed, e.g., for EVs, heat pumps, and white appliances. FlexOffer aggregation and dis-aggregation techniques that take (grid capacity) constraints into account were developed. Techniques for trading FlexOffers as so-called Flexible Orders in the existing NordPool market was developed, along with techniques for modeling, managing, and trading flexibility in a bottom-up way in cellular energy systems. Finally, techniques for adaptive user-oriented optimization of residential devices based on FlexOffers were developed. This work resulted in 10 publications, of which 7 were in ACM e-Energy, the flagship conference of the ACM Emerging Interest Group on Energy (ACM ENERGY) and one of which received the ACM e-Energy 2017 Best Paper award. The group also have a 3 hour tutorial at the key international conference IEEE SmartGridComm. The greatest impact has been in industry innovation projects where FlexOffers have been the key technology in 3 large Danish projects and 6 (very) large FP7/Horizon 2020 projects in the period. Finally, FlexOffers is the foundation for the spin-out company FlexShape.

3.5.3 Web, Text, and Social Media Data Management

This subsection integrates the results from the Web and Social Media Management as well as the Intelligent Web and Information Systems plan sections from the previous research evaluation report (DPW and MI chapters).

Querying semantic data

We have been working on querying semantic data on the Web from a wide variety of heterogeneous sources as well as stored locally on a single instance. To increase efficiency, we have developed advanced federated query processing techniques, such as Odyssey, that exploit source selection, parallelism, as well as heterogeneous interfaces. To increase availability of the data, we have proposed a decentralized architecture (PIQNIC) that uses replication along with appropriate indexing and query processing methods. Furthermore, we have worked on issues of result completeness and skyline queries in this context. Our work in this area resulted in eight publications and an ESWC 2019 best demo award nomination.

Semantic web data warehousing

We have developed techniques that combine semantics and Web data with analytical processing and data warehousing. In this context, our work covered optimizations for processing aggregate queries using materialized RDF views as well as provenance-aware queries on RDF data cubes. Furthermore, we created an extract-transform-load (ETL) framework converting Open Data into RDF and importing it into a Semantic Data Warehouse to enable Online Analytical Processing (OLAP). We have further developed techniques to support spatial data in a Semantic Data Warehouse and spatial OLAP (SOLAP). We have also worked on an RDF representation of a metadata framework to enable metamodel-driven (meta)data exploration. Enrichment of existing QB datasets with QB4OLAP semantics and RDF-based formalization of multidimensional queries have also been focus areas. Our work on Semantic Data Warehousing resulted in17 publications and a best demo award at WWW 2017.

Semantic data integration

Our work on semantic data integration has led to a framework (FrameBase) that taps on linguistic frames to provide a basis for integrating heterogeneous knowledge

and capture complex n-ary relationships. Our work on in this area resulted in more than 6 publications incl. AI Communications, ESWC, and IJCAI.

Knowledge graph exploration and information retrieval

One of the major challenges when working with knowledge graphs is to help users, especially non-expert users, get an overview of the data and retrieve relevant information. We have therefore analyzed the current state of the art in knowledge graph exploration, identified relevant problems and research challenges, and worked on exemplar queries and graph query suggestions in this context. Furthermore, we developed techniques for result diversification of path queries and skyline queries over knowledge graphs. We also worked on personalized page rank on knowledge graphs. Our work in this area has led to 7 publications incl. WWW, EDBT, and ISWC.

Textual data and natural language processing

We have also worked with textual data and news articles. In particular, we have developed an approach that retrieves textual evidence from a given corpus of documents for a given set of facts. Furthermore, we have developed an approach (Event2Vec) that creates neural embeddings for news events and surveyed the coverage of emerging entities in popular knowledge graphs. We have also worked on multilingual representations and published a dataset for subjectivity and review comprehension. We have also worked on entity extraction from medical text in collaboration with Bern University of Applied Sciences contributing to the goal from last evaluation named Understanding and Unlocking the Web in the area of Intelligent Web and Information Systems in Machine Intelligence group. This line of work has resulted in 5 publications at prestigious venues incl. SIGIR, EMNLP, and ESWC.

Social networks and media data

In this area, we have formulated the Top-k Influential Similar Local Query (TkISL) that finds the top-k influential local users from geo-tagged social media data who have interests similar to a query. This work applies machine learning to infer the location of tweets without geo-tags. Key contributions are (i) using words appearing in a user's tweet timeline, (ii) treating tweet timelines as sequenced data to model tweet locations, and (iii) considering collocation of social network

users. This line of work has led to one IEEE MDM publication and two IEEE TKDE publications. Within geo-social networks, we have also worked on multi-source data extraction and entity linkage, optimization of location influence and visitor prediction, and scalable query processing. This resulted in 10 publications in TKDE, KAIS, CIKM, EDBT, WSDM, ECML/PKDD, SSTD, and MDM, plus an SSTD Best Paper Runner-up award.

User centered machine learning/intelligent web and information systems multi view context aware recommender systems

In this area of work, we have studied multi view embedding and factorization models and their impact on recommendation accuracy and time to learn the models. We have explored context as an additional source of information which we embedded in linear and also nonlinear models. We have contributed with a hybrid model for learning which speeds up learning of joint factorization model especially on bigger data sets. The hybrid model is based on combination of multiplicative update rules with Barzilai-Borwein optimization with adaptive learning step. We have also contributed with several neural network models based on convolutional neural network as well adversarial network which in both cases lead to improvements in performance. We have also worked on recommendations over knowledge graph entities with explicit user ratings; with the MindReader movie recommender we have created a dataset that provides explicit ratings not only for items but also knowledge graph entities. This activity resulted in one Ph.D. thesis and one paper awarded with best paper award at FLAIRS 2018 conference, published at ACM RecSys 2019 as well as at ACM CIKM 2020. These contributed to the goal from last evaluation named Web long tail under the area of Intelligent Web and Information Systems in Machine Intelligence group.

Detecting sensitive information in documents

This area of work is a result of collaboration with Data-Intensive Systems group at Aarhus University. There are three main contributions from this area of work: A paraphrasing based recurrent neural network model for detecting sensitive information in documents, improved learning model for classifications based on identifying early not sufficiently recognized splitting regions in the training and using that information to focus training on those areas, and a dataset for testing new algorithms for the identification of sensitive information in documents. The work was published at good venues such as PAKDD conferences. The work also resulted in

one Ph.D. thesis defended at Aarhus University. The tasks in this area contributed to the goal from last evaluation named Understanding and Unlocking the Web in the area of Intelligent Web and Information Systems in Machine Intelligence group since the work directly works on extracting and classifying information in the documents (also web documents).

Explanations of recommendations

In this work we have contributed with a joint neural recommendation model which both recommends items and also explains why the items have been recommended by generating an explanation text based on review information attached to the data the model was learned from. We have explored both, word based and also character based models for learning to generate explanation texts with various success. The work was also published at ACM IUI and WebIST 2018 where it was short listed for the best student paper award. The tasks in this area contributed to the goal from last evaluation named Understanding and Unlocking the Web in the area of Intelligent Web and Information Systems in Machine Intelligence group since the work directly works on models for generating new texts learned on the web data (reviews).

3.5.4 Programming Technology

In the first part of the reporting period, research on Safety Critical Java (SCJ) was consolidated by two journal papers. The work on SCJ had major impact on the JSR 302 standard, which, to a large degree, is based on the initial proposal made by members of DPW and DEIS. The final version of JSR 302 was approved on 26.5.2020. SCJ has had major industrial impact and is now powering more than 20.000.000 devices according to Aicas GmbH (`https://www.aicas.com/`). An effort to persuade Danish industry to adopt SCJ proved fruitless, and a further study, done in collaboration with Epinion, revealed that very few Danish companies had any interest or need for verifiable hard real-time embedded systems development. Furthermore, over the years the small research community around real-time Java diminished, and in 2016 the last of the JTRES workshop series was held. It was therefore decided not to continue this line of research.

Before this reporting period, work had begun on parallelizing spreadsheet computations, a joint effort with members of the DEIS group and researchers at the IT University in the DFF funded Popular Parallel Programming (P3) project. This research builds in part on ideas from SCJ, translating spreadsheet programs into

timed automata analyzable by UPPAAL. Unfortunately, this approach was not scalable. However, dynamic scheduling, based on a dependency graph among cells, gave almost linear speed up on multi-core processors. Recently a cost calculator was completed, based on a formalization of the spreadsheet programming model using structural operational semantics. Work continues on a general framework for abstract interpretation of spreadsheet programs. So far, this has resulted in 3 publications, one of which received the 2019 Journal of Computer Languages Best Paper Award.

Work on Behavioural Types, originating in the DEIS group, but moved to DPW in the spring of 2017, has been pursued in the BEHAPI project, partially funded by the EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement No 778233. Behavioural Types are based on a combination of process algebra and type theory, elevating flat API descriptions to a graph structure of services. This permits automated analyses for correct API compositions so as to provide guarantees such as service compliance, deadlock freedom, dynamic adaptation in the presence of failure, load balancing etc. This line of work aims to bring the existing prototype tools based on these technologies to mainstream programming languages and development frameworks used in industry. This direction has resulted in 7 publications.

The work on research on teaching of programming and technologies and techniques for supporting teaching, has been enlarged to also cover research on Problem Based Learning and digitalization of teaching. 13 papers in this area have been published during the reporting period.

3.5.5 Unanticipated Research Directions

Bio related research

Towards the end of this period we have started collaborating with colleagues in the areas of Bio Science and Medicine. In general our goal is to build upon the the latest state of the art in computer science to solve problems in other domains. For example, in collaboration with AAU's Department of Chemistry and Bioscience, we supported SSI (Denmark's State Serum Institute) in analyzing and organizing data from corona test samples. Moreover, we used machine learning techniques to automatically assign diagnosis codes based on patients' medication history. The first result of our collaboration was published at the AIME conference 2020 and nominated as a best paper candidate.

Energy aware programming

In the research evaluation report from 2010-2015 we wrote: "Finally, we foresee a need for predictability of time, energy and cost of computations". A small subgroup of members of DPW began discussing the need for being able to predict the cost in terms of time and energy of big data analysis and during a departmental event on the UN SDGs a group of researchers from all groups of the department discussed the more general question of energy aware programming and a working group was formed. Funding for starting activities in this area was obtained through the InfinIT miniproject on energy certification of software and software installations. The purpose of this project is to bring researchers, companies and IT-organizations together in a preparatory work analyzing what is needed to establish an energy certification process of software and software installations. The group has also conducted a few experiments and a student programmer has worked on modifying a Python interpreter to enable analyses of both time and energy consumption of Python programs. The work is still in its early phase, but the group has produced a report mapping the state-of-the-art, which turned out to be somewhat more advanced than anticipated, but extremely scattered in terms of outlets for research in this area. We see a great potential for work in this area spanning all research groups in the department.

Planning related research

We have engaged in a collaboration with AAU's Department of Planning on sustainability assessment. Our ultimate goal is to enable an Open Source platform for managing and querying sustainability data in a reliable and unbiased way to compute "product footprints", e.g., computing how much CO2 is generated when producing a particular product. This collaboration has resulted in a publication at ISWC 2020 and a poster at ESWC 2019, which received the best poster award.

3.6 Own Evaluation

3.6.1 Research Areas

Temporal, Spatial, Spatio-Temporal, Spatio-Textual, and Metric Data

The group has conducted research in this general area for several decades, and the area continues to account for a large fraction of the research output. Thus, approximately half of all papers published in the period relate to this area.

At the start of the evaluation period, we highlighted real-time mobility and transportation analytics as a focus area and aimed to study improved real-time data ingestion and online querying capabilities. Thus, main memory processing was expected to play an increasing role.

In keeping with these expectations, a number of studies have targeted main-memory solutions. For example, we have proposed main-memory indexing techniques that aim to enable interactive applications; the UlTraMan and Elite systems for trajectory analytics exploit main-memory technologies; and a number of proposals for specific functionality also exploit main memory.

Next, we expected to be able to utilize additional data sources beyond GPS trajectories for analytics. Here, we have been able to use Bluetooth data from vehicles, and CAN-bus (fuel consumption) data. Towards the end of the period, we also started to use traffic photo data. However, we had hoped for additional progress on this front. The introduction of the General Data Protection Regulation (GDPR) added complications and an increased workload in relation to matters to do with data from users, including GPS data.

We also anticipated three specific outcomes of the research during the evaluation period. First, we anticipated "a state-of-the-art platform for network-based mobility and transportation analytics that supports real-time, on-line analytics." We pursued several directions in relation to this outcome, several of which came to fruition, most notably the Spark-based UlTraMan platform along with the VIP-TRA framework and the peer-to-peer based Elite computational infrastructure, both targeting trajectory analytics. In addition, aSTEP: aau's Spatio-TEmporal data analytics Platform, a student-based project, offered a range of functionalities on trajectory analytics and time series analytics. Further, we extended the existing software stack that handles billions of GPS data records to support real-time feedback for a subset of users when they were speeding and allowed for feedback from drivers, e.g., about the current speed limit and local weather conditions.

Second, we foresaw "a new paradigm for trajectory-based routing." Instead of first reducing trajectory data to travel times that can be associated with road segments, the overall idea was to exploit trajectory data more directly for routing. We have achieved very substantial progress on this and have as a result fleshed out several exciting new paradigms for routing. This research has been the focus of a dozen keynote lectures, in addition to a substantial number of papers.

Third, we foresaw "enablement of a range of new applications developed in collaboration with external collaborators." During the period we developed a series of app prototypes for both the Android and Apple mobile platforms. The apps

were evaluated in two user-studies with a large Danish insurance company. The Android app is still in usage.

In relation to keyword-based querying of spatio-textual data, we aimed to study how to evaluate query result quality and anticipated to enable the use of crowdsourcing for such evaluations. As reported in Section 5.1, we dedicated a Ph.D. study to this topic. We pursued two main directions: that of using GPS trajectories for extracting rankings that could then be used for quality assessment and that of using crowdsourcing to evaluate ranking functions. Overall, while we made good progress, no easy-to-use, reliable proposal exists for the evaluating of rankings in keyword-based querying of spatio-textual data.

Finally, in relation to indoor location-based services, aimed to provide efficient support for functionality needed by indoor services that utilize cloud computing and are delivered to smartphone users. We also aimed to study analytics on indoor mobility data. Several studies address these aims, and the "indoor" research covered a range of additional challenges.

Several noteworthy, unanticipated developments occurred during the evaluation period. Thus, unanticipated research related to metric data resulted in an impressive number of contributions to indexing and processing that were published in top venues. This line of research was driven to a large extent by Lu Chen, who has recently left. Therefore, we do not expect substantial activity related to metric data in the next period. We also note that much of the "indoor" research has been driven by Hua Lu, who has recently taken up a professorship elsewhere, for which reason we also expect decreased activity in this area. The evaluation period has also seen new activity related to time series analytics. As much data, including trajectory data, can be viewed as time series data, we expect work on time series to continue during the coming period. Finally, and perhaps most notably, we are seeing an increased integration of machine learning into our research in this area. We expect this to continue in the coming period.

In addition to accounting for many publications during the reporting period, research in the overall area also contributes very substantially to the listing by csrankings.org of the group as number one in Europe in the reporting period in terms of prestigious publications in the three database systems outlets covered by csrankings.org, which we believe is a strong indicator of academic excellence. This work has also attracted substantial citations, suggesting impact. Further, it has been recognized by, e.g., an impact award, best-paper awards, and about 20 grants, including prestigious excellence-pillar grants and applied grants, during the reporting period.

Advanced analytics and data warehousing

Our work on model-based time series management has achieved very good results and the proposed solution outperforms the current state-of-the-art tools significantly. We have received very positive feedback from industrial practitioners, and we will continue this line of work in the coming years in a recently started H2020 project and the spin-out company ModelarData IVS.

Our work on programmatic ETL has both continued to improve the pygrametl framework and created new solutions. Much this work has been done in collaboration with industrial users and based on their concrete challenges. We thus consider the work successful and will continue it in the coming years.

The work on correlation discovery based on mutual information was a non-planned topic that arrived with a newly hired postdoc. It nicely complemented the work on model-based time series management with a new type of analytics, and we actually aim to combine these two research lines in a new Horizon 2020 project. The results have been very good with a significant improvement over the state-of-the-art.

The work on integrating analytics with decision-making and optimization yielded very good results. The SolveDB and pgFMU engines combined high developer productivity with good efficiency for prescriptive analytics in both data-driven and cyberphysical settings. The planned scaling of SolveDB to a Big Data framework has been postponed due to hiring difficulties and will instead be done in an upcoming Marie-Sklodowska Curie EJD project. The unplanned work on supporting multiple schema versions in a DBMS (INVERDA) yielded a very elegant and effective solution.

Our work on digital energy, energy data management, and energy flexibility was highly successful in both academic and industrial impact terms. The FlexOffer concept has established itself as a key technology within the ever more important task of supporting up to 100% renewable energy in the energy system and is now even used in projects where AAU is not participating and running on a paid basis at customers. This shows that it has graduated successfully into real-world usage. Academically, we are very satisfied with 7 publications and a Best Paper Award at the flagship e-Energy conference, plus a key tutorial.

Semantic web data management

Having been an unanticipated research area in the previous period, this area of our research has become considerably more prominent, and resulted in a good number of publications and new projects.

As planned, we have continued to work on multidimensional (OLAP) support to RDF, Linked Open Data, and SPARQL, including aspects such as modeling and management of metadata, distributed and federated query processing, semantic extract-transform-load, and physical level optimizations. We have also developed techniques and tools for annotating and integrating data sources with additional multidimensional semantics. Several PhD students defended their PhD theses in this area. The demonstration of our work on "Geospatial OLAP on the Semantic Web" received the WWW 2017 Best Demo Award.

Our planned work on guidance and recommendation as well as explaining query results has begun a bit later than expected and started in 2019 within the context of a Marie Sklodowska-Curie Individual Fellowship, which has already led to publications in this area, tutorials on example-based querying at top conferences (SIGMOD 2020, SIGIR 2019, ESWC 2020), as well as publications on personalized page rank on knowledge graphs at EDBT 2020 and graph-query suggestions for knowledge graph exploration at WWW 2020. In 2018, a member of the group also co-authored a Morgan & Claypool book (in the Synthesis Lectures on Data Management series) on data exploration using example-based methods.

Furthermore, as planned, we have considered novel aspects in the area of semantic data management, such as provenance, quality, and context, which are reflected in a number of publications but continue to be part of our ongoing and future work. For instance, in our collaboration on sustainability data (project: A Virtual Laboratory for Sustainability Assessment, together with colleagues from the Department of Planning) we have encoded fine-grained provenance information to assess quality and trace back and explain the origin of the data (ESWC 2020 Best Poster Award). Our work has provided context for knowledge graph facts by retrieving textual witnesses where a human can validate quality and truthfulness of facts contained in a knowledge graph. However, because of changes in personnel the focus has been less strong on quality than originally anticipated.

Techniques for discovering and selecting new sources have been a major part in our work on federated query processing, which includes indexing as well as relevance assessment but not (yet) interestingness. Our work in this area resulted in a number of publications and a demo paper, which was one of three contributions nominated for the the ESWC 2020 Best Demo Award. Work in this area was mainly funded by the QWeb project funded by DFF project.

Social networks

The research on local search for social networks was essentially completed as planned. The research on location inference resulted in very good results and publications, although it was not planned. The anticipated research on food poisoning detection and social media analytics platform did not happen as expected, mainly because efforts to secure funding for this were not successful. The work on data extraction, entity linkage, and location influence in geo-social networks was only partly anticipated but led to very good results, including top publications and an award, so it is considered as successful.

User centered machine learning/intelligent web and information systems

The tasks in this area lead to one PhD defended at Aalborg University and one Ph.D. defended at Aarhus University. The results were published in 16 conference and journal papers. We received a best paper award at the FLAIRS 2018 conference, and we were shortlisted for a WebIST 2018 best student paper award. Among the published conference papers, five at the SA level. In general, we have contributed to the goals from the previous period mostly as planned. We were not able to satisfactorily contribute to the goal on computational methods for web user interfaces because we were unable to retain or hire appropriate staff for that area. Similarly, we were able to contribute to the goal on understanding and unlocking the Web only to some extent because the assistant professor focusing on this area left at the beginning of the evaluation period

Programming technology

Throughout the reporting period members of the programming technology sub-group have had an exceptionally high teaching and administration load. Two members have been full-time teaching/administration, one member has led the department's efforts in establishing a data science education, one member joined the group in the spring of 2017, one joined in the autumn of 2017 and left in the spring of 2018 and one member has been on leave for a total of one year. Thus, despite the sub-group counting 5 members, the mentioned circumstances, mean that most of the time the research efforts of the sub-group were equivalent to 1.5 full time research year pr. year (one funded by the department and the half funded by external funding).

During the reporting period members of the programming technology sub-group

has published 26 papers. We find this an acceptable level of output, especially because one paper received the 2019 Journal of Computer Languages Best Paper Award and one paper received the CONCUR Test-of-Time award.

During the reporting period we acquired funding for 7 externally funded projects; InfinIT Embedded Systems Engineering interest group, Popular Parallel Programming (P3) project, InfinIT project: Status of Danish embedded time-critical software, IT-Vest support for development of research and teaching in Computational Thinking in Computer Science, InfinIT miniproject: Energimærkning af Software og Software Installationer, IT-Vest support for development of research and teaching in Computational Thinking in Data Science, BeHAPI, EU H2020 RISE programme under the Marie Skłodowska-Curie grant agreement.

This constitutes roughly a 50% success rate, as two EU proposals, 2 DFF proposals, 1 proposal for the SparNord Fund and 1 proposal for Willum Experiment program of the Velux foundation were unsuccessful. This is an acceptable level, although we had hoped to acquire funding for one or two PhD students and one or two Postdocs.

Unanticipated research directions

- Bioscience and medicine
  Although not planned, collaboration with researchers from Bioscience and Medicine has emerged as a topic towards the end of the period, not only because of the corona epidemic. Our first results are quite promising and led to a best paper nomination for our work on automatically assign diagnosis codes based on patients' medication history.

- Energy aware programming
  This area has turned out to create a lot of interest both from industry and from the public, including several interviews with journalists and a funded InfinIT miniproject. Although only at its beginning we forsee at huge potential in this area.

- Planning related research
  In collaboration with AAU's Department of Planning we have contributed to developing an Open Source platform for managing and querying sustainability data. The pipeline of extracting information from Open Data sources, converting them into RDF, and enabling transparent use of the information has been successfully completed, led to a number of publications, and is currently used by the BONSAI organization (`https://bonsai.uno/`).

3.6.2 Scientific Output

As in the previous evaluation period, the publication strategy of the group focuses on promoting research quality and visibility of the results. Thus, the group strives to publish its best research results in the best general journals and conferences within the general scientific communities covered by the group (mainly data management and programming languages). Other results are published in the best specialized journals and conferences within specialized subareas, e.g., spatio-temporal data management, mobile services, data warehousing, semantic web, e-energy systems, and embedded or functional programming. The lowest priority is given to publication in general outlets with little quality control and consequent relatively low quality. This strategy aims to increase the impact and visibility as well as the peer-reputation of the group and its staff. The group's ranking of publication outlets is as follows. All those in bold are at the top level in the Danish national bibliometric ranking, BFI.

General Journals

Tier A. **ACM Transactions on Database Systems**, **ACM Transactions on Programming Languages and Systems**, **The VLDB Journal**, **IEEE Transactions on Knowledge and Data Engineering**, **Proceedings of the VLDB Endowment**, **ACM Computing Surveys**, **ACM Transactions on Information Systems**, **Information Systems**, **Journal of the American Society for Information Science and Technology**, World Wide Web.

Tier B. **Data and Knowledge Engineering**, **The Computer Journal**, Journal of Computer Languages, **Software: Practice and Experience**, **IEEE Software**, IEEE Computer, **Acta Informatica**, Eurasip Journal on Embedded Systems, Innovations in Systems and Software Engineering, ACM Transactions on Data Science, AI Communications, **Fundamenta Informaticae**, Future Generation Computer Systems, IEEE Systems Journal, IEEE Transactions on Big Data, Journal of Systems and Software, Knowledge and Information Systems, **Distributed and Parallel Databases**, Communications of the ACM.

Tier C. **Journal of Intelligent Information Systems**, **Journal of Database Management**, Journal on Data Semantics, **Information Sciences**, **SIGMOD Record**, ACM SIGSOFT Software Engineering Notes, Datenbank-Spektrum, Global Journal of Flexible Systems Management, IEEE Access, Applied Sciences, Data Science and Engineering, Electronic Proceedings in Theoretical Computer Science, Elementa: Science of the Anthropocene.

Specialized Journals

Tier A. **Geoinformatica**, **Transactions in GIS**, **Decision Support Systems**, **IEEE Transactions on Intelligent Transportation Systems**, **Semantic Web**, **Concurrency and Computation: Practice & Experience**.

Tier B. Intelligent Transportation Systems, Journal of Object Technology, **Higher-order and Symbolic Computation**, **ACM Transactions on Computing Education**, Int. Journal of Data Warehousing and Mining, Journal of Mobile Communications, Transactions on Large-Scale Data- and Knowledge-Centered Systems, **Energies**, Journal of Web Semantics.

Tier C. Trafik & Veje, Inzinerine Ekonomika, Intelligent Data Analysis, International Journal of E-Collaboration, **International Journal of Engineering Education**, Journal of Problem Based Learning in Higher Education, SIGWEB Newsletter.

General Conferences

Tier A. **ACM SIGMOD**, **OOPSLA**, **POPL**, **PLDI**, **ACM SIGKDD**, **ACM SIGIR**, **WWW**, **IEEE ICDE**, **EDBT**, **ECOOP**, HOPL, **CIKM**, **ICDM**.

Tier B. **DASFAA**, **ADBIS**, **DEXA**, **IDEAS**, IEEE International Conference on Big Data.

Tier C. SAC, BIS, ICEIS, OTM, COMAD, BNCOD, BTW, ETFA, KES, ADC, FedCSIS, ICCIT, CRIWG, OZCHI.

Specialized Conferences

Tier A. **SSTD**, **ACM SIGSPATIAL**, **MDM**, **SSDBM**, TIME, **SOCC**, ICDCS, **ER**, **DaWaK**, **SDM**, **ECML/PKDD**, **PAKDD**, ISMIR, MMM, ICME, SAS, **ICFP**, **OOPSLA/SPLASH**, **ACM e-Energy**, **ESWC**, **ISWC**, **ECIR**, **UbiComp**, WSDM, RecSys, **ECML/PKDD**, **IJCAI**, **PAKDD**, **EMNLP**.

Tier B. MobiQuitous, MobiDE, W2GIS, STDBM, MobiLight, PerTrans, ICMB, ITS, DOLAP, BIRTE, ELS, JTRES, JIST, ISORC, RTAS, Scheme and FP, **LCTES**, BYTECODE, WCET, GIR, ISA, DARE, Cloud-I, DeRiVE, COLD, BIaW, EnDM, DaMoN, ICEL, DLS, NWPT, ASE, Annual Transport Conference at Aalborg University, TAAI, Emerging Topics in Semantic Technologies, DSAA, SmartGridComm, WEBIST, WebDB, DeSemWeb, LDOW, ESWC (Satellite Events), WWW Companion Volume, AIME, ISWC P&D/Industry, DMAH,

Table 3.1: Publications according to ranking category

	General			Specialized			Total
	A	B	C	A	B	C	
Journals	65 (61.9%)	14 (13.33%)	7 (6.67%)	9 (8.57%)	3 (2.85%)	7 (6.67%)	105
Conferences	71 (30.47%)	12 (5.15%)	5 (2.15%)	75 (32.19%)	29 (12.45%)	41 (17.60%)	233

Tier C GeoRich, LDK, eBISS, CBMS, EDOC, CCGrid, IUI, LBS, FLAIRS, APWeb-WAIM, QuWeDa, **CHIIR**, ICATPN, eCOM, LREC, BigSpatial, TPDL, **ECAI**, WebSci, APLAS, EXPRESS/SOS, FASOCC, WS-FM, PLACES, PLATEAU, SETE, **CC**, **ISSTA**, PPDP, European Conference on e-Learning, SIGTYP, BlackboxNLP.

The top-tier outlets are the most selective. In particular, publication in tier A general conferences is highly competitive. Papers are typically 10 to 12 pages long (2-column, 9 or 10 point type), and most accepted papers report on experimental studies that require substantial software development. The acceptance rates for the first tiers of general conferences and for some of the specialized conferences are below those of the best journals. In particular, the acceptance rates at tier A general conferences are typically below 20%. There is a very good match between the group's top-ranked outlets and the top level in BFI ranking. All but one of the tier A (general and specialized) journals are at the top level, and all but two of the tier A and B general conferences and most of the tier A specialized conferences are at the top level. In the evaluation period, 338 refereed publications were co-authored by DPW members. A significant portion of them, 105, are in journals, of which 65 are in tier A general journals. Of the 233 conference and workshop papers, 71 are in tier A general conferences and 75 in tier A specialized conferences. Table 3.1 shows that 70.47% of all journal publications were published in tier A general or specialized journals. Similarly, 62.66% of all conference publications appeared in tier A general or specialized conferences.

The group strives to publish in the very top venues. Table 3.2 gives the breakdown of journal and conference publications according to the CORE ranking categories. A substantial amount of papers are published in the A* venues. Note that some journals and conferences have yet to be included in the CORE ranking system and that the percentages in the table are relative to the total numbers of journal (105) and conference (233) publications.

Table 3.2: Publications according to the CORE categories

	A*	A	B	C
Journals	41 (39.05%)	5 (4.76%)	11 (10.48%)	2 (1.90%)
Conferences	60 (25.75%)	44 (18.88%)	20 (8.58%)	29 (12.45%)

Table 3.3: Publications according to BFI levels

Level 2 journals (the top level)	79 (75.24%)
Level 1 journals	20 (19.05%)
Level 2 conferences (the top level)	157 (67.38%)
Level 1 conferences	55 (23.61%)

Further, Table 3.3 gives the breakdown of journal and conference publications according to BFI categories. Note that some journals and conferences have yet to be included in the ranking system and that the percentages in the table are relative to the total numbers of journal and conference publications. We are very satisfied to see that by far most of our publications are in BFI Level 2 (top level) outlets.

The group's contribution to the BFI-points of the Department of Computer Science has increased steadily during the reporting period. In 2015, DPW contributed 35.01% (37.7 out of 107.75 BFI-points); by 2019, this percentage had increased to 54.68%. Table 3.4 shows the BFI-points of the group and the Department of Computer Science (CS) from 2016 until 2019. We do not include the BFI-points for 2020 because these values are not available before late 2021.

Table 3.4: DPW BFI-points within the department

	2016	2017	2018	2019	Total
DPW	63.9	67.2	99	85.4	315.6
CS	140.97	134.51	198.64	156.24	630.36
% CS	45.36%	49.98%	49.84%	54.68%	50.07%

3.6.3 Projects and Funding

Projects

We describe briefly selected main projects active during the period.

TotalFlex, 2012—2016, DKK 4,034,000. The project significantly developed and demonstrated the FlexOffer concept and techniques for utilizing most of the energy flexibility in an energy system, in collaboration with energy technology, economics, and a range of domain companies.

Arrowhead, 2013—2017, DKK 1,079,042. The project integrated the the FlexOffer concept and techniques into the Arrowhead framework for collaborative automation, resulting in a virtual market of energy, and demonstrated it in a number of industrial trials.

Erasmus Mundus Joint Doctorate in Information Technologies for Business Intelligence - Doctoral College, 2013—2021, DKK 7,080,480. This is a doctoral programme addressing fundamental BI challenges. Upon completion, graduates are awarded with a joint degree from the two universities. Expected to produce a total of 14 Ph.D. degrees in DPW.

Obel Professorship in Data-intensive Systems, 2013–2021, total funding DKK 7,000,000 (external), 7,000,000 (AAU). The key purpose of the grant has been to ensure a strong, sustainable, and internationally oriented research environment in data-intensive systems that conducts world-class research and creates societal value regionally as well as nationally.

Villum Kann Rasmussen Annual Award, 2013–2021, DKK 1,974,226. This grant was used to provide back-up funding, enabling considerable agility in the hiring of temporary staff, notably postdocs, and Ph.D. students. This has had a very positive impact on hiring quality and volume.

Center for Data-Intensive Cyber-Physical Systems (DiCyPS), 2015–2020, Total: DKK 32,500,000. DPW share: DKK 9,000,000. The DPW part of the project worked on management of sensor data from large cyber-physical systems wuch as the transport and energy systems, including management of trajectory and time series data, and prescriptive analytics.

Popular Parallel Programming, 2015–2019, Total funding: DKK 5,680,479. DPW share: DKK 1,420,119.75. The goal of the P3 project is to help spreadsheet users harness the power of their multicore computers.

QWeb: Querying the Web of Data Easily and Efficiently, 2015–2019, DKK 6,447,744. QWeb's main goal is to develop breakthrough technologies based on Semantic Web standards, such as RDF and SPARQL, and enable answering user queries easily and efficiently over the Web of Data.

GOFLEX, 2016–2020, DKK 9,058,734. The project further developed and the FlexOffer concept and techniques including better aggregation, flexibility trading, and AI-based flexibility extraction and grod bottleneck prediction. The resulting GOFLEX system was demonstrated in 3 large trials in DE, CH, and CY with more than 500 prosumers covering the residential, govermental, commercial, transport, and industrial sectors.

NRP75 Relief Support, 2016–2022, DKK 1,085,000. This grant offers relief support in relation to Christian S. Jensen's service as president of the steering committee of the Swiss National Research Program 75 on Big Data.

OptiTruck, (INEA), GA 713788, H2020, EU, 2016–2019, 1,629,352. The purpose of the project is to reduce the CO2 emissions from +40 tons trucks (no platoning). AAU does spatio-temporal data integration.

Unleashing the Potential of Open Data, 2017–2021, DKK 3,000,000 (AAU). The goal of this project is to help unleashing the potential of Open Data, which covers aspects of information extraction, indexing, query processing, and search.

Analytics of Time series in spatial networks, 2018–2022, DKK 2,556,000. The project aims at providing accurate and scalable learning algorithms for correlated time series analytics, such as forcasting and outlier detection.

Professorship in Data-intensive Systems, 2018–2023, DKK 8,500,000 (external), 3,740,000 (AAU). The grant is to strengthen and consolidate the research group in Data-intensive Systems. Specifically, the grant covers Katja Hose's professorship during a five-year period.

RESPONS, EU, Interreg, Baltic sea Region, R101, 2018–2021, 1,940,625.

The purpose of this project is to increase the usage of demand-response transportation (DRT) in rural areas in the Baltic/Scandinavian area.

SEMIOTIC, 2018–2021, DKK 2,045,858. The project aims to integrate concepts and techniques from wireless communication and big data analytics to enable integrated and optimized communication and analytics on massive amounts of IoT sensor data.

A Data-Intensive Paradigm for Dynamic, Uncertain Networks, 2019–2023, DKK 5,882,400. The project aims at fleshing out a data-intensive paradigm to enable dynamic, uncertain networks, with a focus on applications in transportation.

A virtual laboratory for sustainability assessment, 2019–2020, DKK 894,000 (AAU). The ultimate goal of this project is to enable an Open Source platform for managing and querying sustainability data in a reliable and unbiased way to compute "product footprints", e.g., computing how much CO2 is generated when producing a particular product.

Data Management Foundations for Indoor LBS, 2019–2023, total funding: DKK 5,655,534, DPW share: DKK 2,764,551. The project is intended to create a prototype indoor data management system with general, systematic, and solid data management foundations upon which indoor location-based services (LBS) can be built to provide location-dependent information efficiently and effectively to users in a broad range of indoor scenarios.

EDAO: Example Driven Analytics for Open Knowledge Graphs, 2019–2021, DKK 1,544,474. The proposed research project has the goal to provide methods to support the understanding of the information represented by LOD repositories and to provide easy access to their content. To address this need, the researcher aims to develop an Exploratory Analytics System for LOD.

Flexible Energy Denmark, 2019–2023, total funding: DKK 6,400,000, DPW share: DKK 3,128,238. This project aims to create the next generation of intelligent, integrated energy systems based on Big Data and AI to pave the way for the green transition

RelWeb: A Reliable Web of Data, 2019–2023, DKK 5,904,000. To enable reliable query processing over the Web of Data, this project aims at developing techniques that use replication, blockchain, and peer-to-peer (P2P) principles to

relieve data providers from the burden of having to ensure uninterrupted availability of their Web interfaces.

Corona Danica, 2020, total funding: DKK 4,852,313. DPW share: DKK 360,938. In response to the corona pandemic, this project aims at generating and analyzing the coronavirus' genetic material.

Digital Research Centre Denmark (DIREC), 2020–2025, total funding: DKK 100,000,000. DPW share: DKK 4,000,000 (estimated) DIREC aims to improve Danish computer science research and education capacity, to address societal challenges, and to create societal and industrial value.

DNA Genome Analysis of the Coronavirus, 2020, total funding: DKK 5,000,000. DPW share: DKK 315,000. In response to the corona pandemic, this project aims at generating and analyzing the coronavirus' genetic material.

domOS, 2020–2023, total funding: DKK 4,304,000, DPW share: DKK 2,220,460. The project will develop and demonstrate an operating system for smart services in buildings where different IoT platforms and applications can be integrated seamlessly.

DREAMS Digitally supported Environmental Assessment for Sustainable Development Goals, 2020–2023, total funding: DKK 18,239,314, DPW share: DKK 1,035,000. The vision of the DREAMS project is to promote progress on Sustainable Development Goals (SDGs) by digitally transforming the way society accesses and communicates information about environmental impacts of projects and plans in order to enable the best decisions towards green transition in a transparent and inclusive democratic process.

FEVER, 2020–2023, Total: DKK 4,188,000. DPW share: DKK 2,512,800. The project builds on the GOFLEX project and will further develop the FlexOffer concept and techniques in the GOFLEX platform to include more powerful flexibility aggregation and management, both central and P2P (blockchain-based) flexibility trading, and increased use of AI. The resulting system will be demonstrated in 4 large trials in DE, ES, HE, and CY covering the residential, govermental, commercial, transport, and industrial sectors.

Light-AI for Cognitive Power Electronics, 2020—2023, total funding: DKK 2,996,278, DPW share: DKK 1,489,660. This project aims to design and de-

velop light-AI, including computation-light AI and data-light AI, to enable cognitive power electronics.

MALOT: Managing Mobility Data Quality for Location of Things, 2020–2022, DKK 1,636,068. Location of Things is an IoT paradigm for mobility analytics, in which massive mobility data is being gathered, processed, and transmitted among heterogeneous data nodes in a decentralized architecture. This project aims to develop new data quality management techniques adaptive to the Location of Things.

Management of Real-time Energy Data – MORE, 2020–2023, DKK 4,761,803. The project will extend the ModelarDB system to provide efficient model-based time series management on both edge and cloud and provide rich and fast analytics and machine learning functionality.

Time Series Analytics and Spatio-Temporal Data Management, 2020–2022, DKK 1,800,000. This project aims at designing and developing algorithms for time series analytics and spatio-temporal data analytics.

Funding

The group has attracted substantial funding for its research.

Tables 3.5 and 3.6 list funded projects grouped by starting year, reporting for each group the count of projects, the total amount of funding, and the funding categories and sources. The reported figures are the amounts from external grants that go specifically to DPW. As can be seen, the group has received funding from 46 projects during the period, with a total funding of DKK 106.8 million (EUR 14.36 million), an increase of 71.2% over the previous period. Of these projects, 36, totaling DKK 68.73 million (EUR 9.24 million), started during the period. As can be seen, the amounts of projects and funding vary considerably from year to year, but were generally stable over the period.

Robust capacity for maintaining a funding portfolio. As it is important to always have a healthy funding portfolio, the group strives to maintain a robust capacity for maintaining such a portfolio. To achieve robustness, it is important that multiple staff members are able to attract funding as (co-)PI and that staff members are able to attract funding from diverse sources. Figure 3.5 shows the number of projects attracted as (co-)PI for DPW staff. As can be seen, a dozen staff members have attracted funding as (co-)PI, and 9 staff members have at-

Table 3.5: External funding attracted by DPW (–2018)

Start Year	New Projects	Amount in DKK	Funding Categories and Sources
2012–2015	10	38,075,611.75	DK_innovation (Mapicture ApS), DK_public_excellence_pillar (Independent Research Fund Denmark), DK_strategic (Energinet.dk ForskEL, FORSKNINGS- OG INNOVATIONSSTYRELSEN, Innovation Fund Denmark), EU (EC Erasmus Mundus), EU_strategic (European Commission: ARTEMIS Industry Association), Private_foundation (Obel Family Foundation, VILLUM FONDEN)
2016	6	12,224,511	DK_innovation (Region Nordjylland), DK_public_excellence_pillar (The Danish Agency for Science, Technology and Innovation), EU_strategic (EC Horizon 2020), Intl_strategic (VRI Agder, Norway, Swiss National Science Foundation)
2017	5	1,046,375	DK_innovation (Mapicture ApS, Rambøll Danmark A/S), DK_strategic (EUopSTART, InfinIT), EU_strategic (EC Horizon 2020)
2018	7	15,627,108	DK_innovation (Region Nordjylland), DK_public_excellence_pillar (Independent Research Fund Denmark), DK_strategic (IT-Vest), EU_excellence_pillar (EC Horizon 2020 Marie Sklodowska-Curie), EU_strategic (Interreg Baltic sea Region), Private_foundation (Poul Due Jensen Foundation)

Table 3.6: External funding attracted by DPW (2019–2020)

Start Year	New Projects	Amount in DKK	Funding Categories and Sources
2019	7	19,495,498	DK_public_excellence_pillar (Independent Research Fund Denmark, The Danish Agency for Science and Higher Education), DK_strategic (InfinIT, Innovation Fund Denmark), EU_excellence_pillar (EC Horizon 2020 Marie Sklodowska-Curie)
2020	11	20,331,729	DK_strategic (FORSKNINGS- OG INNOVATIONSSTYRELSEN, InfinIT, Innovation Fund Denmark), EU_excellence_pillar (EC Horizon 2020 Marie Sklodowska-Curie), EU_strategic (EC Horizon 2020), Private_foundation (Poul Due Jensen Foundation, VILLUM FONDEN), Intl_strategic (HUAWEI TECHNOLOGIES SWEDEN AB)
Total	46	106,800,833	Sources contributing in excess of DKK 2 million: EC Erasmus Mundus, EC Horizon 2020, EC Horizon 2020 Marie Sklodowska-Curie, Energinet.dk ForskEL, Independent Research Fund Denmark, Innovation Fund Denmark, Obel Family Foundation, Poul Due Jensen Foundation, VILLUM FONDEN
Total new	36	68,725,221	

tracted more than one project as (co-)PI. Thus, many staff members, junior as well as senior, are able to attract their own funding. For example, two young (non-tenured) faculty are PIs on two EC Horizon 2020 Marie Sklodowska-Curie fellowships. Thus, the future success of attracting funding is spread across many staff members, which is very encouraging.

Figure 3.5: Number of projects attracted as (co-)PI across DPW staff during the reporting period

Funding sources. The funding sources are diverse and exhibit balances between Danish versus European sources, public versus private sources, as well as sources that fund independent versus application-oriented research. Figure 3.6 shows the total amount of funding attracted per category.

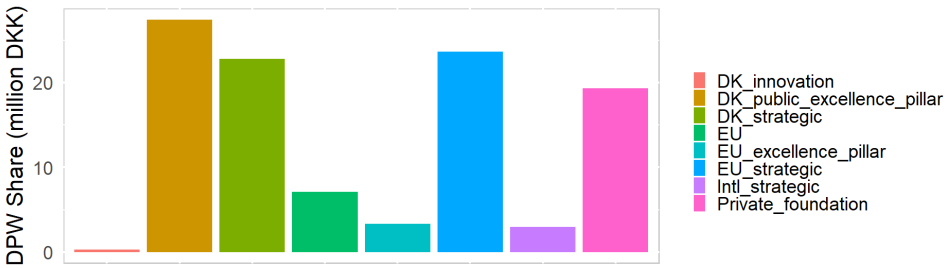

Figure 3.6: Amount of funding attracted per funding category

Danish public excellence pillar councils, including the Independent Research Fund Denmark, account for 25.65% of the funding. Danish strategic councils, including the Innovation Fund Denmark, account for 21.31% of the funding. European excellence pillar sources, including EC Horizon 2020 Marie Sklodowska-Curie, account for 3.14%. European strategic sources, including EU Horizon 2020, account for 22.15%. Other European sources, such as EC Erasmus Mundus, account for another 6.63%. Private Danish organizations account for 18.09%.

There is a reasonable balance between independent and application-oriented research, with the latter category being the largest. Six projects exceed DKK 6 mil-

lion in funding: Querying the Web of Data Easily and Efficiently (QWeb), Center for Data-intensive Cyper-Physical Systems (DiCyPS), the Obel professorat i data-intensive systemer, the Poul Due Jensen Professorship in Data-intensive Systems, GOFLEX, and the Erasmus Mundus Joint Doctorate in Information Technologies for Business Intelligence - Doctoral College (IT4BI-DC). In summary, the group is very satisfied with the level of funding and the fact that it has acquired more large projects from attractive sources than previously. Figure 3.7 shows that the new funding received by the group is overall exceeds the funding spent.

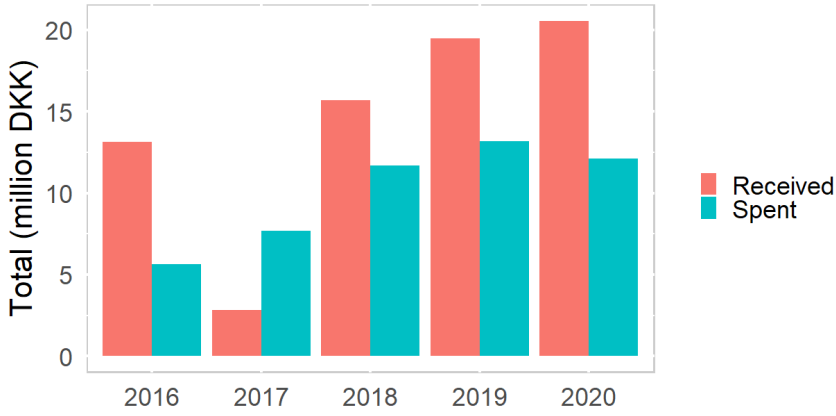

Figure 3.7: Amount of funding received and spent per year

In addition to the externally funded projects, the group has been successful at attracting internal funding. Examples include Unleashing the Potential of Open Data funded by the AAU Talent Programme (DKK 3 million), In-Memory Big Trajectory Data Analysis funded by AAU (DKK 1 million), co-funding of the Obel professorat i data-intensive systemer (additional DKK 7 million), and co-funding of the Poul Due Jensen Professorship in Data-intensive Systems (additional DKK 3.7 million).

3.6.4 Scientific Service and Recognition

Members of the group have served the scientific community in a number of leadership roles that indicate substantial peer recognition, including as Editor-in-Chief of ACM Transactions on Database Systems, one of the top data management journals, and as Associate and Area Editor of several journals and an encyclopedia. They have served as program committee cochair and vice chair and as general

chair for a range of conferences, including DASFAA, VLDB, WISE, NDBC, MDM, ICDE, APWEB-WAIM, and HT, and have co-organized several workshops. This is in addition to other leadership service and a very substantial number of program committee memberships of the top conferences. In the reporting period, members of the group received several best paper, demo, and poster awards, including the ESWC 2019 Best Poster Award, ACM e-Energy 2017 Best Paper Award and WWW 2017 Best Demo Award.

3.6.5 Education of Researchers

Educating new researchers has also been an important activity throughout the period. In total, 25 Ph.D. students started their Ph.D. during the reporting period, 19 Ph.D. students defended their thesis during the reporting period, and there are 17 active Ph.D. students as of end of December 2020. Ph.D. students who have left during the reporting period now hold positions at Pandora, Salling Group, Gothenburg University, Cargill, Grab (Singapore), A.P. Moller – Maersk, Geodatastyrelsen, Amazon Web Services (Germany), Department of Planning (Aalborg University). 18 postdocs have worked in the group in the period. Ten are still at the group, either as postdocs, assistant professors, or associate professors; four hold academic positions at Huazhong University of Science and Technology, Hacettepe University, INRIA/IRISA, Gothenburg University; and other three work at the companies MHI Vestas Offshore Wind, FlexShape, Veovo. In summary, the group finds that it has been successful in educating new researchers and that this area will have increasing importance in the future.

Figure 3.8 shows that ratio between the number of Ph.D. students and the permanent staff has remained above 1.5 during most years in the reporting period.

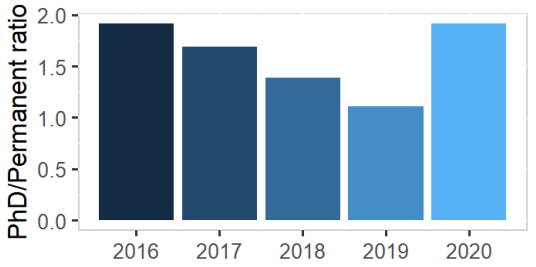

Figure 3.8: Ph.D. students per permanent staff ratio

Table 3.7 shows the number of Ph.D. defenses during the reporting period. This

Table 3.7: DPW Ph.D. defenses within the department

	2016	2017	2018	2019	2020	Total
DPW	4	4	3	5	3	19
CS	9	4	8	7	5	33
% CS	44.44%	100%	37.5%	71.43%	60%	57.58%

number, 19, represents an increase of 46% in comparison to the 13 Ph.D. defenses in the previous period. Moreover, the group's contribution to the number of Ph.D. defenses of the Department of Computer Science has been significant during the reporting period. Table 3.7 shows the number of Ph.D. defenses in the group and in the Department of Computer Science (CS) from 2016 to 2020.

3.6.6 Industrial Uptake and Dissemination

The work on model-based time series management has received much interest from industry practitioners. In the newly started H2020 project MORE, ModelarDB will be used on large amounts of data from renewable energy sources (RES), in particular wind turbines and photo-voltaic modules. Further, the spinoff company ModelarData IVS will help companies to start using ModelarDB.

The work on programmatic ETL is used actively in industry. For example, the company FlexDanmark that handles demand-responsive transport in Denmark uses SimpleETL and pygrametl for loading a data warehouse with around 2 million coordinates collected daily from 5,000 vehicles. We have also been in touch with other industrial users who have told us about their use of pygrametl, but since the source code is freely available both from our GitHub page and non-official repositories, we do not know the exact number of users. On average, the home page for pygrametl has 20 unique visitors every day.

The work on digital energy, in particular FlexOffers, has had significant uptake in industry with use in more than 10 large innovation projects, including the key Danish project in this area, paid installations, and being the foundation of the spin-out company FlexShape.

The work on Safety Critical Java had major impact on the JSR 302 standard, which, to a large degree, is based on the initial proposal made by members of DPW and DEIS. The final version of JSR 302 has been approved in a ballot ending 26.5.2020. Safety Critical Java has had major industrial impact and is now powering more than

20.000.000 devices according to Aicas GmbH (`https://www.aicas.com/`).

3.6.7 Collaboration

The group maintains collaboration with a large number of international colleagues. The coauthorships of papers published during the reporting period offer an accurate description of the collaborations. In particular, the group has substantial academic collaborations with colleagues at several Asian universities (e.g., Peking University, Tsinghua University, East China Normal University, Zhejiang University, University of Science and Technology China, Shandong University, Hong Kong Baptist University, Nanjing University of Aeronautics and Astronautics, Hong Kong Polytechnic University) and Australian Universities (Monash University, RMIT), in addition to US universities (e.g., Auburn University, Tufts University) and European Universities (e.g., University of Hagen, TU Dresden, Free University of Bozen-Bolzano, UPC Barcelona, Nantes University, Universite Libre de Bruxelles). There has also been collaboration with researchers at Danish Universities (ITU, DTU and RUC). Internal collaboration with other groups in the department has continued and even increased over the period, on topics like Energy certification of Software and Software Installations, embeded time-critical software, human computer interaction for flexible energy systems and data-intensive cyber-physical systems. Moreover, collaborations with groups from other departments have been started, for example in analyzing the spread of the corona disease. Such broad collaborations are expected to increase in the coming years.

3.6.8 Hiring

The group has hired 25 Ph.D. students in the reporting period. The group is happy with the Ph.D. students that has hired, but it has been a substantial challenge to hire high quality Ph.D. students. Hiring external Ph.D. candidates is challenging, although we offer very good salaries and conditions. Establishing elaborate ranking and interviewing procedures has helped somewhat, but the challenge remains. One reason is that Ph.D. recruitment often occurs internationally after the B.Sc. degree rather than the M.Sc. degree. A 4+4 Ph.D. program was established in 2017 and currently we have two Ph.D. students in the program. We expect to have more students in that program in the coming years.

At the postdoc level and above, the hiring opportunities have been significantly better, especially in data-intensive systems where four highly qualified associate professors (one external and three internal) and 13 highly qualified assistant pro-

fessors were hired (seven external and six internal). Three of these were promoted to associate professor. In programming technology, one highly qualified assistant professor was hired. In addition, 15 very well qualified postdocs, 11 external and four internal, were hired in data intensive systems; and one very well qualified postdoc was hired in programming technology. Three of these postdocs were promoted to assistant professors, and one of these postdocs was promoted to associate professor during the reporting period. One assistant professor in programming technology was promoted to associate professor. The external hires have generally been Ph.D. students or postdocs from top European or Asian-Pacific universities. During the reporting period, we managed to hire five assistant professors in tenure track positions and three associate professors in ordinary positions. Additionally, two associate professors were promoted to professor, and three associate professors were promoted to professor MSO.

The group's hiring has a good general balance in diversity. For example, the gender of the hired assistant professors, associate professors, and professors is presented in Figure 3.9.

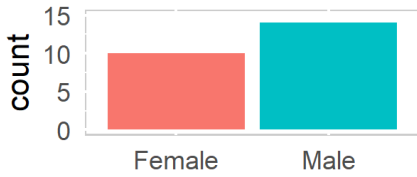

Figure 3.9: Gender—new assistant, associate, and full professors

During the reporting period, the group hired, as assistant, associate, and full professors, graduates from multiple countries. This information is presented in Figure 3.10.

From the 24 hirings in the reporting period for assistant, associate, and full professor positions, 12 are graduates from universities in the top-100 of the Times Higher Education World University Rankings 2021. This includes graduates from very well ranked universities such as Imperial College London (ranked 11th), National University of Singapore (ranked 25th), and University of Hong Kong (ranked 39th).

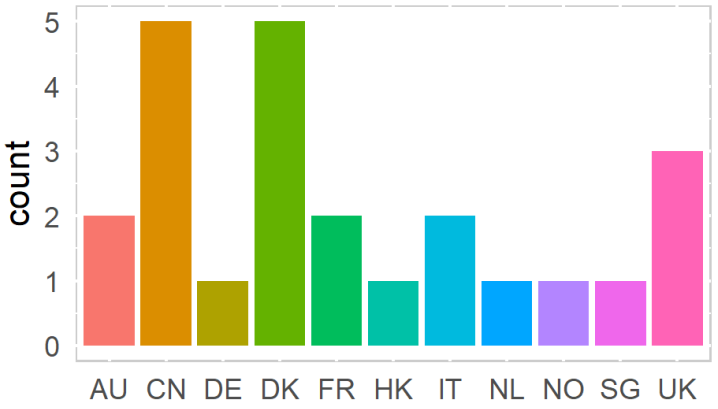

Figure 3.10: Country issuing the latest diploma of the new Assistant Professor, Associate Professors, and Professors

3.6.9 Comments from the Previous Panel

We recall the recommendations of the previous research panel and explain how the recommendations were addressed.

The number of PhD defenses given the size of the unit has been rather low. More external funding is required to appoint more PhD students. In the last phase of the evaluation period, quite a few new PhD students have been appointed. So, this weakness is being addressed properly.

During the reporting period 2016—2020, 25 Ph.D. students started their Ph.D., 19 Ph.D. students defended their theses, and there are 17 active Ph.D. students as of end of December 2020. The number of Ph.D. defenses, 19, represents an increase of 46% in comparison to the 13 Ph.D. defenses in the previous period. More details on Ph.D. education are provided in Section 3.6.5.

The report and the presentation mention "perceived" real world applications. Given the maturity of the computer science field, it is important to look at real applications to be aware of the full extent of the challenges of the applications. There is a number of opportunities. Big Data is getting more and more attention, also outside the academic world. In many companies and organizations, there is a need to handle large volumes of data and gain knowledge from it to improve products and services. The number of possible applications is enormous. It is important to

focus on specific applications while still producing generic solutions or results.

During this reporting period, the group has worked on a variety of (actual) real-world applications, but has also chosen to put focus on particular application areas, most notably digital, or green, energy and mobility analytics. As a result of its application-oriented research, the group has had significant societal impact and has contributed to a number of SDGs, including SDGs 3, 7, 11, and 13. When conducting foundational research, where applications may occur some 5 or 10 years into the future, the group's research often continues to be inspired by envisioned applications that may be possible in the future.

Furthermore, society needs people that know how to handle data, so it is important to educate students in Data Science. The recommendation is to start a Data Science program as soon as possible.

During this period, we have designed and launched a new education in Data Science. The first intake occurred in 2019. The education has a strong focus on analytics, statistics, machine learning, and data management, but involves also, e.g., law and ethics and application domains.

Given the change in funding, that goes more along the lines of applications, there is the threat that Computer Science becomes invisible. Because of this, it is important to claim a strong position of CS in these projects to cover also the CS challenges. For DPT, this means that it should have a clear research agenda for the database technology challenges it wants to address.

In its key application areas, predominantly transport and energy, DPW has been able to combine making significant contributions to real-world applications (see the response to the previous panel comment) while at the same time significantly advancing the scientific state-of-the-art for data management technology. For example, this has been done for traffic analytics, vehicle routing, prescriptive analytics, and model-based time series management. The former two relate to transportation and have provided industrial and academic visibility to DPW. The latter two are rooted in our work in digital energy, making DPW very visible in the energy domain, nationally and internationally.

Given the changes in the requirements coming from society and the advances in CS, it is important to have a flexible organization to adequately support cooperation across units. It is suggested to have an open mind about the positioning of Programming Technology in the departmental structure, e.g., to consider whether a matrix structure can help in this case as well as others.

In September 2020, the sub-group on Programming Technology ceased to be part of the group. To increase cross fertilization and collaboration, the majority of its members joined the DEIS group.

3.7 Plans for 2021–2025

Plans are described for data-intensive systems and semantic data management. Since the sub-group on programming technology is no longer part of the group, plans for programming technology are not included here.

Establishment of a new group

With the growing group size and increased activities within certain areas, the Department of Computer Science considers to establish a new group (incubator) to represent these areas within research and teaching. The ambition is to establish a new group on an equal footing with the three established groups that will reach a sustainable size within the next 3–4 years.

The new group focuses on Web-scale data management, knowledge engineering, and Web intelligence. This includes distributed, decentralized, and heterogeneous data processing, graph data management, knowledge graphs, applied machine learning, data semantics, natural language processing, Semantic Web, recommender systems, personalization and user modeling, and stream processing.

Excellence and impact

The group views fundamental and applied research as equally valuable and and targets excellence in both kinds of research. Further, the group believes that the excellence and impact of both can be improved substantially through effective cross-fertilization and collaboration.

In terms of foundational, or long term, research, performance indicators include publication in top scientific outlets (excellence) and citations (scientific impact), as well as awards, invited talks, and other forms of peer recognition.

In terms of application-oriented research, a more diverse and less established range of indicators exist, such as demonstrations, industry funding, industrial uses of results and software, spinouts, software downloads, patents, and system and service users.

In addition, the societal impact of both fundamental and application-oriented research is often evaluated w.r.t. its contributions to SDGs. In the group, applications are often related to smart transportation, renewable energy, and bioscience, thus contributing to SDGs 3, 7, 11, and 13.

The group intends to keep monitoring its performance with respect to the above indicators and the contributions to SDGs.

Synergies: cross-cutting themes

We aim at realizing latent potentials for synergies in the group by identifying and focusing some of our activities around selected, cross-cutting themes. Cases exist where group members study the same types of techniques in different settings or application domains. By bringing together such group members, we expect to leverage each others' insights. We expect that this will improve research quality and focus, and we believe that the creation of new connections has the potential to improve the innovative nature of our research. Envisioned means include meetings, talks, and workshops and also shared Ph.D. students and postdocs.

As an example, we study machine learning in different applications where the same types of learning are often studied in different settings. For example, deep learning-based time series analytics are studied both in transportation, where traffic time series show strong periodic patterns, and in renewable energy, where energy supply time series are subject to external factors such as wind speeds. Bringing together the staff working on machine learning in different application domains holds the potential for advancing the techniques and sharing application-specific knowledge and insights, thus benefiting all participants.

As another example, we study the theme of ingestion of data into databases in different application domains, including data warehousing, mobility analytics, and energy analytics. Data ingestion covers topics such as extract-transform-load, main-memory index updates, multithreaded processing, and streaming. Specifically, some members of staff have studied techniques for extract-transform-load in the setting of general multi-dimensional data. Bringing this staff together with staff working in the setting of mobility analytics holds the potential for advancing the techniques as well as for benefiting the studies in mobility analytics.

Other types of techniques to be considered encompass: scale-in (scale-up), including main-memory and multi-core processing; scale-out, including Spark and other types of parallel and distributed processing; dimensional query processing; predictive techniques; best practices for experimental studies, including documentation and automation; open data and services practices; open source infrastructure.

Synergies: cross-fertilization

We aim at realizing latent potentials for synergies by means of sharing of data, code, concepts, and knowledge and collaboration between foundational and applied research. Although GDPR sets up barriers for sharing, especially sharing of data, the group will still benefit by sharing of other research artifacts, e.g., code and ideas, while ensuring GDPR compliance. Thus, the group plans to focus on and facilitate sharing in the next period.

No single group member can focus on and excel in foundational research and applied research at the same time. While working in the same general topic area, different members of staff position themselves differently along the foundational-to-applied dimension. A collaboration between foundational and applied research is established by means of collaboration among these group members. The potential benefits are many: Those focusing on more foundational aspects get inspiration for new and relevant research problems and may see their results being implemented in systems that create value for collaborators and society in an accelerated fashion. Those focusing on challenges closer to applications may benefit from the availability of foundational insights and advanced techniques.

The most notable collaborations between foundational and applied research in the group are found in the areas of mobility analytics and energy data analytics. In these areas, the group has a demonstrated ability to publish in the most prestigious scientific outlets, to make prototype systems available, to create value for its external collaborators, and to co-found several spinouts. We aim to maintain close collaboration and to achieve additional synergies.

Robust and diverse funding portfolio

The group strives to maintain a robust and diverse funding portfolio, where many group members are able to obtain funding and where different types of funding, e.g., funding in the excellence and strategic pillars, can be obtained. The group has been very successful at achieving this in the current reporting period and plans to strive to maintain a robust and diverse portfolio in the next reporting period.

Hiring

Hiring highly qualified staffs in the core research areas of the group remains challenging. In addition, some current members may leave for other positions. Thus, the group plans to continue hiring highly qualified staffs at all levels.

Laboratory support

The group is developing, maintaining, supporting, and using increasing amounts of software stacks. This is important for long term research, industrial collaboration, and uptake, and it also contributes to the visibility of the group.

Examples include that the Danish police have used a Daisy speed-map for determining where to do speed controls and that municipalities have been able to study traffic congestion at a more fine-grained level than previously possible. These activities are increasingly labor-intensive and are difficult to sustain via external funding, in part because externally funded projects are relatively short-lived and because they often must address new challenges. In addition, funding agencies generally do not offer means of sustaining their investments after the funding period. Thus, the group would benefit substantially from longer-term laboratory support in terms of both hardware and laboratory technicians, e.g., through the department's laboratory initiatives. For example, software laboratory technician support may help turn the group's large tool portfolio within digital energy, e.g., the FlexOffer stack, into a unified, robust, and well-documented open source resource for larger societal impact in relation to SDGs 7 and 13. The group hopes to achieve various support from the department in the coming period, to the benefit of many research activities in the group.

Ability to seize opportunities

We aim at building our capacity for seizing opportunities for achieving further research activities and results in and beyond our core research areas. For example, this includes our ability to compete for funding in areas where we expect future funding to be available, such as artificial intelligence, machine learning, big data, mobility, smart cities, and renewable energy. It also includes our ability to collaborate with industrial partners.

Means of achieving this aim include a continued focus on relevant networking and collaboration. In addition, we will consider the feasibility of having some group members proactively specializing in areas that show particular promise, such as machine learning.

Next, we cover three clusters of envisioned research topics for the coming period. Each cluster is described briefly in the following.

Spatial, temporal, and sensor data management and analytics.

For the research within spatio and spatio-temporal, the group has a very strong outset in using GNSS (GPS) data both as single points and as trajectories. A goal for the new period is to retain this focus on using GNSS data but extending the research to cover a broader range of vehicle types, e.g., bicycles, e-bikes, and trucks. These new vehicle types are relevant for SDG-11. Further, there are new challenges with the new vehicle types, e.g., map-matching, and varying sampling period. If relevant for externally funded projects looking at AIS (GPS) data from ships is an interesting extension.

A number of new data types are becoming available due to the opening of data sources and very interesting development in new sensor technology. Examples of these data types are fuel/energy consumption, induction-loop, photo, radar, video, and air-pollution data. A goal for the group is to actively include these new spatial/temporal data types and integrating these with existing data, as they become available. The extension with data types raise new research challenges related to large-data management, data mining, query-processing algorithms, and query optimization. Supporting more data types also broadens the possibility of collaboration with companies and universities.

In the coming period, we anticipate growing activities on time series analytics and data-intensive routing with extensive use of machine learning. First, we expect to continue our work on time series outlier detection and forecasting, while focusing on explainability, robustness, and efficiency in resource limited environments. Second, we expect to continue our work on data-intensive routing by means of reinforcement learning and imitation learning. In addition, we expect to work with efficient stochastic routing algorithms under different paradigms that are able to take into account the travel cost distribution dependencies.

Finally, we will continue the work on spatial crowdsourcing, spatial entity linkage, and spatio-textual analytics started in the previous reporting period.

Deep and extreme-scale analytics

Building on our previous work in advanced analytics and data warehousing, we will perform research in the area spanned by two directions: deep analytics and extreme-scale analytics.

Here, “deep” refers to either a deep technology stack with several layers, e.g., predictive and prescriptive analytics, or a deep application domain, mainly digital energy. Within predictive analytics, we will investigate machine learning and deep

learning for IoT timeseries, including time series forecasting and effective model selection and update. Within prescriptive analytics, we will further extend SolveDB functionality and transform it into a Big Data platform based on Spark. Within digital energy, we will continue our successful work on FlexOffers, extend these to handle energy storage and conversion, and develop global ontologies for energy and smart building data, flexibility metrics, and trading strategies for existing and emerging energy/flexibility markets, e.g., blockchain-based P2P trading. We will also develop machine-learning based techniques for flexibility extraction and prediction.

Next, "extreme-scale" refers to either extremely large and fast IoT time series and/or extreme variety data lakes. For the former, we will co-develop novel IoT data transfer and compression algorithms with researchers from wireless communication and extend our work on using mutual information for time series mining. We will integrate the results into our model-based time series management platform ModelarDB, which will be extended further to support edge computing, new model types, and more advanced analytics. For the latter, we will continue our work on the Danish National Data Lake for Energy at CenterDenmark, which will be scaled to many thousands of very diverse and possibly very large datasets. To support this, we will work on explorative analytics using a data cube metaphor and adding deep semantics on top of the data sets.

Semantic data management, graph databases, knowledge engineering, natural language processing, and recommender systems

Building upon our previous work in this area, we will continue our research on decentralized knowledge graphs to enable efficient query processing, availability, provenance, and collaborative updates. As knowledge graphs are evolving over time and old versions typically become unavailable, we will investigate how archiving and versioning can be enabled efficiently while keeping the data queryable. In consideration of knowledge graphs with high update rates, so-called dynamic knowledge graphs, our research will cover aspects such as anomaly detection, vandalism attempts, and embeddings and link prediction. Furthermore, we plan to work on aspects of privacy, such as privacy-preserving querying of knowledge graphs and privacy-aware data integration in general. Moreover, we plan to continue our research on semantic data integration, ETL, (spatial) OLAP, as well as example-driven analytics of graphs, which we will expand to new applications and use cases. We plan to intensify our research efforts in the area of Natural Language Processing by developing methods that can make efficient use of language similarities and

differences in particular to support truly low-resource languages. Finally, we plan to continue the effort on developing new models for recommendation based on embedding methods (both matrix factorization and neural methods) with information from multiple views.

Data science in the context of interdisciplinary research

Our interdisciplinary efforts will be strengthened in the next period in the context of several upcoming or just starting projects. First, together with colleagues from Bioscience we will improve the rate of recovery of microbial genomes by leveraging, expanding, and combining DNA sequencing, machine learning, and graph-based analysis. Second, we plan to continue our work in the Medical domain in the context of entity similarity and graph embeddings. Third, in collaboration with the Department of Planning we will work on extracting, organizing, learning, searching, and using knowledge to assist the process of environmental assessment.

3.8 Sustainable Development Goals

- SDG3: Good Health and Well-Being
  Towards the end of the period, we have started to collaborate with researchers from Bioscience and Medicine. Our initial works on microbial DNA sequencing as well as on patient data and diagnosis codes are quite promising and will be expanded upon in the next period.

- SDG4: Quality Education
  Throughout the period, quality education has been a natural concern for the teaching efforts of the group. Problem based learning (PBL) enhances the quality of the education by ensuring that the students have a high degree of motivation in major parts of their study load. In order to enhance active learning, the flipped classroom approach has been introduced in several courses. The availability of video lectures, as a contrast to traditional lecturing in auditoriums, makes the educations flexible in a way that many students experience as enhanced quality. Digital support of selected educational activities has been boosted by development of prototype tools, which play an important role in first year courses attended by many participants.

- SDG7: Affordable and Clean Energy
  For more than a decade and throughout the period, support for the green energy transition (which in the period was formulated as SDG 7) has been a

key focus for the group, with significant research, innovation, and commercialization efforts within digital energy, energy data management, and energy flexibility. The main associated research topics were the FlexOffer concept and model-based time series management for wind turbines.

- SDG 11: Sustainable Cities and Communities
  The group has a good track record related to handling vehicle data, e.g., GNSS and CANbus data. The efficient management of such diverse data sources lays a strong foundation for contributing to quantifying and monitoring the greenhouse gas and particle emissions in cities. Using the group's background in spatio-temporal query optimization and advanced traffic analytics, we want in the new period to extend the work related to this SDG by looking at new sensor types, support the decarbonization of transport, and consider near real-time data.

- SDG 13: Climate Action
  The work mentioned under SDG 7 above has also had significant implicit impact on SDG 13, as optimal digital support for renewable energy in the energy system will enable massive greenhouse gas reductions.

3.9 AI and Machine Learning

- Data Science Education
  During this period, we have designed and launched a new education in Data Science as recommended by the committee in the previous research evaluation. There is a three-year program for achieving a BSc degree. The first intake was in 2019, and the students are thus now in their 4th semester. When they finish the BSc program, they will be able to continue in the two-year MSc program. The education has a strong focus on analytics, statistics, machine learning, and data management, but involves also, e.g., law and ethics and application domains.

- Machine Learning on Time Series, Spatio-temporal Data, and Graphs
  During this period, we spent significant efforts on machine learning, especially on time series data, spatio-temporal data, and graphs. On time series data, we have proposed new learning algorithms for correlated time series prediction and outlier detection. We studied learning to route techniques where we learn knowledge, e.g., routing preference, travel time uncertainty, from large amount trajectory data to enhance vehicle routing. We also studied

representation learning on paths in graphs, more specifically, road network graphs. The results have been published in top venues in artificial intelligence and data management, such as IJCAI, PVLDB, ICDE, The VLDB journal, and TKDE. We expect that increasing efforts will be put into this area in the next period.

- Machine Learning on Energy Data
  During the period, significant work was done on machine learning for energy data, in particular for time series of energy consumption with the aim of predicting/extracting available energy flexibility. Here, we applied and extended a number of machine learning techniques with good results. The work on discovering correlations in time series based on mutual information also falls under this topic. Results were published in top venues and applied in real-world settings. We expect to increase our efforts in this area in the coming period.

- Recommendation Systems
  During this period we have also spent significant effort on machine learning in recommendation scenarios. We have provided a hybrid learning model for collective matrix factorization and we have also developed deep neural models for joint collective embeddings, models utilizing convolutional adversarial networks as well as models which jointly learn recommendation and explanation function for recommendations from rating data and from reviews in context aware scenarios. The work was published for example at ACM RecSys and at ACM IUI conference.

3.10 Committee Evaluation

The report "Research Evaluation 2016-2020" authored by the Database, Programming and Web Technologies group (in short, DPW group), as well as the accompanying Appendix are quite extensive and informative. The on-line presentation on the 16th of April 2021 was also very informative and helped to clarify several issues. Based on the above the observations are as follows.

Observations The DPW group hosts Daisy, the Center for Data-intensive Systems, and research-wise it is focusing in areas related to Data Management, Programming Technologies and Web Engineering. The area of Programming Technologies was part of the DPW group until Fall 2020, when it moved to the DEIS group. In particular, by using the term Data Management, one includes the fields of

Databases, Data Warehouses, and Data Analytics, including querying, data mining, and machine learning. Incidentally, working in these large areas gives opportunities to tackle problems related to Information Systems, Social Media, Bioinformatics, as well as energy-related, engineering-related, and production-related issues.

Even before the current evaluation period, the DPW group had established itself as a leading group, among the top groups in Europe in its area. The group has improved its performance further, as evidenced by key-performance indicators, including its journal and conference publications in high-ranked outlets, its success at obtaining highly competitive excellence- and strategic-pillar funding, its extensive external/industry cooperation, and its international exposure and recognition (awards, service in leadership roles, and keynote/invited lectures). In particular, the group ranks number one in Europe in its area during the 2016-20 period according to csrankings.org.

Notably, it is not merely quantities that are impressive but quality issues are equally honored. In particular, group members are publishing mostly in the very prestigious Core A*/A categories and in the Danish BFI Level 2. Most of the published works appear in high-prestige journals and conference proceedings. Consequently, the published works succeed in attracting impressive numbers of citations. This way, production, excellence, and impact complement each other. Further, this is reflected in the personal profiles of the group members, which demonstrate from satisfactory to significantly increasing citation profiles (e.g., the newly appointed full professors Katja Hose and Bin Yang as well as Hua Lu and Peter Dolog), up to impressive numbers of citations and h-index values (e.g. Christian S. Jensen and Torben Bach Pedersen).

Further to the previous, it should be mentioned that a significant number of young persons are trained by the senior and junior academic staff, either as researchers in projects or as doctoral students, and contribute to the group performance. The number of Ph.D. dissertations during the evaluation period is indicative of success in this dimension.

As a service to the community, group members have participated as PC members in a huge list of international conferences, plus they have assumed key roles in many prestigious conferences (e.g. VLDB, ICDE, DASFAA, ESWC, SSTD, WISE, MDM, DOLAP, etc. etc.). In addition, they have contributed to the Encyclopedia of Database Systems with a huge number of lemmas, as well as they were heavily involved in editing conference proceedings (e.g. ICDE, EDBT, ESWC, SSTD, ApWeb-WAIM etc).

A remarkable number of complex system prototypes (e.g., GOFLEX, ModelarDB,

pygrametl, SimpleETL, SOLVEDB) have been built as part of project work and successfully transferred to open source systems. This is well received and is a perfect way to raise the visibility and even more the real-world impact of the DPW group in many aspects, personally and technically.

As a final token about the achievements of the group, it should be mentioned that two spinouts have been founded by a DPW member (Torben Bach Pedersen), thus contributing to putting the accumulated research know-how into industrial practice. Former graduates of the group like Kostas Tzoumas and Morten Middelfart have been very successful entrepreneurs.

The DW group very successfully contributed via co-operations and projects to applications in the transportation and energy domains. With this, SDG topics 7, 11, and 13 are directly addressed and significant advancements to clean energy and climate are supported. Contributions to the areas of health and well-being are also remarkable.

There are a lot of successful activities related to machine learning and artificial intelligence (e.g., ML on time series, spatial-temporal data, graphs, ML on energy data or Data science education). Most of these activities are linked to the crosscutting team AI&ML. This way the DW experts join forces with experts from the other groups in order to build strong teams. Many AI and ML topics are truly interdisciplinary; hence, this crosscutting team is an effective and successful measure.

The net result is that the DPW has demonstrated world-class performance in the evaluation period, combining academic excellence and impact with industrial and societal impact. It is one of the top groups in Europe in its area, as demonstrated by its record of publications in prestigious outlets.

Recommendations Being a top group is a challenge with respect to keeping the current level of performance. A few possible paths to maintain the current status follow.

- Efforts should be made towards applying for high-reputable funding schemes, such as ERC grants, either for senior or for younger scientists; we clearly see potential candidates!
- Having heavily published in traditional journals and conference related to Data Management etc., they could try other outlets of this level, such as Nature for instance, or Communications of the ACM or Proceedings of the IEEE, which are historical and archival journals with very broad circulation.

- As data is the oil of IT, the group should further intensify efforts to increase the visibility and keep on attracting industry funding and co-operation as well as interdisciplinary projects with internal partners from other AAU departments and/or external partners from industry or other universities. To accomplish this, a strategy should be defined that exploits on one side existing project experiences and existing partner networks and on the other side invite young and new group members (assistant and associate professors) to join. This way the new colleagues get educated on the job.
- In addition, we fully support the future research topics as mentioned in the available material; that is spatial, temporal, and sensor data management and analytics, deep and extreme scale analytics, semantic data management, graph databases, knowledge engineering, natural language processing, and recommender systems, and finally, data science in the context of interdisciplinary research. Some of these topics will fuel ongoing and future joint activities in the crosscutting team AI&ML.
- Complementary to the previous, members of the group could examine to establish new spinouts, considering this outlook as a natural consequence of their high quality research. The prototype systems mentioned above may be good starting candidates.
- Members of the group should consider educational issues as well. For instance, they could author books, either textbooks or extensive monographs, aiming at further leveraging the outlook of the group at international level.
- Lab staff support is rather necessary to hire, esp. for prototype systems that result from time-restricted projects. These prototypes may shorten ramp-up phases for future projects and it may increase visibility (e.g., if open-sourced) or even spin-out activities.
- Pay attention to the numbers of students at all levels, in particular at the doctoral level. Such a top-performing group (as well as the department as a whole) should provide exchange of junior staff with other European universities.
- By renaming to "Data and Web" (or similar) the group becomes more coherent as the difference between databases and web is not the scale but the nature of the data (structured vs unstructured). Probably, the term "Big Data" or/and "Intelligent Data Management" could be emphasized as already such research directions are followed in the group.

4 Distributed, Embedded & Intelligent Systems

4.1 Executive Summary

The research of the Distributed, Embedded and Intelligent Systems group concerns modelling, analysis, and realization of computer programs, with an emphasis on distributed, embedded systems and intelligent. The contributions of the group include work on semantic foundations, algorithms and tools for verification and validation, analysis and construction of networks, security, model-based methodologies, probabilistic graphical models, machine learning as well as applications in a variety of domains. The research activities are both theoretical/foundational in nature as well as practical/experimental. The group host CISS, a regional ICT competence center conducting world-leading research within embedded systems with a long history of industrial and international collaboration.

During the period DEIS performed very well with respect to key performance indicators such as publications, external funding, Ph.D. education, scientific service, academic and industrial collaborations as well as dissemination.

DEIS has published 317 peer reviewed journal and conference publications, one book and 10 edited proceedings. Out of these 111 are in top-outlets within our field(s) and 98 are in top specialized outlets. DEIS has participated in a total of 39 projects, with some 75.4 MDKK in funding. The type of funding spans very prestigious (and big) personal grants, large industrial collaborative projects, as well as dissemination projects, with the major part of the funding coming from basic and strategic research projects. Several projects address one or more SDGs, and are conducted in collaboration with other groups at the department or the university. Two spin-out companies has been established. The most recent acquired grant is the Danish National Research Center in Digital Technologies (DIREC), which will provide a unique opportunity for DEIS and the entire department to engage in new research collaborations with other CS departments of Denmark. During the period 9 Ph.D. degrees has been awarded, and at this point there are 12

Ph.D. projects in progress. The group has maintained its strong commitment to provide state-of-the-art model checking tools (e.g. UPPAAL and TAPAAL) with new branches emerging, aiming at wide industrial take-up, experiencing extremely many downloads, and winning several gold and silver-medals at model checking competitions.

Over the period DEIS has been success full in promoting new full professors, attracting top candidates in software defined networks, quantitative modelling as well as machine learning. At the same time DEIS suffered a loss of two key members both being called for Full Professorships outside Denmark. In both cases research collaboration continues. At the beginning of the period the former DES (Distributed and Embedded Systems) group and members of the MI (Machine Intelligence) group were merged, and at the end of the period the Programming Technology sub-group of the DPW (Database, Programming, and Web Technology) group joined DEIS. The DEIS group has received substantial recognition from peers and society, with several best paper awards, prizes and honours. Also the group has made extensive services to research community.

In conclusion, we believe that the research activities carried out by DEIS in the period has been excellent. In particular, the DEIS group believes that its publication performance in top outlets, contributions to verification tools, level of industrial collaboration and impact on societal challenges compare favorably with similar groups across Europe.

4.2 Profile

The research of the DEIS group concerns modelling, analysis and realization of computer programs, with the emphasis on distributed, embedded and intelligent systems. The research contributions of the group covers foundational theories, verification and optimization tools, validation methodologies, security, probabilistic graphical models and machine learning focusing on distributed, embedded and intelligent systems.

In 2001 the Centre for Embedded Software Systems (CISS) was established as a collaboration between the former Distributed and Embedded Systems (DES) group and research groups from the Institute for Electronic Systems at Aalborg University, with the ambition to create an industry oriented research center of excellence. After nearly 20 years, CISS has established itself as a preferred partner having coordinated and participated in several European and Danish national projects, including both basic research projects as well as large-scale industrial projects and

focused dissemination projects.

As a result of the last research evaluation, the former Machine Intelligence (MI) and the DES group joined forces into what is now the research group for Distributed, Embedded and Intelligent Systems (DEIS). Leveraging on these two strongholds, the group has in the current evaluation period (2016-2020) broadened its research scope towards intelligent and cyber-physical systems (CPS) in which networked and learning-enabled computing nodes interact with the physical environment, collaborate, and operate in context of other systems and human users.

The research of DEIS spans the entire range from foundational research, strategic research to industrial application and includes the following key research areas:

Foundational Theories Here the group continues its long track record within Concurrency Theory, with focus on mathematical and logical theories for modelling, specifying and verifying concurrent processes, quantitative aspects, in particular timing and stochastic aspects. Our work in this period include axiomatization, decidability and complexity of various quantitative models and logics, as well as revising the semantic foundation of higher-order programming to that of probabilistic programs. Moreover, we have continued our work on behavioural metrics as a way of quantifying degrees of correctness. This research is strongly connected to an effort in giving convergence guarantees of various machine learning methods.

Verification, Validation and Optimization The group has a long track record in model checking, contributing with state-of-the-art tools, e.g. UPPAAL and TAPAAL. This work includes algorithms and data structures for model checking, testing, performance analysis, synthesis and optimization for complex systems. In the period focus have been on applying and developing Dependency Graphs as a generic framework for obtaining efficient on-the-fly model checking, and on extending Partial Order Reduction techniques to real-time systems. The effort on statistical model checking has also continued, with focus on efficient estimation of rare events, and efficient run-time probability estimation of temporal properties. The challenge of synthesis has received particular attention with new learning methods for obtaining optimal yet guaranteed safe control strategies. Finally, adapting model checking towards scheduling and planning is been an ongoing research topic.

Networks Analysis and construction of services and protocols for networks. Particular effort has been made on software-defined networks, resilient and secure routing in network constructions, demand-aware optical networks, network

security and network verification. Finally, wireless and highly decentralized networks has been studied.

Security Besides security in software defined networks, focus has been on security in socio-technical systems. Here leveraging on our stronghold in quantitative model checking work has been made on attack-defence trees as a formalism for describing and analysing potential attacks.

Machine Learning Machine learning using statistical as well as logic and relational-based methods. Historically, and in the current period, members of the group have made several contributions to the field of Probabilistic Graphical Models (e.g. Bayesian networks). In the field of relational learning we are investigating foundational principles and algorithmic solutions for modeling, learning, and inference. Connections between machine learning and formal methods have been explored in several directions.

Methodologies The DEIS group has in previous and in this period made substantial effort towards industrial take-up of model-driven and contract-based development (testing and verification) using our UPPAAL tool suites. Also DEIS has collaborated with the Programming Technology sub-group of DPW ondefining and using Safety Critical Java.

Applications areas include power electronics, drones, satelites, intelligent transport, green energy, health, epidemics, and clean water.

Summarizing, our research benefits companies and organizations that work with safety critical software systems where requirements on monitorability, predictability, security as well as safe and intelligent decision making are crucial. Other beneficiaries include the several academic and industrial users of our award-winning tools UPPAAL and TAPAAL.

Example academic partners include INRIA Rennes, TU Wien, TU Eindhoven, Strathclyde University, Oxford University, Northeastern University, Trento University. Example industrial partners include NTNU, Grundfos, City of Aalborg, Aarhus Vand, Huawei, Nilfisk, Neocortec, ATS, Ambolt, Neogrid, Hardi, Seluxit, and GOMSpace.

4.3 Staff

4.3.1 Current Staff (2020)

We report the current situation of the DEIS group as to January 2021.

The DEIS group currently consists of 2 Full Professors, 2 MSO Professors, 11 full-time Associate Professors, 7 assistant professors (among which 3 with tenured track position), 7 postdocs, and 11 Ph.D. students (see Figure 4.1). These are complemented by 5 Affiliated part-time Professors, 2 full-time Tool developers, and 2 full-time Technicians.

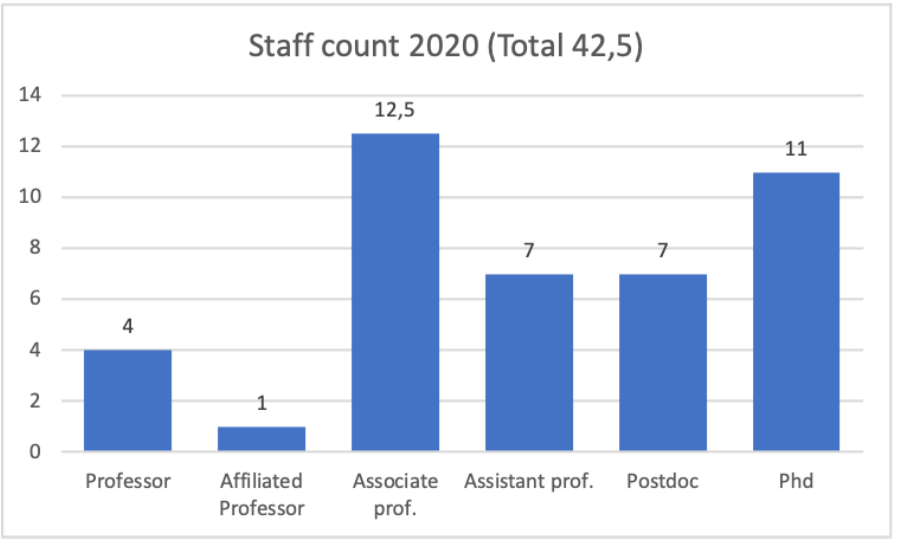

Figure 4.1: Current staff count summary. The counting is weighted w.r.t. to percentage of employment.

Professors: Kim Guldstrand Larsen, Jiri Srba, Thomas Dyhre Nielsen (MSO), Bent Thomsen (MSO).

Associate Professors: Alvaro Torralba Arias De Reyna, Arne Joachim Skou (40%), Brian Nielsen, Giorgio Bacci, Giovanni Bacci, Kristian Grønborg Olesen, Manfred Jaeger, Max Tschaikowski, René Rydhof Hansen, Ulrik Mathias Nyman, Josva Kleist (10%), Kurt Nørmark, Lone Leth Thomsen, Thomas Bøgholm.

Assistant Professors: Michele Albano (tenure track), Peter Gjøl Jensen (tenure track), Florian Lorber (tenure track), Alessandro Tibo, Anders Schlichtkrull (tenure track), Danny Bøgsted Poulsen, Marco Antonio Muñiz Rodriguez.

Postdocs: Daniele Toller, Junior Koffi Eveque Dongo, Martijn Angelo Goorden, Hessam Golmohammadi, Nicolas René Jean Schnepf, Saeed Shakibfar, Shiraj Arora.

Ph.D. students: Andrej Kiviriga, Imran Riaz Hasrat, Martin Kristjansen, Mathias Claus Jensen, Muhammad Naeem, Morten Konggaard Schou, Anders Mariegaard, Napalys Kličius, Søren Enevoldsen, Anton Christensen, Frederik Meyer Bønneland, Jonas Hansen.

Permanent Tool Developers: Marius Mikučionis, Kenneth Yrke Jørgensen.

Technical Staff: John Bjerregaard-Michelsen, Per Printz Madsen.

Affiliated Professors: Anders Læsø Madsen (20%), Axel Bernard E. Legay (20%), Iulian Radu Mardare (20%), Andrea Passerini (20%), Pieter Jan Laurens Cuijpers (20%).

4.3.2 Staff Development

Compared with the previous evaluation period, the DEIS's scientific staff has increased from a grand total of 32 scientific personnel to a current total of 49 scientific personnel (including part-time employees and partially affiliated staff members). The increment in staff was proportionally distributed among all the different positions. Over the period 2016-2020 the staff evolved substantially. The total number of staff members within the DEIS group remained approximately the same from 2016 to 2019, and experienced a notable increment in 2020 (see Figure 4.2). These numbers reflects the economical situation of the department.

In the following we report the most significant staff development events occurred within the period 2016-2020. A more detailed historical account can be found in the Appendix.

Strengthening of Distributed Systems: To strengthen the area of distributed systems and network, in 2015 Stefan Schmid was hired as Associate Professor.

Merge with Machine Intelligence: After the end of the previous evaluation period, the DES (Distribuited and Embedded Systems) group was in June 2016

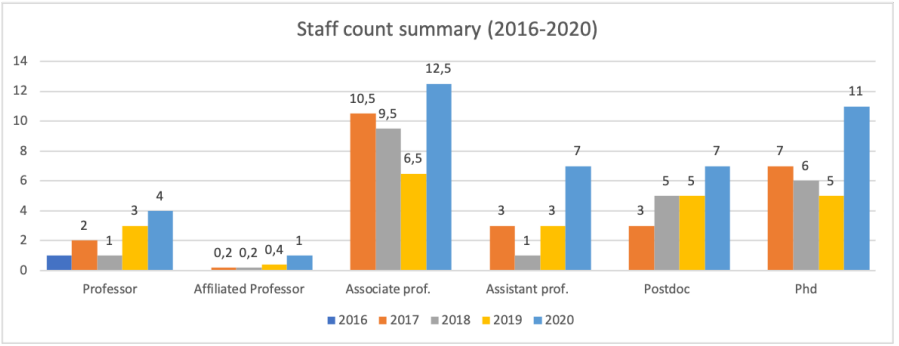

Figure 4.2: Staff summary for the evaluation period 2016-2020. The counting is weighted w.r.t. to percentage of employment.

merged with part of the MI (Machine Intelligence) group, increasing the group by 2 Associate Professors (Manfred Jaeger and Thomas D. Nielsen). To account for this structural change the group was renamed to *Distributed, Embedded and Intelligent Systems* (DEIS).

New Professors: Following the recommendations of the evaluation panel from the previous Research Evaluation, in 2019, the DEIS group increased the number of professors within the group by promoting Radu Mardare and Jiri Srba to MSO Professors. In late 2019, Jiri Srba got promoted to Full Professor and Thomas D. Nielsen to MSO Professor.

Loss of Key Members: The DEIS group suffered two important losses of personnel. At the end of 2018 Stefan Schmid left the group accepting a Full Professorship position at the University of Vienna (Austria); and at end of 2019 also Radu Mardare left for a Full Professorship position at the University of Strathclyde (Glasgow, Scotland).

The collaboration with Stefan Schmid, is still active through the DFF project *Quantitative Analysis and Synthesis of Network Protocols* as well as through the Vienna Research Project *WHATIF*. The same is with Radu Mardare, who was offered a part-time Professorship position (20%) sponsored by the ERC-Advanced grant LASSO.

Poul Due Jensen Professorship: Max Tschaikowski joined the group as *Poul Due Jensen* Associate Professor, a position specially funded by the Poul Due Jensen/Grundfos foundation for supporting the hiring of a new international

researcher with an additional "starting package" support for the hiring two full-time postdoc positions.

Strengthening of Machine Intelligence: To strengthen the area of Machine Intelligence, in 2020 Alvaro Torralba (for the area of planning) was hired as Associate Professor and Alessandro Tibo (for the area of deep learning) as Assistant Professor.

Merge with Programming Technologies: In September 2020, another major structural group change occurred. The *Programming Technology* sub-group of DPW, joined the DEIS group. This resulted in and increment of the DEIS group staff by 1 MSO Professor (Bent Thomsen) and 3 Associate Professors (Kurt Nørmark, Lone Leth Thomsen, and Thomas Bøgholm).

In conclusion, the permanent staff increased considerably with respect to the previous evaluation period 2011-2015. The actual increment of the numuber of staff members, however, has to be contextualised with the merge with the Machine Intelligence group and the Programming Technology sub-group, which did not contribute to the increment of the overall departmental staff.

With respect to the previous evaluation period, the DEIS group hired considerable number of new staff: 2 Associate Professors, 3 Assistant Professors and 2 full-time tool developers and several Postdocs. The goal of the group for the coming years is to continue hiring new scientific and technical staff to strengthen the group.

4.4 Goals (2016-2020)

As a result of the previous research evaluation the current DEIS group was formed from the former DES group and part of the former MI group. Here we recall the stated goals for 2016-2020 for DES and MI.

4.4.1 Goals for DES: Distributed and Embedded Systems

Research Directions We will strive to maintain our strong international collaboration and leading role within Embedded and Cyber-Physical Systems, and our high international visibility and recognition in the areas of model checking, logics and semantics for concurrent systems. We will strengthen further our focus on tool-development. In order to scale to complexity of CPS, techniques from other areas such as SMT and machine learning will be exploited, as well as taking full advantage of the recently established MCC cluster. The current trend in our

foundational research in replacing Boolean verdicts with metric notions and quantitative results will be expanded. We will invest substantial new effort in research on cloud computing, (wireless) networks and IoT.

Collaboration and Projects We want to maintain current balance between basic research and application-oriented projects. Our strong international collaboration must be sustained. We have been extremely successful in the ARTEMIS program, but significant effort is required to ensure similar performance in the successor ECSEL, though we are already involved in two new projects ECSEL MANTIS and ENABLE-S3. We also intend to be partner/coordinator of future FET projects and continue our active participation in (future) COST actions. In terms of national collaboration, the directing role of the Innovation Network InfinIT provides us several possibilities. In collaboration with other groups of the department we want to pursue new collaboration with SDU and ITU within the area of Smart Society. Within the department, significant collaboration with other groups is already planed in the context of various projects, e.g. DiCyPS, PPP as well as the Innovation Network InfinIT. Also, the ERC Advanced Grant LASSO provides an opportunity for closer collaboration with researchers from the MI group. Within the university, we want to pursue new, strategic collaborations with Department for Electronic Systems within WSN, IoT and Security.

Research Management and Staff To minimize the overhead on researchers to engage in industrial collaboration with SMEs, the CISS VIP staff must be increased, and should be involved more directly in research projects. We wish to identify and establish better ways of managing our several and interconnected projects. Furthermore, we wish to strengthen the alignment and synergy between MSc/PhD/post-doc projects with industrial research projects, and initiate exploitation of existing financial sources for industrial PhDs/post-docs. Finally, we encourage members of the research group to plan sabbaticals in order that new ideas from the outside may be brought in.

4.4.2 Goals for MI: Machine Intelligence

General Plans The main objective for the next 5 year period will be to overcome the currently critically low staff level, and to re-establish a group size at which the MI group can fulfill its teaching and administrative obligations, while also pursuing new research initiatives. The short term goal is to find a replacement for the

associate professor who recently resigned from his position. In the longer term, the target is a moderate growth by an additional 1-3 assistant and associate professors.

Apart from growth through regular appointments, we will try to enable additional research activities through continuing efforts to obtain external funding, and through collaboration in externally funded research projects of other research groups at the department.

Concrete Plans We will continue our methodological research on learning from relational data. An important application domain that we plan to pursue within a European research project is modeling, learning, and decision support for sustainable management of water and land resources.

Our recent interdisciplinary work on automata learning and model checking has opened a new line of research that looks more broadly at the integration of machine learning and formal system analysis. In collaboration with the DES group, we will work on practical and theoretical aspects of this integration. Practical aspects that we will consider include the use of reinforcement learning techniques for the synthesis of optimal controllers, and the use of Bayesian networks for system diagnostics.

We plan to continue our research on latent variable models, focusing on regular model structures inspired by the model classes that we have previously established for collaborative filtering. For this application domain, in particular, we will consider methods for integrating different types of data sources, including both structured and unstructured data as well as network data representing, e.g., social relations among customers.

Our activities within streaming data was initiated in 2013/2014 with the start of the EU FP7 project AMIDST. We wish to extend and build on top of the current results achieved within this project.

4.5 Activities and Results

We report on the research of the DEIS group during the period 2016-2020 within the following focus areas: *Foundational Theories; Verification, Validation and Optimization; Networks; Security; Machine Learning; Methodology and Applications; Software Tools.*

4.5.1 Foundational Theories

This section collects and summarises the works the DEIS group contributed to in several areas of the foundation of computer science.

Behavioral Metrics

For the comparison of the behavior of quantitative systems there is a need for replacing classical Boolean verdicts of correctness and equivalences, with quantitative notions formalized as metrics able to measure the degree of "adequacy" between the system and its specification. The DEIS group has a long record of publications on behavioral metrics which has continued from previous reporting periods. In this reporting period we studied expressiveness, axiomatizability and computational aspects of behavioral metrics.

Expressiveness and Computational Aspects. We studied behavioral metrics for different types of quantitative systems, ranging from continuous-time Markov chains [25], semi-Markov decision processes [220], weighted Kripke structures [91, 137], and probabilistic automata [29]. In [33] we studied the approximate-minimization problem for labelled Markov chains, i.e., the problem of finding the automaton of a given size with minimal distance from the original one. We showed that the complexity of problem lays in between the classed NP and PSPACE and developed heuristics that use a variant of the Baum–Welch algorithm for providing sub-optimal approximants in polynomial time.

In [29] we studied the connection between the bisimilarity pseudometrics and Condon's simple stochastic games, providing an effective algorithm for computing the distance employing a variant of the simple policy iteration algorithm. Despite the problem of computing the distance lays in the intersection of NP and co-NP, experiments showed that our algorithm performs well in practice.

We investigated the question of how to compare two semi-Markov processes with respect to their time-dependent behavior [220, 219]. To this end, we have introduced the relation of being "faster than" and a hemimetric between processes measuring the least acceleration factor by which a process needs to speed up its actions in order to behave at least as fast as another process. We have studied these two notions from an algorithmic perspective, studying their algorithmic complexity and their connections to probabilistic automata.

Axiomatizability. In [191] we proposed a quantitative extension of the of the equational reasoning framework, called *quantitative equational logic*. This work followed up with a series of papers and applications. On the purely model theoretic perspective we provided a quantitative extension of the Birkhoff variety theorem [192] providing a better understanding of the nature of quantitative algebraic effects and their axiomatizability. Successively, we investigated the compositional aspects of qualitative equational theories [32], showing that the disjoint union of two quantitative theories corresponds to the coproduct of monads as computational effects. The application of quantitative algebraic reasoning has been further investigated in a series of work aiming at showing that many natural metrics used in the context of approximate verification of probabilistic systems can be axiomatised using the framework of quantitative equational theories. In [193] we showed how Wasserstein metrics arise naturally as the free algebras satisfying the axioms of Stone barycentric algebras and an extra quantitative interpolative axiom. In [24, 27] we provided a complete axiomatization of the probabilistic bisimilarity distance on Markov chains. In [26] we showed that the total variation distance (a.k.a., probabilistic trace distance) on Markov chains is obtained just by adding an extra distributive axiom.

Logics

We continued the work started in the previous reporting period regarding the development of proof systems for logics aiming at achieving sound-complete axiomatizations against various stochastic semantics. In [96] we provided an unrestricted Stone-type duality for Markov processes against an alternative definition of Aumann algebras. In [201] we investigated a modal logic for expressing properties of Markov processes with real-valued semantics based on the mathematical theory of Riesz spaces. The key theoretical contribution is the use the duality theory of Riesz spaces that provided the connection between Markov processes and the logic.

In a series of work [168, 167, 111] we took the challenge of solving long-standing open problems regarding weighted logics involving clocks (including the alternation-free weighted μ-calculus), and their decidability and axiomatization.

As for multi-weighted logic in [136], we have shown that both the model checking and synthesis problems are decidable and respectively EXPTIME-complete and NEXPTIME-hard. Instead, as for the multi-weighted games in [152] we showed that the model checking problem for the full logic is undecidable with already three weights. By restricting the bounds to constant upper- or lower-bounds on the individual weights we showed the same problem becomes decidable and that

the model checking problem is PSPACE-complete.

Semantics of Probabilistic Programs

Probabilistic programs add to the usual programming constructs the possibility to perform probabilistic choices during computation. There has been a renewed interest in studying probabilistic extensions of the λ-calculus. The challenge is to combine higher-order features with a probabilistic choice [95]. In [31] we developed a Boolean-valued model for stochastic λ-calculus. This allowed us to interpreted equality in a measure algebra and thus to reason about the semantics of λ-terms up-to "almost everywhere equality". This provides a new framework for reasoning about probabilistic programs and their higher-order properties.

Timed Automata, Energy Timed Automata and Energy Games

Historically, the group has made serveral contributions to the theory of timed automata. The release of the Handbook of Model Checking in 2018 is the most recent witness of the standing of DEIS in this field: the Handbook being an effort over more than 5 years involves more than 70 researcher, including members of DIES co-authoring the chapter on Timed Model Checking [56]. Also members of the DEIS group has contributed with an invited contribution of the special volume 10.000 of LNCS on the use of timed and hybrid automata for system design [10], the contribution being co-authored with Rajeev Alur and Tom Henzinger. Also, in this period we have contributed Priced and game extension of timed automata are a powerful formalism for modelling a number of dynamic resource allocation and optimization problems. In this line of research, we have studied the problem of synthesizing optimal and robust schedulers for energy games, average-energy games and energy timed automata [160, 173, 57, 35] allowing the models to run indefinitely while ensuring that the energy consumption will always remain within some given safety bounds. This problem is called energy-constrained infinite-run problem and it was proven to be undecidable for generic energy timed automata. In [35] we proved the decidability of the energy-constrained infinite-run problem in settings with both certainty and uncertainty of the energy rates for the subclass of flat segmented energy timed automata. Notably, this work received the *best paper award* at Formal Methods 2018 (FM'18).

Approximation of Euclidean Markov decision processes

Euclidean Markov decision processes are a powerful tool for modeling control problems under uncertainty over continuous domains. In [124] we developed a theoretical framework for the approximated analysis of Euclidean Markov decision processes, showing that when the Euclidean process is approximated by finite state approximations induced by increasingly fine partitions of the continuous state space the cost functions over finite time horizons become arbitrarily precise.

Highlights

Awards and Special Mentions. Among the publications recall above, two received best paper award: at at FM 2018 [35] and Computation Tools 2016 [90].

Notably, many the conference publications have been selected for the publication into special issues: ICTAC 2015 [28], CONCUR 2016 [27], PST 2016 [34], LICS 2016 [191], LICS 2017 [192], LICS 2017 [97], CONCUR 2019 [30], FM 2018 [35], FM 2019 [248], QEST 2019 [36].

We also mention the invited contributions to the Handbook of Model Checking [56] and the special volume 10.000 of Lecture Notes in Computer Science [10].

External Collaborators. Dana Scott (Carnegie Mellon University, USA –Turing Award), Dexter Kozen (Cornell University, USA), Prakash Panangaden (McGill University, Canada), Gordon Plotkin (University of Edinburgh, UK), Luca Cardelli (University of Oxford), Sangiovanni-Vincentelli (UC Berkeley, USA), Thomas Henzinger (IST Austria), Patricia Bouyer (LSV, CNRS & ENS Cachan, France), Nicolas Markey (Univ. Rennes, IRISA, CNRS & INRIA Rennes, France), Franck van Breugel (York University, Toronto, Canada), Nathaniel Fijalkow (CNRS in LaBRI, Bordeaux, France), Uli Fahrenberg (Ecole Polytechnique, Palaiseau, France), Matteo Mio (CNRS/ENS–Lyon, France), Qiyi Tang (Oxford University). Joel Ouaknine (Max Planck Institute, Saarland), James Worrel (Oxford University), Rajeev Alur (Penn University).

4.5.2 Verification, Validation and Optimization

Dependency Graphs

Dependency graphs (DG) introduced in late 1990 by Liu and Smolka provide a graph-representation of Boolean equation-systems. Building on earlier attempts,

this seminal work offers the first optimal (linear-time) on-the-fly algorithm for fixed-point computation. This result is utmost important, since numerous verification problems can be reduced into the problem of computing a minimum or maximum fixed-point assignment on DG.

While the original definition of DG focused on Boolean values, works at DEIS have extended the assignment domains to more general settings, thus allowing to capture quantitative logics and models. Already in 2005, UPPAAL TIGA introduced a zone-based version of dependency graphs for computing winning strategies for timed games. More recently, another symbolic extension of DG allowed for model checking of weighted Kripke structures with respect WCTL [135]. Most recently a sequence of papers have considered further DG extensions pursuing model checking of Markov reward models (and games) with respect to PWCTL [194, 138, 78].

In all the above extensions of DG, the on-the-fly approach demonstrates significant improvements of performance with respect to global fixed-point computation. However, each extension required a separate correctness proof and one-purpose implementation. This issue has been tackled by abstract dependency graphs that provide a uniform theory and a uniform implementation for all existing approaches [80, 77]. This work received the *TACAS'19 best paper award*. More recently [81, 79] offer a comprehensive survey of DG, underlying algorithms, extensions and applications. Finally, in the sequence of papers distributed algorithms for DG were implemented showing nearly linear speed-up [71, 70, 69]. This work received the *Petri Nets'17 best paper award*.

Petri Nets

Petri nets are an established high-level modeling formalism and appear in the analysis of qualitative and quantitative systems. By combating the infamous state-space explosion in the context of Petri nets, several efficient analysis algorithms could have been obtained. Specifically, methods for truth-preserving reductions of (temporal) logical formulae over Petri nets were devised [48], while [69] introduced on-the-fly and partial-order reduction algorithms for Petri nets. The results have been made available in the verification tool TAPAAL [142] leading it to receive a number of gold medals in the prestigious Model Checking Contest in 2018 (two gold medals), 2019 (four gold medals), and 2020 (two gold medals).

Partial Order Reduction

Partial order reduction techniques are some of the most beneficial methods for state-space reduction. For timed systems, partial order reduction is however more challenging because of the dependencies among timed events. In [50] an efficient stubborn set reduction for timed systems with urgent behavior has been proposed. The key idea was to employ classical untimed partial order reduction techniques as long as urgent behavior is enforced. The approach has been implemented in the tool TAPAAL and its applicability has been demonstrated on real-world models. Most recently, the exploitation of urgency has been applied to obtain efficient partial order reduction for networks of (extended) timed automata [169]. This is implemented in UPPAAL observing substantial reduction on real-world models. Another line of work proposed partial order reductions for 2-player games with reachability/safety objectives. There, stubborn reduction allowed to prune the interleaving behavior of both players [49]. For the special case of weighted Petri net games with inhibitor arcs, an efficient implementation in TAPAAL has been provided.

Statistical and Random Model Checking

Statistical and randomized model checking describe simulation-based methods for settling properties of behavioural models.

In randomized model checking, the given model may posses non-deterministic choices, that the underlying exploration algorithm repeatedly resolves randomly. Here the objective is to efficiently find counter examples. During the period, we have developed memory-efficient tactics for random LTL model checking, improving the efficiency of previous results [171]. In [155] a new randomized refinement checking (falsification) method for timed I/O automata is developed. The resulting implementation finds violations up to 600 times faster than ECDAR.

In statistical model checking the random generation of runs is based on a stochastic semantics of the given (Markovian) model and the objective is to estimate the probability of logical properties. In 2011 the first SMC engine was implemented in UPPAAL. The tutorial on UPPAAL SMC, published in 2015 [73], is by now cited 324 times. During the previous 5-year period, UPPAAL has been applied in several case-studies. This includes several of our own applications (for more details see Application section) ranging from long battery lifetime prediction [123], nano-satetelites [205], power-electronicts [210], attack-defense trees [112], wireless sensor networks, COVID-19 [139] and several others. In addition, UPPAAL

SMC has been applied outside DEIS in numerous case-studies, e.g. the work on maintenance analysis and optimization via UPPAAL SMC by Marielle Stoenligna et al.

Contributions to UPPAAL SMC during the 5-year period include an importance sampling framework that combines symbolic analysis and simulation to estimate the probability of rare reachability properties in stochastic timed automata [132]. Also, and most recently, we have proposed a framework for monitoring and updating, at run-time, the probabilities of temporal properties of stochastic timed automata [126]. Our method – which is demonstrated useful in various real-time applications – is based on Bayesian networks, and has been implemented by exploiting UPPAAL SMC.

DEIS has facilitated the research activities in SMC through the organization of specialised tracks at the ISOLA symposium [162, 161, 176, 163].

Synthesis

Rather than perform posterior verification of manually constructed control programs, synthesis offers a disruptive method where correct-by-construction control strategies are generated automatically.

Synthesis of the parameters of a quantitative model ensuring a given property is a first step from verification to synthesis. Here we have applied abstract dependency graphs to paramameterized weigthed Kripke structures [91]. Also implementatoin of trace-abstraction has lead to synthesis for paramterized timed automata outperforming existing tools (IMITATOR) [65, 64].

Synthesis for 2-player games on graphs includes work on Petri Net games [151], weigthed games with branching winning conditions [152]. In the timed setting, a subclass of timed-arc Petri Net games is identified for which discrete and continous strategies coincide, leading to performance improvement for several timed (automata) games compared with symbolic methods applied in UPPAAL-TIGA [140]. Also, controller synthesis for timed games with respect to objectives in a subset of MITL has been developed and implemented in UPPAAL-TIGA. Finally, a method for synthesising control strategies for continuous dynamical systems has been developed combining UPPAAL-TIGA with a set-based Euler method for guaranteeing that the synthesis is safe [175].

Work on UPPAAL-*STRATEGO* has continued in several directions. In short, by combining symbolic techniques with reinforcement learning, UPPAAL-STRATEGO allows to synthesize guaranteed safe and near-optimal strategies. Already in 2016,

UPPAAL-STRATEGO was applied to the industrial control of floor-heating systems [170]. Here the synthesis method was applied in an on-line and compositional manner. More recently, we have developed two new reinforcement learning techniques that tackle the problem of continuous state spaces via online partition refinement techniques [125]. The work in [124] provides insight into the convergence of these learning methods. A crucial aspect when synthesizing strategies is memory-complexity of their representation. While deep neural networks offer an attractive representation of strategies they do not guarantee safety. In [17] we suggest the use of decision trees and demonstrate that they provide both compact representation as well as safety safety.

UPPAAL-STRATEGO has been applied to several case-studies of our own ranging from floor-heating, traffic control, adaptive cruice control as well as safe and optimal control of railways. In addition UPPAAL-STRATEGO has also been applied outside DEIS, e.g. on energy-optimal multiprocessors, power management of wireless sensor networks and maritime autonomous systems.

The community interest in UPPAAL-STRATEGO is witnessed by several invited talks and papers, e.g. TACAS 2017 [174] and CONCUR 2019 [159].

Timed Systems and Uppaal

Timed automata are an established formalism for the modeling of timed systems and at the heart of the DEIS group. During the last two decades, they have been successfully used to model industrial applications and more recently, to provide forecasts to the COVID19 pandemic. Key to this was and is the software suite UPPAAL that has been downloaded, during the last year, more than 13000 times. The leading role of UPPAAL in the field of quantitative modeling is maintained by ensuring that the most recent scientific advances of the DEIS group are implemented and thus made available to nonexpert users. For instance, [50] introduced the first efficient partial order reduction technique for timed systems with urgent behavior. The effectiveness of the technique was demonstrated on larger case studies via an efficient implementation in UPPAAL [169]. Similarly, [80], that won the best-paper award, introduced abstract dependency graphs which unify all previous approaches while maintaining efficiency. In [65] our new trace-abstraction based method provides solutions to problems for (parameterized) timed automata which are unsolvable by the current state-of-the-art tools. In [141] the datastructure PTrie is introduced for storing sets of binary strings of arbitrary length. The data-structure has improved signicantly the performance of our model-checkers, and won the best paper award at ICTAC2017. In [177], a review of 20 years of

significant industrial application of UPPAAL was presented. The work highlighted a number of selected case studies and discussed the pitfalls in achieving industrial impact as well as tool sustainability.

Model-based Testing

Model-based testing of real-time systems have been extended with mutation-based test case generation techniques that are able to generate test cases from model mutations (fault-model) that could potentially detect corresponding implementation faults. The developed algorithms use model-checking and refinement checking to distinguish equivalent mutants, and to generate adaptive test-strategies. This has been implemented, evaluated and published in the tools Ecdar [166], ECDAR 2.0 [107], and directly in UPPAAL [181, 206].

In addition, we show how testing of real-time systems can be done compositionally by combining conformance testing, refinement checking by verification and simulation into a coherent theoretical framework with tool support [164].

Scheduling, planning and optimization

To scale up analysis of timed automata with larger state-spaces, we developed distributed algorithms for time optimal reachability analysis and implemented those as extensions to UPPAAL. One approach uses swarm verification where a number of independent instances follows different search strategies to find a solution and to improve its cost. As an extension to this approach, we developed fully distributed algorithms that compute better solutions faster by exploiting a shared state-space and different parallel search orders [260, 259].

We have extended the reachability algorithm for UPPAAL to support multi-objective optimizations for scheduling and planning tasks. The new algorithm incrementally computes the Pareto front for simply priced timed automata (i.e. an extension that allows discrete prices on transitions for multiple cost variables but no price rates). The search is pruned based on a dominance relation between the cost vector. The applicability and efficiency of the algorithm is demonstrated on task graph scheduling problems, and a real-life case-studies: Task scheduling of GomX-3 nano-sattelites where the objective is to maximize the number of all executed jobs, data collection and downlink jobs, and remaining battery level [261].

We have explored how planning can be used to automatically determine the right level of abstraction in instruction giving systems [156]. When generating technical instructions, it is often convenient to describe complex objects in the world at

different levels of abstraction. A novice user might need an object explained piece by piece, while for an expert, talking about the complex object (e.g. a wall or railing) directly may be more succinct and efficient. We introduce the use of hierarchical planning to this end, a method from AI planning which can capture the structure of complex objects neatly. We apply this to a Minecraft instruction-giving system and show that our approach is flexible and can generate different explanations at different levels of abstraction.

Traditionally, planning tools are designed as applications that users need to install, taking care of the specificities of their operating system, packages, etc. This often places a significant burden on users, especially due to the academic nature of the software. To remedy this situation, we proposed for planning to jump onto the growing trend of light-weight software use; demonstrating the possibility to run planning tools directly in the browser [256].

Highlights

Awards and Special Mentions. Among the publications above, three received best paper award: at TACAS 2019 [80], at ICTAC 2017 [141] and at Petri Nets'17 [70]. Among the many invited talks we highlight those at TACAS 2017 (unified speaker) and CONCUR 2019 (key note speaker).

We mention that TAPAAL won a number of gold medals at the Model Checking Contest in 2018 (2), 2019 (4) and 2020 (2). We also highlight the very large number of downloads of UPPAAL (60.000 over the 5 year period). During the spin-off company DEIS-Tools (now VeriAAL) was created on the basis of the tool-effort of the group.

External Collaborators (selected) Patricia Bouyer (ENS Chachan), Doron Peled (Bar Ilan University), Sean Sedwards (University of Waterloo), Axel Legay (INRIA, Rennes), Jan Kretinsky (Technical University Munich), Bernhard Aichernig, Martin Tappler (Graz Techical University), Holger Hermanns (Saarland University).

4.5.3 Networks

Software-Defined Networks (SDNs) and Virtual Networks

Software defined networking is a new paradigm in controlling network protocols that, contrary to the traditional distributed network operation, assumes the existence of a centralized controller that orchestrates the network tasks in a glob-

ally optimal manner, while still dealing with a highly distributed underlying system of independent routes. We studied the problem of update synthesis, when changes in routing tables of the individual routers must be scheduled so that any intermediate network configuration transiently satisfies a number of properties including loop-freedom [86, 37, 263, 262] and waypointing [13, 184, 14, 12]. Another important topic was the optimization of virtual network functions, middleboxes and routing in virtual networks [186, 249, 87, 185] including the clustering techniques [258, 93, 19, 46, 238] for virtual networks. We developed novel analysis techniques and dealt with scalability issues in software defined networking [153, 120, 42, 62, 45, 149, 38] and studied congestion-free, adaptive and randomized routing [11, 264, 202, 232, 43, 61, 118, 41, 237]. A virtual network design centered on machine learning techniques was also investigated [44, 238].

Resilient and Secure Routing in Network Construction

In nowadays networks with high dependability requirements, it is essential that networks operate correctly and without security leaks even in case of link and node failures. To account for these problems, we studied fast (local fail-over) rerouting mechanisms in networks [51, 88, 245, 223] and data plane connectivity and node placement in dependable networks [52, 224, 225]. We also designed a routing algorithm [8] in order to reach all devices in an area of interest; the protocol is both resilient against churning of the devices and limits the information that flows in the specified part of the network for security reasons. Finally, we implemented a web-service API for our model checking tool UPPAAL running in a cluster environment, in order to analyze M2M communication in IoT applications [89].

Demand-Aware Optical Networks

In emerging data center architectures, the current trend is towards the construction of reconfigurable optical networks, and employed solutions are usually hybrid in the sense that a part of the network is reconfigurable while other parts are static. We studied algorithms for demand-aware network reconfiguration and self-adjustment with a focus on optimality for different network topologies and routing policies [85, 22, 20, 241, 21, 222].

Network Security and Network Verification

Network verification combines the fields of networking and formal methods with the aim of providing fully automated methods to reason about network properties, including security aspects. This research area covers our studies on security implications of introducing virtual switches [252, 117], analysis of various security attacks [147, 68, 122, 251] and other general security considerations [1, 250, 246]. We worked on dynamic resource scheduling and credit-based shaping for improving network performance with provably optimal guarantees [94, 116] also with the focus on timing aspects [157, 99] and stochastic aspects [148]. We pioneered methods for the analysis of MPLS networks with unbounded number of packet headers [134, 242] and implemented our approaches in the AalWines MPLS analysis tool [143] (available online at `http://demo.aalwines.cs.aau.dk/`). Aalborg University decided to file a patent application [240] due to a potential industrial application of the method. Our work in [216] studied the use of formal contracts to specify communication capabilities and requirements, forming a mathematical framework allowing for both security and safety in communication. We designed and implemented a set of tools [221, 9] to support the creation of secure webservices and their integration in advanced local clouds.

Wireless and Decentralized Networks

Finally, wireless and highly decentralized networks bring further challenges with respect to resource allocation and coordination. We studied here efficient sharing algorithms for the wireless medium under interference where we introduced a distributed MAC protocol [211] and we suggested a novel WiFi architecture based on software defined wireless networks [244]. We evaluated the problems of robot evacuation and convergence using wireless communication [218, 217] and considered further theoretical foundations [119]. In the context of a study on vehicular networking, we developed a tool [63] to configure a mainstream traffic simulator (SUMO) by generating realistic population, industry and school densities.

Awards and Special Mentions. Our novel approach on the verification of MPLS networks turned into a patent application [240] and we received two best paper awards IFIP NETWORKING 2018 [231], and ACM SOSR 2018 [252].

External Collaborators. Chen Avin (Ben Gurion University of the Negev), Anja Feldmann (Max-Planck-Institut für Informatik Saarland), Yvonne-Anne Pignolet

(ETH Zürich), Matthias Rost (TU Berlin), Jie Wu (Temple University, Philadelphia), Riko Jacob (ITU Copenhagen), Wolfgang Kellerer (TU Munich), Marco Canini (King Abdullah University), Marcin Bienkowski (Univerisity of Wroclaw).

4.5.4 Security

Security in Software Defined Networks.

Work in these areas, using mainly tools and techniques from *language-based security*, has lead to an early analysis of the attack surface of P4-based SDN infrastructure [1] and the development of a novel tool, a compiler fuzzer, called *P4Fuzz*[1], for finding bugs and vulnerabilities in the P4 compiler, a critical part of the SDN infrastructure. Basic fuzzing is essentially stress testing an application, typically with either random or "problematic" input, in order to discover bugs and potential security vulnerabilities. Compiler fuzzing is a more involved version of basic fuzzing: a compiler fuzzer works by generating random (but valid) programs exercising the compiler and the generated code. The compiler fuzzer is ongoing work [246]. In other work, various aspects of security for SDNs have been investigated yielding, among other things a number of attack(er) taxonomies, vulnerabilities, and possible attacks against network infrastructure. [117, 147, 157, 68, 122, 251, 250, 183, 235, 236].

Security in Socio-Technical System.

The work in this area started in the previous reporting period in connection with the TREsPASS project, and has been further explored in this reporting period, both in the context of the TREsPASS project and also as separate strands of research. The main perspective has been on applying formal methods to modelling and analysis of systems (mainly complex organisations) with a significant human component, in order to uncover novel attack vectors, potentially exploiting human behaviour and other social factors such as insider knowledge and social engineering. The modelling and analysis provides a foundation for better risk analysis and risk management. This work has lead to the development of several novel modelling formalisms and concomitant analyses and tools [67, 226]. The TREsPASS project was concluded during this reporting period.

[1] `https://github.com/andrei8055/p4-compiler-fuzzer`

String Solving.

Many attacks involve some form of "injecting" code into a program, and that injection is only possible due to unsanitised user inputs (strings) used in insecure ways. String Solving is thus an important research area as it (combined with symbolic execution) allows analysing which string patterns are used in different parts a program. That is, we can determine if the programs allows an unescaped string to be used in a SQL-query.

In the current reporting period, we have developed the tool Woorpje [74, 84] that solves Word Equations with length constraints (the decidability of which was uncertain for long) and showed how SMT-solvers (CVC4 and Z3) could benefit from solving techniques as a pre-processing step. In parallel with developing Woorpje we also created the (to our knowledge) first framework [158] for automating the comparison of String Solvers.

Attack-Defense Trees.

In part motivated by the work socio-technical systems, work on *attack-defense trees* and related formalism has resulted in novel techniques for modelling and analysing potential attacks and attack-scenarios using timed-automata and model checking (with UPPAAL), most notably for a quantitative analysis using stochastic timed automata. This work also resulted in the development of a domain specific language for interactive development, exploration, and analysis of attack-defense trees [113, 112, 98].

Language-Based Security.

In addition to the work mentioned under SDN security, work in this area has progressed in developing models and static analyses of programs that may experience *bitflips*, e.g., through directed software attacks such as ROWHAMMER or through physical attacks, e.g., voltage glitching of IoT devices. Also in this area a tool has been developed for verification of LLVM programs, an essential component in many state-of-the-art languages and compilers, e.g., Rust and Clang [180].

Furthermore a model for analysing and verifying *adaptive* security policies for mobile contexts was proposed [204].

General Security.

In addition to the already mentioned topics, work has progressed in a number of other areas of security, e.g., blockchain security and formal methods for security.

Blockchain technology is probably best known for its use in cryptocurrencies, e.g., Bitcoin, but has potential for many other application areas, e.g., smart contracts. Finally, a *randomness beacon* has been designed, analysed, and developed, for generating random numbers in a scalable and secure way inspired by blockchain technology [76, 75].

Results have also been obtained in areas that are more closely related to classic network security, such as a data-driven approach to bot detection [147]; a method for safe and secure communication in autonomous vehicles [216]; and a web-crawler to collect data on web-sites for security analysis [215].

Other notable results include a method for automatically verifying stateful protocols within an LCF-style interactive theorem prover implemented for Isabelle/HOL. This includes a simple language for specifying transaction-based protocols; an approach to making phishing detection more robust by taking adversarial behaviour into account [213]; and a formal model based on the *ambient calculus* has been proposed for analysis of overlay networks and their security [115].

Finally, a virtualised networking platform for security education [214] was designed and implemented.

Highlights

In 2019, Aalborg University hosted the the 24th Nordic Conference on Secure IT Systems (NordSec 2019) [18]; the paper [212] was awarded the EuroSys *Test of Time Award* (for the most influential EuroSys 2008 paper) at EuroSys 2018.

4.5.5 Machine Learning

Probabilistic Graphical Models

In the field of probabilistic graphical models [133], our main focus has been on scalable inference and learning techniques for latent variable models. A significant part of these investigations have been conducted as part of the AMIDST project (see Section 4.6) and implemented in the AMIDST toolbox [197]. The AMIDST toolbox is a standalone Java-based library for scalable probabilistic machine learning based on probabilistic graphical models with special focus on massive data streams.

Inference in probabilistic graphical models Inference (or belief updating) is fundamental to the applicability of probabilistic graphical models [234]. In this area, we have focused on scalable approximate inference techniques exploiting distributed computing resources managed by processing frameworks like Apache Spark and Apache Flink. Based on the family of conjugate exponential models, we have investigated approximate inference based on variational techniques. As part of the investigation, we have proposed a method for performing variational inference based on a projected natural gradient ascent approach [195, 196]. The method establishes a theoretical foundation for the parallelization of variational methods, realized by a distributed variational message passing framework.

We have also considered techniques for scaling up sampling-based inference methods. Specifically, we have developed a scalable importance sampling-based algorithm for estimating posterior distributions in conditional linear Gaussian Bayesian networks (a subfamily of the conjugate exponential models) [228]. For these types of models, we have also developed methods for scalable maximum a posteriori (MAP) inference, covering both static [229] as well as dynamic domains [227].

Finally, in the context of decision analysis, we have analyzed and developed methods for finding anytime solutions for unconstrained influence diagrams [187]. These types of models provide an explicit representation of (unordered) decisions and utilities, hence a solution to the problem will not only consists of an optimal decision strategy but also a conditional specification of the order in which the decisions should be made.

Learning probabilistic graphical models In the area of learning probabilistic models, we have investigated methods for scaling up Bayesian network learning. This includes the development of a filter-based feature selection method [233] that distributes the required calculations between the available computing units using balanced incomplete block designs, a strategy first developed within the area of statistical design of experiments. The use of balanced incomplete block designs have also been investigated for directly scaling up Bayesian network structure learning [188]. In particular, we have proposed a variant of the PC algorithm, where the computationally intensive conditional independence tests are parallellized using block designs providing significant speedups. Techniques for scalable learning have also been investigated for multi-dimensional classification within this family of models[16].

Data streams are pervasive in many domains, but learning from such data sources introduces new challenges such as continuous model updating and the ability to adapt to changes and drifts in the underlying data generating process [58, 198].

By approaching this learning problem from a Bayesian perspective, we have made use of non-conjugate priors over the model parameters to explicit model temporal changes in the data [199, 200]. To ensure efficient inference and learning, we have derived a novel variational inference technique that accommodates the non-conjugate model components while maintaining the computational efficiency offered by variational methods over conjugate models.

Relational Learning

In the field of relational learning we are investigating foundational principles and algorithmic solutions for modeling, learning, and inference with multi-relational graph or network data.

Foundational Issues The *statistical relational learning (SRL)* approach to relational data consists of defining probabilistic generative models for relational structures, i.e., modeling the data by full probability distributions over all possible relational structures for a given finite domain of entities. These generative models provide general templates that can be instantiated for any finite domain. A potential problem of most commonly used statistical relational modeling frameworks is *domainsize dependence*: the result of a query related to a fixed set of entities does not only depend on the known properties of these entities, but on the size of the domain they are embedded in. We have conducted an in-depth study of *projective* models in which query results are independent of the domainsize. In [128] we have provided rigorous definitions of projectivity, clarified the implication for probabilistic inference and statistical learning, and identified for several prototypical SRL frameworks syntactic fragments that are guaranteed to specify projective models. [129] introduces a new type of latent variable probabilistic graphical model that is shown to capture precisely the class of projective relational models.

Relational learning frameworks In the area of algorithmic frameworks and implementations, we have taken a long-standing line of research on type extension trees in several new directions: in [127] we use the counts-of-counts semantics of type extension trees to define a flexible class of metrics for entities in relational structures, which in experiments is shown to outperform similarity-based predictions using traditional kernel-based similarity measures. Combining the recursive type-extension-tree modeling approach with a neural network architecture, we have introduced a framework of multi-multi instance learning [253, 254], that in experiments outperforms other neural architectures for graph data.

Machine Learning and Formal Methods

Connections between machine learning and formal methods have been explored in several directions. In [125] a reinforcement learning approach has been developed to construct controllers for cyber-physical systems with continuous state spaces, as modeled in the UPPAAL-STRATEGO platform. The reinforcement learning methods have proven to outperform an earlier simulation based optimization strategy implemented in UPPAAL-Stratego. The work in [125] gave rise to a number of theoretical questions related to the convergence of approximations of continuous state spaces by nested sequences of finite state space partitions. These questions were addressed in-depth in [124]. In this work partition-based approximations are formalized as imprecise Markov decision processes, for which then value-iteration algorithms are developed to compute upper and lower bounds of a cost function. An analysis of the behavior of these bounds then leads to approximation guarantees for the original continuous state space.

In a second line of work a framework has been developed for learning Bayesian networks as a model checking and online analysis tool for stochastic timed automata [126]. The proposed approach consists of learning a Bayesian network model for a fixed set of temporal properties from sampled system traces. Conditional probability computations in the Bayesian network then provide a flexible tool for online updating of probability estimates for future system behaviors. Experiments showed that due to the Bayesian network learning's ability to recognize conditional independencies, given a fixed amount of sampled traces, this approach often leads to more accurate probability estimates than statistical model checking.

External Collaborators. Helge Langseth (NTNU, Norway), Andrés Masegosa and Antonio Salmerón (University of Almería, Spain), Jose A. Gamez (University of Castilla-La Mancha, Spain).

4.5.6 Methodologies and Applications

This section collects research that is focused on specific methodologies or more focused on the application domain.

Model-Based Methodology

The Uppaal Tool Suite. The group has developed the UPPAAL tool suite for model-based development of CPS for more than 20 years. It has expanded into a

full-grown tool suite that covers a number of different areas, including verification, testing, scheduling and controller synthesis. This development has from the outset been driven by a series of research projects on one hand and industrial case-studies and applications on the other. These have not only led to a strong tool suite, but also valuable insights into what it takes to sustain such tool development, industrial collaboration, and the challenges of industrial uptake [177, 165].

Modal transitions systems and contract based design Compositional system design is essential in the development of complex systems. In this line of research, the highly cited book [39], co-authored by internationally renowned researchers. The framework is inspired by modal transition systems, one of the core competences of DEIS, in which common transitions are extended by so-called must and may transitions. The book illustrates the use of contracts on two real world examples: requirement engineering for a parking garage management and the development of contracts for timing and scheduling in the context of the AUTOSAR methodology in use in the automotive sector. In a similar line of work, [178] proposed strategies enabling designers to synthesize or refine a set of contracts so that their composition satisfies a given contract, while [179] introduced heterogeneous composition, refinement, dominance and compatibility between contracts, altogether enabling a formalized and rigorous design process for heterogeneous systems. Most recently, [40] presents a first Interface Theory – Modal Mixed Interfaces – for systems exhibiting both non-determinism and randomness in their behaviour.

Analysis of Hierarchical Scheduling Systems Some embedded systems can be structured as a hierarchy of components. This is an approach that is mainly used in the automotive and aviation industries for now. The main motivations are stability of safety critical systems, but one motivation can also be the ease with which several legacy systems can be integrated into one system. The hierarchical nature of the systems can also be used in the analysis of the systems.

As a specific example of a hierarchical systems we have developed a modelling framework (Uppaal stopwatch automata) and compositional analysis method for Distributed Integrated Modular Avionics Systems based on the ARINC-653 standards. The model-based approach allows partition deployments on multiple cores as well as distributed processors interconnected by AFDX networks [110, 109, 108]. Several papers have been published with the aim of describing a more general software engineering approach to the analysis of hierarchical scheduling frameworks [66, 53, 154], thus aiming to make the research more accessible for the schedu-

lability community. In [54] a new *Generic Formal Framework for Compositional Analysis of Hierarchical Scheduling Systems* is proposed.

A single paper suggesting a new scheduling policy for mixed-criticality systems, in a non-hierarchical setting, has also been published [55].

Uppaal Stratego has been applied to the problem of synthesising strategies for the robots playing soccer (RoboCup)[121] and an architecture based on the precomputation of strategies was able to cope with the problem of long computational time caused by state space explosion.

Testing Using formal models for testing as well as for verification is also the focus in five papers. A framework for requirement-driven test generation is proposed in [5]. Application of formal methods to aid in the certification of a wireless fire alarm systems as a case study [15], with an emphasis on testing. In [164] compositional testing of real-time systems is explored by combining the capabilities of several Uppaal tool versions, namely Uppaal Tron, Uppaal Ecdar and Uppaal SMC. Using a mutation analysis approach to find the location of errors in a system under test: Effortless Fault Localisation: Conformance Testing of Real-Time Systems in Ecdar [107]. Mutation analysis and conformance testing has also been implemented in the Uppaal toolsuite [206].

Methodology A single invited talk focuses more directly on suggesting a development approach for Dependable and Optimal Cyber-Physical Systems [172].

Safety-critical Java for embedded systems We have collaborated with the DPW group on defining and using Java for development of safety critical Java [243, 230]. In the first part of the reporting period, research on Safety Critical Java was consolidated by a journal paper [182] on the hard real-time Java Virtual Machine (HVM) and a journal paper on the timing and schedulability tool TetaSARTS, which translates Safety Critical Java programs into timed automata models suitable for analysis by various versions of the UPPAALmodel checker. The work on Safety Critical Java had major impact on the JSR 302 standard, which, to a large degree, is based on the initial proposal made by members of DPW and DEIS. The final version of JSR 302 has been approved in a ballot ending 26.5.2020. Safety Critical Java has had major industrial impact and is now powering more than 20.000.000 devices according to Aicas GmbH (https://www.aicas.com/). A huge effort to persuade Danish industry to adopt the result of this research turned out fruitless, and a further study, done in collaboration with Epinion, revealed that very few

Danish companies, at the time of the study, had any interest, nor perceived any need for verifiable hard real-time embedded systems development, especially not in Java. Furthermore, over the years the small research community around real-time Java diminished and in 2016 the last of the JTRES workshop series was held. It was therefore decided not to continue this line of research.

Popular Parallel Programming At end of the last reporting period research work had begun on parallelising spreadsheet computations in a joint effort with members of the DPW group and a team of researchers at the IT University in Copenhagen in the DFF sponsored project Popular Parallel Programming (P3) project. This research builds in part on ideas from the work on Safety Critical Java by translating spreadsheet programs into timed automata models analyzable for parallel schedules by the UPPAALmodel checker. Unfortunately, one of the major findings of the project was that this approach was not scalable to large spreadsheets. However, results of the project showed that dynamic scheduling, based on a graph of dependencies among cells in a spreadsheet, gave almost linear speed up on multi-core processors, especially for very large spreadsheets (millions of cells and up to 100.000 formulae). In the latter part of the project, work on a cost calculator, i.e. a system that can analyze the computational steps needed for evaluation of formulae, as well as partial and full recalculations of spreadsheets, was completed. This work is based on a formalization of the spreadsheet programming model using structural operational semantics. So far, this direction has resulted in 3 publications, one of which received the 2019 Journal of Computer Languages Best Paper Award. Although the DFF sponsored project has finished, work has continued on constructing a general framework for abstract interpretation of spreadsheet programs, based on the ideas from the cost calculator and the formal semantics. Funding for scaling spreadsheet calculations to big data analysis has been sought, but so far unsuccessfully, but funding for investigating how the spreadsheet model may help in bringing computational thinking to non-computer savvy users has been obtained from IT-Vest. This work will formally start in 2021.

Application Areas

We categorize our applications according to the societal challenges of energy, transport, health, production, and others.

Intelligent Transport A notable application to the intelligent transport area is optimization of the traffic flow through of signalized intersections by combining

better control algorithms as synthesized by UPPAAL Stratego with modern radar technology for observing approaching vehicles [82]. In [83] the new controller is tested against the real controller deployed in the intersections using the micro simulation program VISSIM. The simulation shows that in comparison with the existing controller, this controller provides a reduction of between 30% and 50% in average delays, queues and number of stops. The fuel consumption and total travel time of cars in the coordinated section are reduced by about 20% in the simulation study. All these reductions are achieved without making the situation worse for the side roads. A company (ATS: Advanced Traffic Systems) has been spun of to commercialize the technology. Also real-life test show good results (13-16% less traffic stops, and 2.5 seconds less travel time). Although less than the simulations, the gain is still significant. The reduction is mainly attributed to imperfect sensors. Hence, there is a still a huge improvement potential if better sensing technology can be developed. Inspired by this, Uppaal stratego has also been applied optimized and safe control in the rail-domain [150]. Finally, we have worked on coordination of autonomous robots: In [239], we suggest a timed game formulation of the collision-free planning problem for UAV fleets accommodating the uncertainty of flight duration and providing the synthesis of a time optimal discrete event controller. In [217] we develop distributed algorithms for coordinating robots with limited cognitive abilities.

In the context of road networks, we have developed a new machine learning techniques for machine learning on road network graphs [145, 146, 144]. We developed Relational Fusion Networks (RFN), which is a novel type of Graph Convolutional Networks (GCN) designed specifically road networks. We show that RFNs outperform conventional GCNs and is able to improve the qualitative information such as road type and speed limits. A delay-robust touristic plan recommendation system for real-World public transportation information is developed in [23].

Green Energy The increased penetration of renewable resources (RES), such as wind and solar energies, whose production is volatile, and difficult to accurately predict, challenge our energy distribution and consumption systems. Further, to exploit the available energy in the best possible way, the energy systems are becoming increasingly integrated (e.g., district heating, cooling,grid, industrial excess energy, energy storage). This calls for the development of advanced optimised control of multiple RES according to price or availability signals from dynamic energy markets [130, 103], and formal verification of such energy management systems [247]. It also calls for exploiting demand side flexibility where consumption is shifted in time or amount according to available energy supply. To support this, the po-

tential flexibility of devices in different sectors have been characterized along with methods for aggregating this flexibility [92, 72, 102, 100, 101, 6].

The UPPAAL tool-suite has been applied for modelling and analysing energy systems [247]. Particular effort has been on smart houses [106]. Notably, Uppaal-stratego has been applied to synthesizing controllers for intelligent floor-heating control. To combat the complexity of the control problem, two main techniques are used: compositionality where a controllers are synthesized for a subset of rooms under certain assumptions of the control of the rest, and online synthesis where the control strategies are recomputed periodically for a certain time horizon, possibly taking the weather forecast into account. The resulting control system have been demonstrated in laboratory and real-life environment [170, 3], showing a significant improvement over a conventional bang-bang controller on both reduced energy consumption and discomfort. A controller for photo-voltaic comfort cooling has been developed and demonstrated [4]. As part of this, a new concept of probabilistic flexoffers has been developed [2]. Also, new frameworks for feature interactions [105] and smart home scenarios [104] have been developed.

Energy Aware CPS: Cyber-physical devices are often energy constrained by available battery and energy harvesting capabilities, and hence using the battery carefully to maximize utility or lifetime becomes essential. In a case-study provided by the company GomSpace, we study how to optimize the utility of low-earth-orbit (LEO) nano-satellites under energy harvesting constraints dictated by the eclipses of its orbit. The objective is to complete as many tasks while guaranteeing energy for mandatory tasks for attitude control. This achieved by combining realistic kinetic battery models, stochastic task models, and cost optimal reachability analysis of priced time automata to construct day-ahead schedules for the satellite's tasks [205, 261]. In the DONUT project we have developed an stochastic timed automata model for a Sigfox based wireless sensor node for measureing the water depth in manholes in urban water sewage systems [203]. The model is analysed using Uppaal SMC to estimate the resulting battery lifetime given different data-sampling, transmission, adaptivity, and data compression techniques. Preliminary results shows a potential for increasing lifetime from 1 year to 3 years. Scheduling under energy constraints is also examined in [123, 257].

Health: We have contributed to a number of projects in health care.

In a project on prediction of home care deliverables with a high risk of deviations, we analyze data from the home care sector in order to detect deviations between

granted and actual deliveries [189].

We worked on prediction patient flow in the emergency department at Aalborg University hospital in order to improve the quality of service by reducing waiting time and provide a better foundation for planning staff schedules [190].

A dual-chamber implantable pacemaker monitoring a human heart has been formally modelled using hybrid automata and verified using a combination of Uppaal and SpaceEx [10].

COVID-19 Within the project BEOCOVID, sponsored by Poul Due Jensen Foundation, significant work has been made to apply UPPAAL to the modelling, predicting and control of the COVID-19 pandemic. In [139] we focus on the SEIHR epidemiological models and show how the risk of viral exposure, the impact of super-spreader events, tracing-apps as well as other scenarios can be modelled, estimated and controlled[139]. Most recently, an agent-based model of Northern Jutland has been made (comprising 1 million interacting automata, 2 per citizen of the region) allowing to evaluate effect of the lock-down of the various municipalities during the autumn of 2019. To allow for simulation of such an extremely large model, several special purpose improvements of the SMC engine had to be made. The work has raised quite some awareness at SSI (Statens Serum Institute).

Clean Water: We have contributed to the clean water agenda through the DONUT project. The aim of this project is to provide a holistic view of the urban water system by making continuous water level measurements using a low-cost sensor network. We have worked on ensuring the reliability of the network by developing methods for detecting anomalous events/observations, e.g., due to faulty sensors or anomalous water conditions [255]. In the recently started CLAIRE project, we pursue the same agenda with the goal of synthesizing safe and sustainable control strategies for water resource management based on sensor measurements.

Production: The works in [59, 60] define an open-source tool-suite (based on Arrowhead and other tools) for remote condition-based maintenance, which feature a low entrance barrier for SMEs. Based on this, [7] describes a set of pilots from the MANTIS project regarding condition-based maintenance, and it leverages on them to categorize how different sensors are used in real-world advanced maintenance. The chapter [131] discusses how ubiquitous information is the stepping stone to allow for modern maintenance operations, and in particular for condition-based

and proactive maintenance.

Other: We have applied our techniques and tools to **static analysis of programs**. Very large spreadsheet programs can take hours to compute; In [47] we translate spreadsheet programs into Timed Automata Models, which are analyzed by the UPPAAL model checking tools, with the purpose of finding schedules for parallel execution of the spreadsheet calculations. On example spreadsheets, linear speed up was possible with up to 4 cores, but also that the scalability of UPPAAL becomes a limitation.

We develop a semantic framework formal modelling of a large variety of single event upsets for ARM assembly programs [114], and use this to statically analyse and enhance their fault tolerance. The semantics is further used to extract a timed automata model that is subjected to statistical model-checking to quantify the program behavior in presence faults and to show that the enhanced analysis method significantly decrease the probability of such faults going undetected. In [180], we develop a tool (Lodin) for checking (potentially concurrent) programs given as LLVM intermediate code using explicit-state model-checking supplemented with simulation-based model-checking to combat the state explosion problem.

Finally, [246] improves the effectiveness of a fuzzing based test generator by using static code analysis, and demonstrates that the resulting technique and uncovers several new vulnerabilities in network security tools.

In the area of **power electronics**, the technique of Statistical Model Checking in the form of UPPAAL SMC has been applied to the validation of Model Predictive Control for DC-AC power converters. This was a close collaboration with the Department of Energy Technology from Aalborg University. This collaboration led to four publications [208, 209, 210, 207] tackling two different types of power converters. The work focused both on estimating good parameters for the configuration of the systems and on validating the stability of Model Predictive Control power converters, something which traditional methods within control theory has not been able to handle. The two final publications were in a high level journal [208] and the wider reaching IEEE Industrial Electronics Magazine [209].

4.5.7 Software Tools

The DEIS group has a renowned reputation on application of formal theoretical models for tools development. This fame has been achieved in the years by investing considerable effort in the development and maintenance of our software tools.

Next we report the main software tools which have been developed or received maintenance during the evaluation period. Details about the work done for each tool within the period 2016-2020 can be found in the Appendix.

UPPAAL is an integrated tool environment for modeling, simulation, validation and verification of real-time systems modeled as networks of timed automata.

TAPAAL is a tool for modelling, simulation and verification of Petri Nets and Timed-Arc Petri nets.

CAAL (Concurrency Workbench, Aalborg Edition) is a open source web-based tool for modelling, visualization and verification of concurrent processes expressed in the CCS language (Calculus of Communicating Systems) and its timed extension TCCS (Timed CCS).

AalWiNes (AALborg WIen NEtwork verification Suite) is a open source network verification suite performs fast (polynomial time) verification of forwarding rules in MPLS routing tables in the presence of link failures. AalWiNes is the result of the collaboration between Aalborg University and Vienna University.

AMIDST is an open source Java software for scalable probabilistic machine learning with a special focus on (massive) streaming data. The toolbox supports a flexible modelling language based on probabilistic graphical models with latent variables. AMIDST provides parallel and distributed implementations of scalable algorithms for doing probabilistic inference and Bayesian parameter learning in the specified models.

Arrowhead The Arrowhead Framework is the result of a large European effort to apply Service Oriented Architectures (SOA) to industrial applications, and DEIS was part of the core development group since the beginning. The framework supports the development of distributed applications with capabilities such as service registration and lookup, service orchestration, security tokens, Quality of Service setup, inter-cloud communication and service search, online generation of interfaces for service interoperation, and service-oriented management of security certificates.

Homeport is an open-source middleware that enables home automation applications to be developed on top of an Internet service layer. Homeport translates from low level wireless or wired end-device protocols to services that are accessed by conventional service calls on the Internet.

OpenAPI-Generator Code generator for REST APIs. It is one of the leading code generators adopted by the industry and it covers the largest numbers

of programming languages. DEIS members are part of its founding members group, and are maintainers of the code generators for a few programming languages.

ERODE is a tool for the solution and reduction of systems of ordinary differential equations developed in collaboration with the University of Oxford and IMIT Lucca.

Highlights

In the period 2016–2020, UPPAAL has been downloaded more than 1000 times per month, totaling over 60,000 downloads. TAPAAL participated yearly in the Model Checking Contest winning a combined of 8 gold medals, 6 silver medals, and 1 bronze medal.

4.6 Projects and Funding

4.6.1 Projects

In order to describe the DEIS project activities, we divide them into the following categories:

- *Industrial Dissemination* projects serve as settings for contributing to industrial innovation, to competitiveness, and to stimulate further industrial collaboration. The projects often include demonstrator work packages and generate a number of so-called *demonstrators* which have been developed and presented at industrial exhibitions. Also, dissemination has been pursued through focused industrial interest groups and through industrial courses that have been taught both as parts of accredited life-long learning activities.
- *Strategic Research* projects are based on specific societal challenges and often include one or more industrial cases which serve as milestones for the scientific work. In the period, a number of application domains have been part of the strategic research projects, e.g. Energy, Transport, Production and Epidemics. Also, in order to strengthen the impact on society, the projects often include domain experts and international partners.
- *Basic Research* projects take their offset from specific scientific challenges.

Below, we briefly describe the main projects that have been run in the period.

- The ERC Advanced Grant LASSO (Learning, Analysis, SynthesiS and Optimization of Cyber-Physical Systems) granted to Kim G. Larsen aims at developing a new generation of scalable tools for cyber-physical systems through combining advanced model-checking techniques with machine learning. Validation is to be made from case in a variety of application domains. The project is running for the period 2015-2021.

- DiCyPS: The work of DiCyPS – a national Danish strategic research Centre for Data-Intensive Cyber-Physical Systems – focuses on utilizing software and data from the IT management of complex physical systems for the development of smarter and more user-friendly solutions for society and individuals.

  The research centre is funded by Innovation Fund Denmark. It is based at Aalborg University and constitutes a unique collaboration between internationally leading research groups, Danish and international companies, researchers and public authorities on the creation of a solid basis for IT structures in the smart society of the future.

  The research centre involves all research groups at the CS-department as well as AAU-researchers from the energy and transport domains. The center is coordinated by DEIS.

- InfinIT: This project is a large industrial dissemination innovation network involving all Danish CS-Departments. DEIS is directing the innovation network and its activities and is also in charge of one of the two project offices. Industrial dissemination is carried out through a series of national events (e.g., InfinIT Summits) including keynote talks, panel sessions, focused parallel tracks and technology demonstrations. Also, DEIS is leading several of the special interest groups of InfinIT and has also been in charge of a number InfinIT mini projects, each involving at least two companies.

- AMIDST: In the H2020 AMIDST strategic research project (Analysis of massive Data streams), we explore and develop scalable algorithms for learning probabilistic graphical models within a data streaming context. In particular, we have considered scalable and distrbuted Bayesian learning algorithms focusing on data streams exhibiting concept drift.

- FED (Flexible Energy Denmark): The ambition of this national industrial strategic research project is to enable the development of digital solutions that are capable of adjusting the power consumption to fit the power production – among other things by use of Machine Learning and different tools for handling Big Data.

FED involves 24 participants from the four technical universities, utility companies, businesses and public organizations. The solutions will be demonstrated at seven Living Labs.

- MANTIS: The overall concept of this H2020/ECSEL strategic research project is to provide a proactive maintenance service platform architecture based on Cyber Physical Systems that allows to estimate future performance, to predict and prevent imminent failures and to schedule proactive maintenance.

  The project included data collection/ demonstration activities in Denmark that were provided by Vestas and the city of Aalborg involving the technology SME's Xtel and Neogrid.

- DONUT: This national Danish strategic research project results will accelerate a paradigm shift within the water sector towards data-driven decision-making, improving planning, monitoring and operation. The DONUT project will help move the water sector into the big data era as it will develop and mature the basis of the forthcoming IoT and ICT infrastructures for blue, green and climate-adapted cities.

- FEVER and DomOS: These two H2020 strategic research projects address the development of flexibility DSO services and building services respectively - including their IT infrastructure. DomOS includes two Danish demonstration activities and SME technology providers Neogrid and SUNTHERM. Both projects are examples of joint efforts across the research groups in the department.

- CLAIRE: is a Villum Syngergy Project. The overall aim of CLAIRE (ControLing wAter In an uRban Environment) is to combine machine learning and formal model checking techniques with water engineering in order to optimize water resources while also ensuring a safe and sustainable management of these resources.

- BEOCOVID: is a project funded by the Poul Due Jensen / Grundfos foundation. The overall aim is to use UPPAAL SMC and UPPAAL Stratego to model, predict and control the COVID-19 pandemic. The project has been conducted in close collaboration with the a the sister-project CTRL-Covid funded by the Novo Nordisk Foundation, involving mathematicians and control theorist from Aalborg University and Technical University of Denmark. Close collaboration with SSI (Statens Serum Institute) on agent-based modelling has taken place leading to "myndighedsbetjenings"-contracts being negotiated.

- IoT and CPS Associate Professorship. The position is enabled by a generous grant from the Poul Due Jensen Foundation, and includes additional funding of a PostDoc. The objective of the grant is to strengthen the department's activities in the broad range from theory to applications in the area of distributed, embedded, and intelligent systems.

- ASAP (Approximate Reasoning for Stochastic Markovian Systems): is a DFF project that proposes to investigate the logical foundations of the metric behavioural theory for stochastic systems.

- MULTI-CORE SAFETY: is a DFF project that proposes to utilize modern compositional analysis techniques to analyze the safety of multi-core systems. The results are obtained by applying more detailed models of the individual parts of the temporal behavior of the system and then analyzing these models compositionally so that the complexity does not exceed what can be analyzed.

- Quantitative Analysis and Synthesis of Network Protocols: is a DFF project with the goal of significantly advance the state-of-the-art of automatic network analysis and synthesis, with focus on accounting for the possibility of failures and supporting efficient what-if analysis, as well as for the quantitative consequences on quality-of-service.

4.6.2 Fundings

Table 4.1 shows the number and grant size of projects that are acquired in the evaluation period, with an addition from large and long-term projects (e.g. DiCyPS, LASSO and the DFF project ASAP) that were granted in the end of 2015 proportional to their spending in 2016-2020.

We believe that the DEIS group has continued to be active and successful in obtaining external funding in the present period. The funding ranges from a few very large (and prestigious)grants (e.g. LASSO, DiCyPS and DIREC) to several small and medium size grants for basic and strategic research, as well as industrial dissemination. The group has obtained grants for 39 projects, and the total external funding has been 72.3 M. DKK. For the previous period, the numbers were 44 projects and external funding 75.6 M. DKK (with the funding from DiCyPS, LASSO and ASAP only included proportional to the spending in the period). Figure 4.3 shows the external funding per year for the current period.

Compared to the previous 5 years evaluation period, there has been some 5%

Year	Projects	Total Grant	Significant Projects
2016	16	30.5 M. DKK	DiCyPS, LASSO, ENABLE-S3, Grundfosprisen 2016
2017	6	4.5 M. DKK	REACHI, Productive4.0, MULTI-CORE-SAFETY
2018	4	17.1 M. DKK	DONUT, InfinIT, GRUNDFOS professorship
2019	4	7.1 M. DKK	FED, QASNET
2020	9	13.1 M. DKK	DomOS, BEO-COVID, DIREC
Sum	39	72.3 M. DKK	

Table 4.1: Funded Projects by year

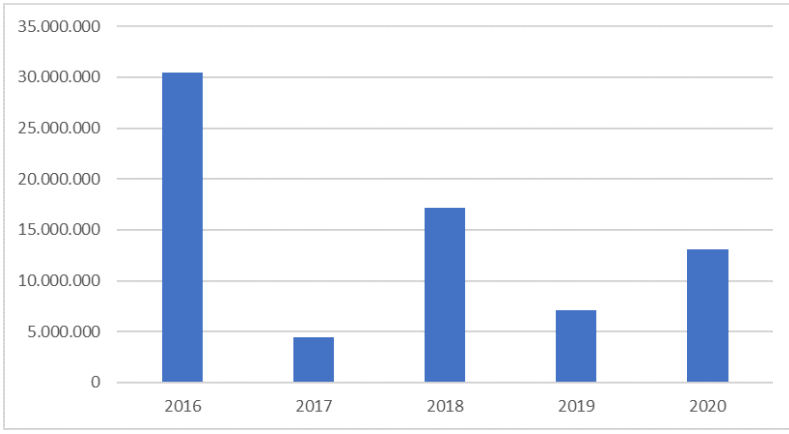

Figure 4.3: Yearly Aquired External Funding Profile 2016-2020

reduction in the external funding. However, no reduction in the overall activity level has been observed. This is largely due to the fact that three large projects have been active throughout the period —LASSO (basic research), DiCyPS (strategic research) and InfinIT (dissemination).

Table 4.2 show the distribution of projects granted in the period categorized by their degree of basic research, strategic research, and industrial dissemination. It can be seen that basic and strategic research projects are the major categories with respect to funding, and that the dissemination projects have significantly minor funding. We expect this profile to be kept for the coming period due to the

current Danish and European funding policy.

Project Type	Projects	Total Grant	Av. Grant/Project
Basic Research	12 (1 intl.)	41.6 M. DKK	3.5 M. DKK
Strategic Research	14 (8 intl.)	23.6 M. DKK	1.7 M. DKK
Dissemination	13 (6 intl.)	7.1 M. DKK	0.6 M. DKK
Total	39 (15 intl.)	72.3 M. DKK	1.9 M. DKK

Table 4.2: Projects by Type

For the current and near future funding situation we are especially happy about the recent acquired, substantial funding of DIREC, the new Danish National Research Center in Digital Technologies. We definitely expect that new collaborative research and development projects will emerge from this collaboration. Also, we believe that our generic research competences in CPS, optimization and AI will provide a competitive edge for several of the upcoming calls – both from private as well as public funding agencies – where a variety of societal challenges (e.g. reduction of CO2 emission) are to be addressed.

4.7 Own Evaluation

In our opinion the research activities carried out by the DEIS unit in the period 2016-2020 have been very satisfactory both in quantity and quality. The number of publications have increased since the last period totalling 348 publications and 1 patent, including 99 réfereed journal papers, 210 refereed conference/workshop papers, 8 book chapter, 1 book and 10 edited proceedings.

We consider these figures as an indication of the high quality and impact of our unit, which is further witnessed by the ERC Advanced Grant won by Professor Kim G. Larsen (2015), who also was awarded the Grundfos prize (2016), appointed international chair at INRIA/Rennes, appointed Distinguished Professor at North Easter University and Foreign expert of China. Also witnessing the excellence an impact is the large number of yearly downloads of the UPPAAL tool (14.000 pr year), the several gold (and silver medals) won by the TAPAAL tool in the Model Checking competition, the several best paper awards received by members of the unit, as well as several invited key note lectures.

In the following we detail how the unit has met the strategic goals of the period in terms of research directions, collaboration and projects as well research

management and staff hve been met.

Research Directions

- *Maintain strong international collaboration and leading role in Embedded and Cyber-Physical Systems*

  During the period, several of the major projects have been focusing on CPS. This include the ERC Advanced Grant LASSO (Learning, Analysis, Synthesis and Optimization for CPS) as well as the Innovation Fund Denmark center DiCyPS (Data-Intensive CPS). However, there have been a shift to more focus on national collaboration (with several Innovation Fund Denmark projects) and less EU projects. The later is partly caused by the fact that the ECSEL program has largely been without Danish participation during the period.

- *Maintain our high international visibility and recognition in the areas of model checking, logics and semantics for concurrent systems.*

  During the period series of papers on Dependency Graphs respectively Partial Order Reduction – demonstrating substantial performance of model checking in various settings – have been published at top-outlets earning the best-paper award at TACAS 2019. A decidable, yet useful subclass of Energy Timed Automata – a formalism introduced and popularized by DEIS researchers in previous periods – has been identified and applied to automatic synthesis of optimal control of hydraulic systems. This work received the best paper award at FM 2018. Several research results concerning various types of probabilistic systems, their semantics, characterizing logics, axiomatizations has been obtained, resulting in a record high number of publications at LICS. Also, the long-standing research on compositional specification theories have lead to a collaborative book on *Contracts* co-authored by several renowned international researchers. Finally, our leading position with respect to model checking real-time systems is witnessed by the co-authoring of a chapter in the Handbook of Model Checking released in 2018 after 10 years of joint effort by the community.

- *Strengthen further our focus on tool-development. This will sustain the bridge between foundational research and applications. To scale to the complexity of CPS, machine learning techniques as well as taking full advantages of the recently established MCC cluster will be pursued.*

The high standing in the research community is witnessed by several gold-medals at model-checking competitions won by the tool TAPAAL, record-high number of downloads of the tool UPPAAL (14.000 downloads per year) with supporting dissemination at several invited key-notes and lectures at PhD schools. The industrial recognition of UPPAAL is most clearly witness by the Grundfos Prize being awarded to Kim G Larsen in 2016.

To increase the quality and impact of our tool-development a permanent tool laboratory has been established lead by two research engineers, supported by varying numbers of student programmers.

During the period, DIES has acquired new significant investments and engaged in collaboration with the IT Support Staff at CLAAUDIA, in order to advance the High Performance and GPU infra-structure for large-scale model checking and Neural Network learning capabilities available for DEIS researchers.

Within UPPAAL Stratego new machine learning algorithms have been developed for the synthesis of control strategies. Within UPAAAL SMC recent improvements of the simulation engine enable by a prior static dependency analysis, has allowed for simulation of very large-scale Agent-Based models – e.g. more than 500.000 components within 3 hours and with provable linear-speed-up (for analysis of COVID-19).

- *The current trend in our foundational research in replacing Boolean verdicts with metric notions and quantitative results will be expanded.*

  Several new results has been obtained focusing on expressiveness, axiomatizability and computuation aspects and published in top outlets (e.g. several publications at LICS, CONCUR, LMCS, ao.). In several cases conference papers have been invited to publication in special-issues. Our long-term collaboration with leading international researchers has continued.

- *Invest substantial new effort in research on cloud computing, (wireless) networks and IoT.*

  The hiring of Stefan Schmid provided an enormously boost to the activity within this area with particular focus on Software Defined Networks.

Collaboration and Projects

- *We want to maintain current balance between basic research and application-oriented projects.*

The balance is still there, perhaps with a bit more focus on Basic Research.n Also the project portfolio over the period shows a good mix of large projects with collaboration between many partners and smaller more individual projects.

- *To involve SMEs, we want to maintain a fair number of smaller, short-term projects.* The interaction with SMEs are witnessed from the several short-term INfinIT Miniprojects (see Appendix).

- *Our strong international collaboration must be sustained (ECSEL, FET, COST actions, ..)*

  Partly due to the Danish withdrawal from the ECSEL EU program and the increased competetiveness of the FET program due to its move from the ICT program to the Excellence pillar, this has not fully been achieved. However, joint effort with DPW has acquired significant EU funding in the area of Energy.

- *National, University and Departmental Collaborations*

  DEIS has been heading the national innovation network InfinIT. As a main outcome of this 10 years effort, strengthened research collaboration between Danish CS department is ensured for the future by the newly funded Danish National Research Center for Digital Technologies, DIREC, as well as the ICT cluster DigitalLead.

  Also, DEIS has been engaged in the national effort on Flexible Energy in several national projects, including FED involving all 4 technical universities in Denmark. The Innovation Center DiCyPS marks a significant collaboration involving all research units in the department, addressing challenges within the transport and energy sectors.

  Significant Collaborations with other parts of the university have been made, e.g. CLAIRE (with the Engineeering Faculty), Reachi and BEOCOVID (with Department of Electronic Systems)

Research Management and Staff

- *Increase CISS VIP staff to minimize OH of projects.*

  The size of the CISS VIP staff has been maintained, and seems stabilized with recent recruitment. The effort made has been vital for matchmaking and implementation of demonstrators, not only for DEIS but for the entire

department, as witnessed by Tina Sihm winning the university matchmaking award. We are currently waiting for a departmental technicians to be hired.

- *Identify better ways of managing our several interconnected projects.*

  In the middle of the period we introduced more systematic bi-weekly meetings on our several Energy related projects have been made – some being co-meetings with the research unit DPW. Also, a new series of coordination meetings between full professors, group coordinator, financial and PR staff of the unit has recently been adopted in order to increase overall awareness.

- *Alignment and synergy between MSc/PhD/post-doc projects with industrial research projects.*

  Most of the DEIS research projects provide cases studies for our PhD students and PostDocs and has in several cases been accompanied by one of more student projects, e.g. the Reachi project with NEOCORTEC (a number of student projects, leading to patent) and DiCyPS (several student projects, even a pre-mega project with the research unit HCC, as well as leading to the spin-out company ATS).

- *Exploit existing financial sources for industrial PhDs/post-docs.*

  This line of applications has been pursued (unsuccessfully) twice.

- *Encourage members of the research unit to take sabbaticals.*

  Has only been pursued once.

Comments from Panel We also recall the recommendations of the previous research panel and detail below the take-up of these:

- *Continue to develop all areas across the spectrum from foundations through to applications. The expansion to consider the role of machine learning in model checking and the continued development of work on Cyber-Physical Systems are important developments*

  The work on UPPAAL Stratego has developed new learning algorithms for Euclidan Markov Decision processes, and has been applied to heating, transport and water-management systems. At the same time contributions to the theoretical foundation wrt. the convergence of the algorithms have been made. Also UPPAAL SMC has been used to derive Bayesian Networks from

Stochastic Hybrid models providing a novel link between model checking and machine learning.

The theory of Priced Timed Automata has found its way into efficient distributed implementations, which has been applied to scheduling of nanosatellites (with the company GOMSPACE) and optimal parallel evaluation of spreadsheets (in the PPP project).

- *Consider the introduction of a mentoring scheme for PhD students - separate from the scientific supervision.*

  During the period some four PhD students have left prematurely. It seems that the use of MUS (employee-development-talks) is not enough to prevent or catch these cases. New measures need to be identified to lower this number, e.g. in case of multiple supervisor, have more clearly defined agreements of responsibilities and expectations, engage PhD students in common reading groups etc.

- *Make better use of all funding schemes for PhD students and consider developing publicity material targeted at prospective students and emphasizing some of the applications work.*

  A new web-page has been made intending to presents DEIS in a more persuasive manner for future employees. For internal recruitment of PhD students targeted information-meetings have been made, with a reasonable degree of success. Still it is difficult to attract well qualified PhD candidates from EU. Attraction of PostDocs and visiting PhD students from EU – utilizing personal networks – has been significantly more successful.

- *Continue the reflection started with the SWOT analysis and decide which threats are significant and should be treated.*

  Loss of key persons was an identified treat, leading to the generous grant from Poul Due Jensen for hiring Max Tchaikovsky. Also the newly announced departmental position as "erhvervsambassador" may serve to fill DEIS's need for industrial liaison. More proactive in retaining key people.

  Concerning the lack of machine learning staff the hiring of Alvaro Torralba (planning) and Alessandro Tibo (deep learning) are first steps.

- *Work with the Department and other groups to develop a strategy for tool support.*

  We have now a staffed DEIS Tool Laboratory with two full time research-qualified developers, supported by several student programmers. In addition,

DEIS has contributed to defining the qualifications of the newly made call for technicians for the Departmental Demonstrator Laboratory.

4.7.1 Scientific Output and Impact

The publication list of the DEIS unit includes 338 publications and 1 patent, distributed among the different categories as follows:

- 96 - Refereed Journals
- 203 - Refereed Conferences and Workshops
- 8 - Book chapters
- 1 - Books
- 10 - Edited publications
- 20 - Technical Reports
- 1 - Patents

Below we report the number of unit's publications ranked according to a system that is consistent with both the Danish National Ranking of journals (BFI ranking) and CORE-ranking[2]. Full details about the adopted ranking can be found in the Appendix.

	A	B	C	**Total**
JOURNALS	38 (46,91%)	33 (40,74%)	10 (12,35%)	**81**
CONFERENCES	72 (42,60%)	63 (37,28%)	34 (20,12%)	**169**

The table above shows that the publication strategy of the unit favors to a great extent highly ranked outlets, both for journals and conferences.

We observe that the top-tier outlets are the most selective. In particular, publication to A ranked conferences is highly competitive, with acceptance rates typically below those of the best journals. Moreover, the acceptance rate at tier A conferences is always below 20%.

Comparison with previous evaluation period

Next we report the number of publications of the old DES and MI units from the previous evaluation period 2011-2016.

[2]In partial agreement with the other units, we used only three categories, A, B and C. Conferences and journals ranked as A* in CORE-ranking are reported under the category A.

DES	A	B	C	Total
JOURNALS	29 (46,03%)	29 (46,03%)	5 (7,94%)	**63**
CONFERENCES	38 (21,71%)	93 (53,14%)	44 (25,14%)	**175**

MI	A	B	C	Total
JOURNALS	16 (45,71%)	16 (45,71%)	3 (8,57%)	**35**
CONFERENCES	12 (13,48%)	50 (56,18%)	27 (30,34%)	**89**

A fair comparison can be done only qualitatively, because by considering the sum of the total publications from two old units would not take into account that the reported numbers for MI include also publications from researchers that did not join DEIS.

Notably, the number of publications in A-raked journals and conferences increased without lowering the overall total numbers of publications. This is consistent with the publication strategy pursued by the unit in the last period.

It is worth mentioning that until 2019, due to some cuts to the economy of the department, all permanent staff has been required to increase their teaching load. Remarkably, this did not affect substantially the research output that increased in quality. This was possible only thanks to the extraordinary effort put by the single researchers, which should be praised and admired. In 2020, the departmental economy considerably improved and allowed the unit to hire new scientific staff while lowering the teaching load.

4.7.2 Scientific Services and Recognition

Members of the DEIS provided substantial services to the research community, indicating a significant leadership and recognition. This include serving on the editorial boards of journals (e.g., *Formal Methods in System Design*, *Theoretical Computer Science*, and *Artificial Intelligence Journal*) and serving as PC-chairs for the conferences NWPT'16, RP'16, SynCoP'16, WATA'16, MFCS'17, SETTA'17, FMICS'19, HSB'20, PGM'20, and TACAS'21.

Moreover, members of the unit serve on the steering committee on a number of leading international conferences including ETAPS, TACAS and CONCUR, as well on EMSIG the Embedded Systems Special Interest Group. This is in addition to a very significant number of program committee memberships.

The international recognition is also witnessed from the fact that the unit organized 8 international conferences and workshops: WATA'16, NWPT'16, RP'16, KimFest, MFCS'17, NordSec'19, PMG'20, HSDIP@ICAPS'20.

Within the reporting period 2016-2020, one member of the unit was the recipient of the prestigious ERC-Advanced Grant, won the Grundfos prize 2016, and has been appointed distinguished Professor at Northeastern University of China. In addition he has been serving as member of the ERC evaluation panel for PE6 Consolidator Grants, as INRA international chair member in 2016, as well as a member of EATCS annual award committee and CAV Annual Award committee. In the period he served as member on several advisory boards and was appointed Foreign Expert of China.

In the reporting period, members of the unit received several other awards including, a number of gold medals in the prestigious Model Checking Contest in 2018 (two gold medals), 2019 (four gold medals), and 2020 (two gold medals) for software tool TAPAAL, and several best paper awards at conferences such as PetriNets'17, ICTAC'17, FM'18, IFIP NETWORKING'18, ACM SOSR'18, and TACAS'19.

In addition, several conference publications from the unit have been invited for journal publication in special issues. Also, M.Sc. students of the unit are performing at a high level witnessed by several publications during the study.

4.8 Plans

In the following we detail the plans for future research in terms of research directions, possible collaborations and projects for the group, as well as measure to be taken for future management and development.

4.8.1 Research Directions

We will strive to maintain our strong international collaboration and leading role within Embedded and Cyber-Physical Systems, and our high international visibility and recognition in the areas of model checking, tools and foundational theories covering logics and semantics for concurrent systems. The merge with the Programming Technology researcher provides new opportunities for our transfer of methodologies.

Tools and Model Checking In particular we will strengthen our focus on tool development as this will sustain the bridge between foundational research and applications. We shall continue the effort in combining symbolic model checking with machine learning to obtain improved performance. Also we shall consolidate

and expand the DEIS Tool-Development Laboratory, aiming for joint effort and knowledge sharing across individual tools, as well as further exploit the Commercialization effort.

Foundational Theories Our influential work on behavioural metrics and their logic characterization will be expanded to deal with infinite-state dynamical systems. In particular, and continuing initial research of the current period, we want to use metrics in the theoretical study of guarantees that can be made by various methods for learning behavioural models from sampled observations (time-series) even in settings of partial observability. We are currently collaborating with Reikjavik University on using EM (expectation maximization) for learning probabilistic behavioural models. It is also planned to continue our successful line of work on the reduction of nonlinear biochemical models. Ongoing research is concerned with on-the-fly and optimality-preserving reduction algorithms for nonlinear differential equations and involves collaborators from the University of Oxford, IMT Lucca and Sant'Anna Pisa.

Programming Technology Research on Programming Technologies will now be part of the groups profile as members of the programming technology subgroup has joined the group. There has for many years been strong collaborations on embedded systems development and more recently on parallelization of spreadsheet computations and energy aware software development. Work will build on semantic foundations, static analysis and abstract interpretation, as well as constructive and imperative approaches, which can be applied on programs but also on models.

We expect to pursue research in energy aware software development and ICT Infrastructure (e.g. data center), where models can be used e.g. by viewing data centers as a component in a complete energy system.

Security This is an important research area with significant industrial interest and teaching responsibilities given the new Master Edication in Copenhagen on Cyber Security. So far the contributions from DEIS has been offered by a few core members. Within the next period we want to grow substantially in this area, in terms of staff and research productivity. Iniital focus will be on security of IoT and embedded systems exploiting our existing strong competences in model checking and static analysis. We expect to also involve theorem proving exploiting the competence of our new staff member Anders Schlichtkrull.

Artificial Intelligence and Machine Learning We intend to increase staff and research activity significantly within this area. The recent advent of Alvaro Torrabella adds core competences of planning to the group. In the next period we plan to investigate various combinations of planning and planning and reach out to other departments in the university working with operative research. Concretely we have agreed to give a joint PhD course on planning with the operations research group from the Department of Materials and Production.

Other directions to be pursued includes verification (of robustness) of learning-enabled systems (e.g. deep Neural Networks), explainability through learning of explainable models (e.g. Bayesian networks, Decision Trees, Automata of various types). Complementing the symbolic methods we will further pursue the use of reinforcement learning as a way of synthesizing strategies for (infinite-state) Markov Decision Processes. Here collaboration with Professor Bin Yang and his group from the DPW group is envisaged.

Network Modelling and Verification At the end of 2018 Stefan Schmid left the group accepting a Full Professorship position at the University of Vienna. The collaboration with Stefan Schmid, is still active through the DFF project *Quantitative Analysis and Synthesis of Network Protocols* as well as through the Vienna Research Project *WHATIF* and we intend to strengthen our activities and staffing in this area in the next period.

Methodology and Applications We intend to continue to engage in having our generic methods and tools being applied within the ICT-sector itself (e.g. embedded systems) but also in other application areas. To achieve a higher level of impact additional effort needs to be made to make our methods and tools accessible for researcher and practitioners from other fields, possibly requiring domain specific customization – expected to be done with the DEIS Tool-Development Laboratory – and demonstrators – expected to be done with the Departmental Demonstrator Laboratory. Effort needs to be taken to ensure that these activities produce a significant number of publications.

4.8.2 Research Collaborations and Projects

We want to maintain the current balance between basic research and application-oriented projects. Also, we will continue to pursue successful big collaborative projects but still have attention to individual projects as these are crucial for the academic careers of young research members. We also intend to balance between

applying for (hard to get) individual prestigious grants (ERC, Villum YIP, DFF Sapere Aude) and more accessible funding opportunities. We hope to benefit from the Funding support of the Department to identify relevant opportunities. Also we hope to benefit from the coming Erhvervsambassador of the Department to serve as a liaison to danish companies, in particular for our several international staff members.

Collaboration Concrete collaboration and funding plans include: exploit DIREC the new National Research Center for Digital Technologies, the ICT Cluster Digital-Lead. We want to continue the collaboration with other CS researcher on Energy and Transport initiated in the DiCyPS innovation network. We want to continue collaboration with researchers from Civil Engineering on water management. In particular we want to revitalize our international collaborations in Embedded and Cyber-Physical Systems – exploiting that Denmark is now again a contributing partner of the ECSEL program. Finally, efforts toward ERC Proof of Concept applications and ERC Synergy collaborations are envisaged.

4.8.3 Research Management and Staff

In the next hiring rounds we hope to attract more highly qualified applicants at all levels, with a particular focus on attracting an increased number of highly qualified applicants to our PhD position. We aim to make more aggressive use of personal networks and advertise more clearly the "tenure track" possibility.

Concerning future growth of the (new) DEIS group, we have particular focus on the areas of Artificial Intelligence and Machine Learning, Security, Network and Programming Technologies. The staffing of the DEIS Tool Laboratory needs to be made permanent, and as for the staff of the Departmental Demonstrator Laboratory we seek two main competences 1) digital twins and virtual demonstrators and 2) IoT and physical demonstrators.

We intend to identifies ways to decrease the drop-out of PhD students, by introducing mentoring support functions, more extensive use of study groups, as well as social arrangements. In the likely situation of a continued lock-down due to Corona this should have even more attention.

For Assistant Professor that do not hold a tenure-track position, we will try to establish a principle of a 1-2 year extension based on involvement in external projects. First agreements of this type has already been made with the Department. In the same manner we want the mentoring offered to tenure-track Assistant Professors to be extended to non-tenure track Associate Professors.

The research activities are increasingly relying on the availability of the computing facilities, in particular the presence of the Model Checking Cluster. Whereas the staff of CLAAUDIA provides ample support for the general set and operation of these computing facilities, it is important that we within DEIS have own permanent staff, whose responsibility include interfacing to CLAAUDIA.

The new addition of the Programming Technology research group fit in many ways ideally into the research of DEIS. Also the research competence on Theorem Proving coming from the expansion in Copenhagen complements nicely the existing verification competences on Model Checking and Static Analysis. However, these additions also adds significantly to the size of the group, which may eventually become too big. To counter this effort must be taken to maintain social coherence of the group.

4.9 Sustainable Development Goals

Here we summarize the contributions of DEIS to a variety of SDGs. For more in-depth descriptions we refer to subsection of Application Areas in Section **??**.

- SDG 3: Good Health and Well-Being
  The contributions here include work on prediction of home care deliverables with a high risk of deviations based on analysizes of data from the home care sector. Also prediction of patient flow in the emergency deparment at Aalborg University hopsital in order to improved the quality of service (reduction of waiting tiem, better planning of staff).

  At the end of the period, we have worked intensily on applying UPPAAL to modelling, prediction and control of the COVID-19 pandemic. In particular an Agent-based model of Northern Jutland (and most recently the entire country) has been made allowing for estimation of the impact of various lock-down or opening initiatives. These findings have been made available to the governments Official Reference Group on COVID-19. Most recently, the group has become member of the Modelling Expert Group at SSI (Statens Serum Institute).

- SDG 6: Clean Water and Sanitation
  We have contributed to the clean water agenda. In particular development low-cost sensor networks and the data resulting these make a holistic real-time view of urban water systems possible. Most recently we are applying

our techniques for synthesizing safe and sustainable control strategies for water resource management (initial focus on water ponds).

- SDG 7: Affordable and Clean Energy
  Through several years – including the entire evaluation period – the group has been particular active in applying our tools (wireless sensor networks, grid computing, performance analysis and optimal control) in various part of the energy sector ranging from energy-aware building to smart grids with an overall ambition of digitizing the energy sector, resulting in increased comfort, reduced CO2 emission as well as overall energy consumption.

- SDG 11: Sustainable Cities and Communities
  Substantial effort has been made to optimize traffic flow.In the context of road networks, we have developed a new machine learning techniques for machine learning on road network graphs. In cities our model checking and machine learning techniques for synthesizing safe and optimal controllers has been applied to on-line control of traffic lights resulting in 20%-30% reductions of waiting time and CO2 emission. As a particular impact, ATS (Advanced Traffic System) has emerged as a spin-out company.

- SDG 13: Climate Action
  The work mentioned under SDG 7 and SDG 11 above has also had significant implicit impact on SDG 13, as optimal digital support for renewable energy in the energy system as well as reduction in fuel-consumption from the transport sector will enable massive greenhouse gas reductions.

4.10 AI and Machine Learning

As described earlier in this report, the DEIS unit was at the beginning of the period the result of merging the former DES (Distributed and Embedded Systems) research unit with the former MI (Machine Intelligence) research Unit. As such AI and Machine Learning has been one of the defining research directions of DEIS. Particular focus has been on securing the existing stronghold on probabilistic graphical models, and to establish connections between machine learning and formal methods. Here we give a short summary of the work on AI and Machine Learning and refer to Section 4.5.5 for a more detailed description.

- In the field of probabilistic graphical models, substantial effort has been made towards scalable inference and learning techniques for latent variable models.

- In the field of relational learning we have been investigating foundational principles and algorithmic solutions for modeling, learning, and inference with multi-relational graph or network data.
- Connections between Machine Learning and Formal Methods have been explored in several directions. This includes new reinforcement learning methods for constructing controllers, theoretical investigation towards understanding the learning methods and convergence properties, and a framework for learning Beyesian networks from Stochastic Timed Automata.
- A traditional area in AI is that of planning. During the period substantial work on applying (and modifying) model checking towards optimal planning has been made, with a number of industrial applications. Also work on specifying hierarchical planning problems has been made.

4.11 Committee Evaluation

Observations The research of the DEIS group is mainly concerned with the modelling, analysis and realization of distributed, embedded and intelligent systems. The report written by the group shows that it has continued its excellent work in the last evaluation period (2016–2020). DEIS is well known for the development, implementation, and application of state-of-the-art model checking tools, with an emphasis on the quantitative properties. It has kept and consolidated its standing as one of the world-leading groups in this area.

The group has shown an excellent performance with respect to the key performance indicators publications (both w.r.t. quality and quantity), external funding, scientific service, academic and industrial collaborations as well as dissemination of results. It has published more than 300 papers in mostly highly ranked refereed journals, conferences, and workshops. In the evaluation period, the group received funding of more than 70 million DKK for overall 39 projects, among them a very prestigious ERC Advanced Grant. Several of these projects addressed sustainable development goals (e.g., in the context of power consumption and water management). The group has organized eight international conferences and workshops in Aalborg, several group members have served as PC chairs of international conferences and workshop (Larsen, Srba, Bacci, Tschaikowski, Jaeger, Nielsen, Jensen) and group members have been in the PCs of a high number of workshops and conferences. Other outstanding services include Kim Larsen's membership of the panel for ERC Consolidator grants in 2017, 2019 and 2021. DEIS members have collaborated with a high number of well-known international scientists (e.g.,

Dana Scott, Dexter Kozen, Gordon Plotkin, and Luca Cardelli). In addition, the group has collaborated both with industry and with public administrations, and has established two spin-out companies (ATS, VeriAAL ApS).

What the reviewers find most impressive regarding the research performed by DEIS is that the group is excellent in all three areas:

1. Theoretical foundations and algorithm development
2. Development of high-performance tools for the analysis and synthesis of systems
3. Real world applications of theoretical results and systems

Success in 1) is demonstrated by a high number of publications in excellent venues such as LICS, ICALP, CONCUR, and TACAS; the h-index of the scientists (with an outstanding Kim Larsen, but also very good numbers for other members of the group, such as Jiri Srba, Thomas D. Nielsen, and Manfred Jaeger), best paper awards in highly ranked conferences, such as TACAS, Formal Methods (FM), PetriNets, ICTAC, IFIP NETWORKING, and ACM SOSR, and prestigious personal research grants (mainly Kim Larsen) as well as other foundational grants.

Success in 2) is shown by winning many medals in system competitions (notably for the TAPAAL tool) and very high download numbers for the UPPAAL system, as well as the successful application of these and other systems in different projects.

Success in 3) is shown by industrial and government applications projects such as the GOMspace satellite application and modeling of the Covid19 spreading using UPPAAL. As a result of the latter, several members of DEIS are now officially members of the SSI Expert Modeling Group.

The work in these areas is not done separately. Instead, these areas complement each other in an ideal way, with new theoretical results being integrated in the existing systems, thus enabling new applications, and demand from applications triggering the investigation of new theoretical problems.

The group is rather large, but appears to be nevertheless very coherent, with a lot of collaborations within the group. We recommend to actively keep on promoting this coherence, and find the size unproblematic, i.e., we do not see a reason for splitting the group into subgroups. In particular, we want to point out that the Programming Technology sub-group fits well into DEIS, though the reasons for this merger and the expected benefits could have been made clear in the documents provided for the reviewers. There has already been successful cooperation between DEIS and this group before the merger, and we expect that this will intensify now.

Of the wide range of excellent research results listed in the report and shown in the

presentation, we mention a few, to highlight the broad and excellent contributions of the group in different directions:

- The world leading research on quantitative analysis of systems based on behavioral metrics was continued.
- The research regarding the analysis of energy bounded infinite runs (with pump and satellite application as motivation) has determined a decidable subclass of an in general undecidable problem. This resulted in a speedup of several orders of magnitude compared to other modeling approaches in the motivating applications.
- The work on abstract dependency graphs allows (among other things) very fast integration of new approaches into the existing high-performance systems.
- In Machine Learning, the group has not only successfully continued its research on probabilistic graphical models and relational learning, but has started to investigate the combination of Formal Methods and Machine Learning approaches. This is a very interesting new direction, which has already produced several high-ranked publications.

Recommendations

- Keep up the good work!
- The combination of Formal Methods (FM) and Machine Learning (FM) is a very promising and topical direction, which should be supported by hiring additional researchers with ML expertise.
- It is very good that the department has provided support to hire people to support tool development. We strongly support this since is essential to keep up the excellent work on tool development and applications.
- Combination of Model Checking and Theorem Proving (in particular, interactive theorem proving, with new expertise in Copenhagen) should be intensified, which also requires hiring additional researchers with expertise in Theorem Proving.
- The group has many very good researchers. We recommend that also the younger staff members are encouraged to and supported when applying for

prestigious personal grants such as ERC and Villum Young Investigators to further enhance their scientific reputation.

- PhD students: the group has increased the number of finished PhDs compared to the last evaluation period; hiring more PhD students appears to be hard due to structural and financial constraints. The group has improved the supervision of PhD and thus increased the success rate, but it is important to keep on monitoring the situation and develop further improvements if necessary.

- Diversity & inclusion is an issue for the whole department, but in particular for the DEIS group. Therefore, we strongly recommend implementing the departmental recommendations on this topic also at group level.

- Long term consideration: the success of the group currently appears to depend strongly on the excellence of the group leader and a few key staff members. While this is not an immediate problem for the next evaluation period, it should at the latest be addressed in the one after. We recommend developing strategies for finding equally excellent successors (build up internally or hire from outside).

Bibliography

[1] Andrei-Alexandru Agape, Madalin Claudiu Danceanu, René Rydhof Hansen, and Stefan Schmid. Charting the security landscape of programmable dataplanes. *CoRR*, abs/1807.00128, 2018.

[2] Mads Kronborg Agesen, Søren Enevoldsen, Thibaut Le Guilly, Anders Mariegaard, Petur Olsen, and Arne Skou. Energy consumption forecast of photovoltaic comfort cooling using UPPAAL stratego. In *Models, Algorithms, Logics and Tools*, volume 10460 of *Lecture Notes in Computer Science*, pages 603–622. Springer, 2017.

[3] Mads Kronborg Agesen, Kim Guldstrand Larsen, Marius Mikučionis, Marco Muñiz, Petur Olsen, Thomas Pedersen, Jirí Srba, and Arne Skou. Toolchain for user-centered intelligent floor heating control. In *IECON*, pages 5296–5301. IEEE, 2016.

[4] Mads Kronborg Agesen, Arne Joachim Skou, and Keld Lotzfeldt Pedersen. Preliminary report: Controller prototyping and validation for photo-voltaic comfort cooling. In Per Kvols Heiselberg, editor, *CLIMA*, Denmark, 2016. Department of Civil Engineering, Aalborg University.

[5] Bernhard K. Aichernig, Klaus Hörmaier, Florian Lorber, Dejan Nickovic, and Stefan Tiran. Require, test, and trace IT. *Int. J. Softw. Tools Technol. Transf.*, 19(4):409–426, 2017.

[6] Michele Albano, Rodrigo Castineira, Chloe Desdouits, Luis Lino Ferreira, Thibaut Le Guilly, Inge Isasa, Jani Jokinen, Kaspars Kondratjevs, Nadezhda Kunicina, Lorenzo Manero, Aitor Milo, Javier Monge, Claude Le Pape, Per D. Pedersen, Torben Bach Pedersen, Petur Olsen, Laurynas Siksnys, Arne Joachim Skou, Radislav Smid, Rafael Socorro, Petr Stluka, and Anatolijs Zabasta. Application system design - energy optimisation. In Jerker Delsing,

editor, *IoT Automation: Arrowhead Framework*, pages 211–247. CRC Press, February 2017.

[7] Michele Albano, Luis Lino Ferreira, Giovanni Di Orio, Pedro Maló], Godfried Webers, Erkki Jantunen, Iosu Gabilondo, Mikel Viguera, and Gregor Papa. Advanced sensor-based maintenance in real-world exemplary cases. *Automatika*, 61(4):537–553, 2020.

[8] Michele Albano, Matteo Mordacchini, and Laura Ricci. Aoi-based multicast routing over voronoi overlays with minimal overhead. *IEEE Access*, 8:168611–168624, 2020.

[9] Michele Albano and Brian Nielsen. Interoperability by construction: code generation for arrowhead clients. In *3rd IEEE International Conference on Industrial Cyber-Physical Systems (ICPS 2020)*, pages 5519–5524. IEEE, 2020.

[10] Rajeev Alur, Mirco Giacobbe, Thomas A. Henzinger, Kim G. Larsen, and Marius Mikučionis. Continuous-time models for system design and analysis. In *Computing and Software Science*, volume 10000 of *Lecture Notes in Computer Science*, pages 452–477. Springer, 2019.

[11] Saeed Akhoondian Amiri, Szymon Dudycz, Stefan Schmid, and Sebastian Wiederrecht. Congestion-free rerouting of flows on dags. In *ICALP*, volume 107 of *LIPIcs*, pages 143:1–143:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[12] Saeed Akhoondian Amiri, Klaus-Tycho Foerster, Riko Jacob, Mahmoud Parham, and Stefan Schmid. Waypoint routing in special networks. In *Networking*, pages 496–504. IFIP, 2018.

[13] Saeed Akhoondian Amiri, Klaus-Tycho Foerster, Riko Jacob, and Stefan Schmid. Charting the algorithmic complexity of waypoint routing. *Comput. Commun. Rev.*, 48(1):42–48, 2018.

[14] Saeed Akhoondian Amiri, Klaus-Tycho Foerster, and Stefan Schmid. Walking through waypoints. In *LATIN*, volume 10807 of *Lecture Notes in Computer Science*, pages 37–51. Springer, 2018.

[15] Sergio Feo Arenis, Bernd Westphal, Daniel Dietsch, Marco Muñiz, Ahmad Siyar Andisha, and Andreas Podelski. Ready for testing: ensuring conformance to industrial standards through formal verification. *Formal Aspects Comput.*, 28(3):499–527, 2016.

[16] Jacinto Arias, José A. Gámez, Thomas D. Nielsen, and José Miguel Puerta. A scalable pairwise class interaction framework for multidimensional classification. *Int. J. Approx. Reason.*, 68:194–210, 2016.

[17] Pranav Ashok, Jan Kretínský, Kim Guldstrand Larsen, Adrien Le Coënt, Jakob Haahr Taankvist, and Maximilian Weininger. SOS: safe, optimal and small strategies for hybrid markov decision processes. In *QEST*, volume 11785 of *Lecture Notes in Computer Science*, pages 147–164. Springer, 2019.

[18] Aslan Askarov, René Rydhof Hansen, and Willard Rafnsson, editors. *Secure IT Systems - 24th Nordic Conference, NordSec 2019, Aalborg, Denmark, November 18-20, 2019, Proceedings*, volume 11875 of *Lecture Notes in Computer Science*. Springer, 2019.

[19] Chen Avin, Louis Cohen, and Stefan Schmid. Competitive clustering of stochastic communication patterns on a ring. In *NETYS*, volume 10299 of *Lecture Notes in Computer Science*, pages 231–247, 2017.

[20] Chen Avin, Alexandr Hercules, Andreas Loukas, and Stefan Schmid. *rDAN*: Toward robust demand-aware network designs. *Inf. Process. Lett.*, 133:5–9, 2018.

[21] Chen Avin, Kaushik Mondal, and Stefan Schmid. Demand-aware network designs of bounded degree. In *DISC*, volume 91 of *LIPIcs*, pages 5:1–5:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[22] Chen Avin and Stefan Schmid. Toward demand-aware networking: a theory for self-adjusting networks. *Comput. Commun. Rev.*, 48(5):31–40, 2018.

[23] Victor Anthony Arrascue Ayala, Kemal Cagin Gülsen, Anas Alzogbi, Michael Färber, Marco Muñiz, and Georg Lausen. A delay-robust touristic plan recommendation using real-world public transportation information. In *RecTour@RecSys*, volume 1906 of *CEUR Workshop Proceedings*, pages 9–17. CEUR-WS.org, 2017.

[24] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. Complete axiomatization for the bisimilarity distance on markov chains. In *CONCUR*, volume 59 of *LIPIcs*, pages 21:1–21:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[25] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. On-the-fly computation of bisimilarity distances. *Log. Methods Comput. Sci.*, 13(2), 2017.

[26] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. Complete axiomatization for the total variation distance of markov chains. In *MFPS*, volume 341 of *Electronic Notes in Theoretical Computer Science*, pages 27–39. Elsevier, 2018.

[27] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. A complete quantitative deduction system for the bisimilarity distance on markov chains. *Log. Methods Comput. Sci.*, 14(4), 2018.

[28] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, and Radu Mardare. Converging from branching to linear metrics on markov chains. *Math. Struct. Comput. Sci.*, 29(1):3–37, 2019.

[29] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, Radu Mardare, Qiyi Tang, and Franck van Breugel. Computing probabilistic bisimilarity distances for probabilistic automata. In *CONCUR*, volume 140 of *LIPIcs*, pages 9:1–9:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[30] Giorgio Bacci, Giovanni Bacci, Kim G. Larsen, Radu Mardare, Qiyi Tang, and Franck van Breugel. Computing probabilistic bisimilarity distances for probabilistic automata. *Log. Methods Comput. Sci.*, 2019. to appear.

[31] Giorgio Bacci, Robert Furber, Dexter Kozen, Radu Mardare, Prakash Panangaden, and Dana S. Scott. Boolean-valued semantics for the stochastic λ-calculus. In *LICS*, pages 669–678. ACM, 2018.

[32] Giorgio Bacci, Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. An algebraic theory of markov processes. In *LICS*, pages 679–688. ACM, 2018.

[33] Giovanni Bacci, Giorgio Bacci, Kim G. Larsen, and Radu Mardare. On the metric-based approximate minimization of markov chains. In *ICALP*, volume 80 of *LIPIcs*, pages 104:1–104:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[34] Giovanni Bacci, Giorgio Bacci, Kim G. Larsen, and Radu Mardare. On the metric-based approximate minimization of markov chains. *J. Log. Algebraic Methods Program.*, 100:36–56, 2018.

[35] Giovanni Bacci, Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Pierre-Alain Reynier. Optimal and robust controller synthesis - using energy timed automata with uncertainty. In *FM*, volume 10951 of *Lecture Notes in Computer Science*, pages 203–221. Springer, 2018.

[36] Giovanni Bacci, Mikkel Hansen, and Kim Guldstrand Larsen. Model checking constrained markov reward models with uncertainties. In *QEST*, volume 11785 of *Lecture Notes in Computer Science*, pages 37–51. Springer, 2019.

[37] Arsany Basta, Andreas Blenk, Szymon Dudycz, Arne Ludwig, and Stefan Schmid. Efficient loop-free rerouting of multiple SDN flows. *IEEE/ACM Trans. Netw.*, 26(2):948–961, 2018.

[38] Arsany Basta, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. Logically isolated, actually unpredictable? measuring hypervisor performance in multi-tenant sdns. *CoRR*, abs/1704.08958, 2017.

[39] Albert Benveniste, Benoît Caillaud, Dejan Nickovic, Roberto Passerone, Jean-Baptiste Raclet, Philipp Reinkemeier, Alberto L. Sangiovanni-Vincentelli, Werner Damm, Thomas A. Henzinger, and Kim G. Larsen. Contracts for system design. *Found. Trends Electron. Des. Autom.*, 12(2-3):124–400, 2018.

[40] Albert Benveniste, Kim Larsen, and Jean-Baptiste Raclet. Mixed nondeterministic-probabilistic interfaces. Technical report, Inria Rennes Bretagne Atlantique, 2020.

[41] Marcin Bienkowski, Jan Marcinkowski, Maciej Pacut, Stefan Schmid, and Aleksandra Spyra. Online tree caching. In *SPAA*, pages 329–338. ACM, 2017.

[42] Marcin Bienkowski, Nadi Sarrar, Stefan Schmid, and Steve Uhlig. Online aggregation of the forwarding information base: Accounting for locality and churn. *IEEE/ACM Trans. Netw.*, 26(1):591–604, 2018.

[43] Andreas Blenk, Arsany Basta, Laurenz Henkel, Johannes Zerwas, Wolfgang Kellerer, and Stefan Schmid. perfbench: A tool for predictability analysis in multi-tenant software-defined networks. In *SIGCOMM Posters and Demos*, pages 66–68. ACM, 2018.

[44] Andreas Blenk, Patrick Kalmbach, Wolfgang Kellerer, and Stefan Schmid. o'zapft is: Tap your network algorithm's big data! In *BigDAMA@SIGCOMM*, pages 19–24. ACM, 2017.

[45] Andreas Blenk, Patrick Kalmbach, Johannes Zerwas, Michael Jarschel, Stefan Schmid, and Wolfgang Kellerer. Neurovine: A neural preprocessor for your virtual network embedding algorithm. In *INFOCOM*, pages 405–413. IEEE, 2018.

[46] Marcel Blöcher, Malte Viering, Stefan Schmid, and Patrick Eugster. The grand CRU challenge. In *HotConNet@SIGCOMM*, pages 7–11. ACM, 2017.

[47] Thomas Bøgholm, Kim G. Larsen, Marco Muñiz, Bent Thomsen, and Lone Leth Thomsen. Analyzing spreadsheets for parallel execution via model checking. In *Models, Mindsets, Meta*, volume 11200 of *Lecture Notes in Computer Science*, pages 27–35. Springer, 2018.

[48] Frederik Bønneland, Jakob Dyhr, Peter Gjøl Jensen, Mads Johannsen, and Jirí Srba. Simplification of CTL formulae for efficient model checking of petri nets. In *Petri Nets*, volume 10877 of *Lecture Notes in Computer Science*, pages 143–163. Springer, 2018.

[49] Frederik M. Bønneland, Peter Gjøl Jensen, Kim G. Larsen, Marco Muñiz, and Jirí Srba. Partial order reduction for reachability games. In *CONCUR*, volume 140 of *LIPIcs*, pages 23:1–23:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[50] Frederik M. Bønneland, Peter Gjøl Jensen, Kim Guldstrand Larsen, Marco Muñiz, and Jirí Srba. Start pruning when time gets urgent: Partial order reduction for timed systems. In *CAV (1)*, volume 10981 of *Lecture Notes in Computer Science*, pages 527–546. Springer, 2018.

[51] Michael Borokhovich, Yvonne Anne Pignolet, Stefan Schmid, and Gilles Trédan. Load-optimal local fast rerouting for dense networks. *IEEE/ACM Trans. Netw.*, 26(6):2583–2597, 2018.

[52] Michael Borokhovich, Clement Rault, Liron Schiff, and Stefan Schmid. The show must go on: Fundamental data plane connectivity services for dependable sdns. *Comput. Commun.*, 116:172–183, 2018.

[53] Abdeldjalil Boudjadar, Alexandre David, Jin Hyun Kim, Kim G. Larsen, Marius Mikučionis, Ulrik Nyman, and Arne Skou. Statistical and exact schedulability analysis of hierarchical scheduling systems. *Sci. Comput. Program.*, 127:103–130, 2016.

[54] Jalil Boudjadar, Jin Hyun Kim, Linh Thi Xuan Phan, Insup Lee, Kim G. Larsen, and Ulrik Nyman. Generic formal framework for compositional analysis of hierarchical scheduling systems. In *ISORC*, pages 51–58. IEEE Computer Society, 2018.

[55] Jalil Boudjadar, Saravanan Ramanathan, Arvind Easwaran, and Ulrik Nyman. Combining task-level and system-level scheduling modes for mixed criticality systems. In *DS-RT*, pages 1–10, 2020.

[56] Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, Joël Ouaknine, and James Worrell. Model checking real-time systems. In *Handbook of Model Checking*, pages 1001–1046. Springer, 2018.

[57] Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen. Average-energy games. *Acta Informatica*, 55(2):91–127, 2018.

[58] Rafael Cabañas, Ana M. Martínez, Andrés R. Masegosa, Darío Ramos-López, Antonio Salmerón, Thomas D. Nielsen, Helge Langseth, and Anders L. Madsen. Financial data analysis with pgms using AMIDST. In *ICDM Workshops*, pages 1284–1287. IEEE Computer Society, 2016.

[59] Jaime Campos, Pankaj Sharma, Michele Albano, Luis Lino Ferreira, and Martin Larra naga. An open-source framework approach to support condition monitoring and maintenance. *MPDI Applied Sciences*, 10(18):6360, 2020.

[60] Jaime Campos, Pankaj Sharma, Michele Albano, Erkki Jantunen, David Baglee, and Luis Lino Ferreira. Arrowhead framework services for condition monitoring and maintenance based on the open source approach. In *CoDIT*, pages 697–702. IEEE, 2019.

[61] Marco Canini, Iosif Salem, Liron Schiff, Elad Michael Schiller, and Stefan Schmid. A self-organizing distributed and in-band SDN control plane. In *ICDCS*, pages 2656–2657. IEEE Computer Society, 2017.

[62] Marco Canini, Iosif Salem, Liron Schiff, Elad Michael Schiller, and Stefan Schmid. Renaissance: A self-stabilizing distributed SDN control plane. In *ICDCS*, pages 233–243. IEEE Computer Society, 2018.

[63] Falke B. Ø. Carlsen, Jacob J. Rasmussen, Mathias M. Sørensen, Nicolaj Ø. Jensen, and Michele Albano. Generation of realistic activity scenarios for sumo. In *Proceedings of the 17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services (MobiQuitous 2020)*. EAI, 2020.

[64] Franck Cassez, Peter G. Jensen, and Kim G. Larsen. Verification and parameter synthesis for real-time programs using refinement of trace abstraction. *Fundamenta Informaticae*, 178, 2020.

[65] Franck Cassez, Peter Gjøl Jensen, and Kim Guldstrand Larsen. Refinement of trace abstraction for real-time programs. In *RP*, volume 10506 of *Lecture Notes in Computer Science*, pages 42–58. Springer, 2017.

[66] Mounir Chadli, Jin Hyun Kim, Kim G. Larsen, Axel Legay, Stefan Naujokat, Bernhard Steffen, and Louis-Marie Traonouez. High-level frameworks for the specification and verification of scheduling problems. *Int. J. Softw. Tools Technol. Transf.*, 20(4):397–422, 2018.

[67] Lizzie Coles-Kemp and René Rydhof Hansen. Walking the line: The everyday security ties that bind. In *HCI (22)*, volume 10292 of *Lecture Notes in Computer Science*, pages 464–480. Springer, 2017.

[68] Levente Csikor, Christian Esteve Rothenberg, Dimitrios P. Pezaros, Stefan Schmid, László Toka, and Gábor Rétvári. Policy injection: A cloud dataplane dos attack. In *SIGCOMM Posters and Demos*, pages 147–149. ACM, 2018.

[69] Andreas Engelbredt Dalsgaard, Søren Enevoldsen, Peter Fogh, Lasse S. Jensen, Peter Gjøl Jensen, Tobias Skovgaard Jepsen, Isabella Kaufmann, Kim G. Larsen, Søren M. Nielsen, Mads Chr. Olesen, Samuel Pastva, and Jirí Srba. A distributed fixed-point algorithm for extended dependency graphs. *Fundam. Informaticae*, 161(4):351–381, 2018.

[70] Andreas Engelbredt Dalsgaard, Søren Enevoldsen, Peter Fogh, Lasse S. Jensen, Tobias Skovgaard Jepsen, Isabella Kaufmann, Kim G. Larsen, Søren M. Nielsen, Mads Chr. Olesen, Samuel Pastva, and Jirí Srba. Extended dependency graphs and efficient distributed fixed-point computation. In *Petri Nets*, volume 10258 of *Lecture Notes in Computer Science*, pages 139–158. Springer, 2017.

[71] Andreas Engelbredt Dalsgaard, Søren Enevoldsen, Kim Guldstrand Larsen, and Jirí Srba. Distributed computation of fixed points on dependency graphs. In *SETTA*, volume 9984 of *Lecture Notes in Computer Science*, pages 197–212, 2016.

[72] Mohamad K. Daryabari, Reza Keypour, and Hessam Golmohamadi. Stochastic energy management of responsive plug-in electric vehicles characterizing parking lot aggregators. *Applied Energy*, 279:115751, 2020.

[73] Alexandre David, Kim G. Larsen, Axel Legay, Marius Mikucionis, and Danny Bøgsted Poulsen. Uppaal SMC tutorial. *Int. J. Softw. Tools Technol. Transf.*, 17(4):397–415, 2015.

[74] Joel D. Day, Mitja Kulczynski, Florin Manea, Dirk Nowotka, and Danny Bøgsted Poulsen. Rule-based word equation solving. In *FormaliSE@ICSE*, pages 87–97. ACM, 2020.

[75] Samvid Dharanikot, Michael Toft Jensen, Sebastian Rom Kristensen, Mathias Sass Michno, Yvonne Anne Pignolet, René Rydhof Hansen, and Stefan Schmid. Breeding unicorns: Developing trustworthy and scalable randomness beacons. In *Blockchain*, pages 99–106. IEEE, 2019.

[76] Samvid Dharanikota, Michael Toft Jensen, Sebastian Rom Kristensen, Mathias Sass Michno, Yvonne-Anne Pignolet, René Rydhof Hansen, and Stefan Schmid. Breeding unicorns: Developing trustworthy and scalable randomness beacons. *PloS ONE*, 15(4), 2020. e0232261.

[77] S. Enevoldsen, K.G.Larsen, and J.Srba. Extended abstract dependency graphs. *STTT: Software Tools for Technology Transfer*, 2021. Accepted for publication.

[78] Soeren Enevoldsen, Mathias Claus Jensen, Kim Guldstrand Larsen, Anders Mariegaard, and Jiri Srba. Verification of multiplayer stochastic games via abstract dependency graphs. In *LOPSTR*, 2020. to appear.

[79] Søren Enevoldsen, Kim G. Larsen, Anders Mariegaard, and Jirí Srba. Dependency graphs with applications to verification. *Int. J. Softw. Tools Technol. Transf.*, 22(5):635–654, 2020.

[80] Søren Enevoldsen, Kim Guldstrand Larsen, and Jirí Srba. Abstract dependency graphs and their application to model checking. In *TACAS (1)*, volume 11427 of *Lecture Notes in Computer Science*, pages 316–333. Springer, 2019.

[81] Søren Enevoldsen, Kim Guldstrand Larsen, and Jirí Srba. Model verification through dependency graphs. In *SPIN*, volume 11636 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2019.

[82] Andreas Berre Eriksen, Chao Huang, Jan Kildebogaard, Harry Lahrmann, Kim Guldstrand Larsen, Marco Muniz, and Jakob Haahr Taankvist. Uppaal stratego for intelligent traffic lights. In *12th ITS European Congress*. ERTICO - ITS Europe, 2017.

[83] Andreas Berre Eriksen, Harry Lahrmann, Kim Guldstrand Larsen, and Jakob Haahr Taankvist. Controlling Signalized Intersections using Machine Learning. In *World Conference on Transport Research-WCTR 2019*, 2019.

[84] Pamela Fleischmann, Mitja Kulczynski, Dirk Nowotka, and Danny Bøgsted Poulsen. On collapsing prefix normal words. In *LATA*, volume 12038 of *Lecture Notes in Computer Science*, pages 412–424. Springer, 2020.

[85] Klaus-Tycho Foerster, Manya Ghobadi, and Stefan Schmid. Characterizing the algorithmic complexity of reconfigurable data center architectures. In *ANCS*, pages 89–96. ACM, 2018.

[86] Klaus-Tycho Foerster, Arne Ludwig, Jan Marcinkowski, and Stefan Schmid. Loop-free route updates for software-defined networks. *IEEE/ACM Trans. Netw.*, 26(1):328–341, 2018.

[87] Klaus-Tycho Foerster, Mahmoud Parham, and Stefan Schmid. A walk in the clouds: Routing through vnfs on bidirected networks. In *ALGOCLOUD*, volume 10739 of *Lecture Notes in Computer Science*, pages 11–26. Springer, 2017.

[88] Klaus-Tycho Foerster, Yvonne Anne Pignolet, Stefan Schmid, and Gilles Trédan. Local fast failover routing with low stretch. *Comput. Commun. Rev.*, 48(1):35–41, 2018.

[89] Peter Fogh, Thomas C. Hald, and Brian Nielsen. Wuppaal: A web-service for the uppaal model-checker. In *CPS Data*, pages 1–6. IEEE Computer Society, 2016.

[90] Louise Foshammer, Kim G. Larsen, Radu Mardare, and Bingtian Xue. Logical characterization and complexity of weighted branching preorders and distances. In *Computation Tools*, IARIA, pages 12–18. Think Mind, 2016.

[91] Louise Foshammer, Kim Guldstrand Larsen, and Anders Mariegaard. Weighted branching simulation distance for parametric weighted kripke structures. In *Cassting/SynCoP*, volume 220 of *EPTCS*, pages 63–75, 2016.

[92] Davide Frazzetto, Bijay Neupane, Torben Bach Pedersen, and Thomas Dyhre Nielsen. Adaptive user-oriented direct load-control of residential flexible devices. In *e-Energy*, pages 1–11. ACM, 2018.

[93] Carlo Fuerst, Maciej Pacut, and Stefan Schmid. Data locality and replica aware virtual cluster embeddings. *Theor. Comput. Sci.*, 697:37–57, 2017.

[94] Carlo Fuerst, Stefan Schmid, P. Lalith Suresh, and Paolo Costa. Kraken: Online and elastic resource reservations for cloud datacenters. *IEEE/ACM Trans. Netw.*, 26(1):422–435, 2018.

[95] Robert Furber. Scott continuity in generalized probabilistic theories. In *QPL*, volume 318 of *EPTCS*, pages 66–84, 2019.

[96] Robert Furber, Dexter Kozen, Kim G. Larsen, Radu Mardare, and Prakash Panangaden. Unrestricted stone duality for markov processes. In *LICS*, pages 1–9. IEEE Computer Society, 2017.

[97] Robert Furber, Radu Mardare, and Matteo Mio. Probabilistic logics based on riesz spaces. *Log. Methods Comput. Sci.*, 16(1), 2020.

[98] Olga Gadyatskaya, René Rydhof Hansen, Kim Guldstrand Larsen, Axel Legay, Mads Chr. Olesen, and Danny Bøgsted Poulsen. Modelling attack-defense trees using timed automata. In *FORMATS*, volume 9884 of *Lecture Notes in Computer Science*, pages 35–50. Springer, 2016.

[99] M. Glavind, N. Christensen, J. Srba, and S. Schmid. Latte: Improving the latency of transiently consistent network update schedules. In *38th International Symposium on Computer Performance, Modeling, Measurements and Evaluation 2020*. ACM, 2020. To appear.

[100] Hessam Golmohamadi. Agricultural demand response aggregators in electricity markets: Structure, challenges and practical solutions- a tutorial for energy experts. *Technology and Economics of Smart Grids and Sustainable Energy*, 5:1–17, 2020.

[101] Hessam Golmohamadi and Amin Asadi. Integration of joint power-heat flexibility of oil refinery industries to uncertain energy markets. *Energies*, 13(18), 2020.

[102] Hessam Golmohamadi and Amin Asadi. A multi-stage stochastic energy management of responsive irrigation pumps in dynamic electricity markets. *Applied Energy*, 265:114804, 2020.

[103] Hessam Golmohamadi, Reza Keypour, and Pouya Mirzazade. Multi-objective co-optimization of power and heat in urban areas considering local air pollution. *Engineering Science and Technology, an International Journal*, 2020.

[104] Thibaut Le Guilly, Michael Nielsen, Thomas Pedersen, Arne Skou, Jesper Kjeldskov, and Mikael B. Skov. User constraints for reliable user-defined smart home scenarios. *J. Reliab. Intell. Environ.*, 2(2):75–91, 2016.

[105] Thibaut Le Guilly, Petur Olsen, Thomas Pedersen, Anders P. Ravn, and Arne Skou. Model checking feature interactions. In Pascal Lorenz, Jorge Cardoso, Leszek A. Maciaszek, and Marten van Sinderen, editors, *Software Technologies*, pages 307–325, Cham, 2016. Springer International Publishing.

[106] Thibaut Le Guilly, Arne Skou, Petur Olsen, Per Printz Madsen, Michele Albano, Luis Lino Ferreira, Luís Miguel Pinho, Keld Pedersen, Miquel Casals, Marcel Macarulla, and Marta Gangolells. Encourageing results on ICT for energy efficient buildings. In *ETFA*, pages 1–8. IEEE, 2016.

[107] Tobias R. Gundersen, Florian Lorber, Ulrik Nyman, and Christian Ovesen. Effortless fault localisation: Conformance testing of real-time systems in ecdar. In *GandALF*, volume 277 of *EPTCS*, pages 147–160, 2018.

[108] Pujie Han, Zhengjun Zhai, Brian Nielsen, and Ulrik Nyman. A compositional approach for schedulability analysis of distributed avionics systems. In *MeTRiD@ETAPS*, volume 272 of *EPTCS*, pages 39–51, 2018.

[109] Pujie Han, Zhengjun Zhai, Brian Nielsen, and Ulrik Nyman. A modeling framework for schedulability analysis of distributed avionics systems. In *MARS/VPT@ETAPS*, volume 268 of *EPTCS*, pages 150–168, 2018.

[110] Pujie Han, Zhengjun Zhai, Brian Nielsen, Ulrik Nyman, and Martin Kristjansen. Schedulability Analysis of Distributed Multicore Avionics Systems with UPPAAL. *Journal of Aerospace Information Systems*, 16, 2019.

[111] Mikkel Hansen, Kim Guldstrand Larsen, Radu Mardare, Mathias Ruggaard Pedersen, and Bingtian Xue. A complete approximation theory for weighted transition systems. In *SETTA*, volume 9984 of *Lecture Notes in Computer Science*, pages 213–228, 2016.

[112] René Rydhof Hansen, Peter Gjøl Jensen, Kim Guldstrand Larsen, Axel Legay, and Danny Bøgsted Poulsen. Quantitative evaluation of attack defense trees using stochastic timed automata. In *GraMSec@CSF*, volume 10744 of *Lecture Notes in Computer Science*, pages 75–90. Springer, 2017.

[113] René Rydhof Hansen, Kim Guldstrand Larsen, Axel Legay, Peter Gjøl Jensen, and Danny Bøgsted Poulsen. ADTLang: A programming language approach to attack defense trees. *International Journal on Software Tools for Technology Transfer*, 2020. To appear.

[114] René Rydhof Hansen, Kim Guldstrand Larsen, Mads Chr. Olesen, and Erik Ramsgaard Wognsen. Formal modelling and analysis of bitflips in ARM assembly code. *Inf. Syst. Frontiers*, 18(5):909–925, 2016.

[115] René Rydhof Hansen and Christian W. Probst. Modelling and analysing overlay networks by ambients with wormholes. In *ICMU*, pages 1–4. IEEE, 2018.

[116] Hamid Hassani, Pieter J.L. Cuijpers, and Reinder J. Bril. Work-in-progress: Layering concerns for the analysis of credit-based shaping in ieee 802.1 tsn. In Luis Almeida, Ramez Daoud, Ahlem Mifdaoui, Frank Golatowski, Hassanein Amer, and Pedro Santos, editors, *16th IEEE International Conference on Factory Communication Systems (WFCS)*, page 4, United States, June 2020. Institute of Electrical and Electronics Engineers. 16th IEEE International Conference on Factory Communication Systems, WFCS 2020 ; Conference date: 27-04-2020 Through 29-04-2020.

[117] David Hausheer, Oliver Hohlfeld, Stefan Schmid, and Guofei Gu. Security and performance of software-defined networks and functions virtualization. *Comput. Networks*, 138:15–17, 2018.

[118] Mu He, Patrick Kalmbach, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. Algorithm-data driven optimization of adaptive communication networks. In *ICNP*, pages 1–6. IEEE Computer Society, 2017.

[119] Juho Hirvonen, Joel Rybicki, Stefan Schmid, and Jukka Suomela. Large cuts with local algorithms on triangle-free graphs. *Electron. J. Comb.*, 24(4):P4.21, 2017.

[120] Oliver Hohlfeld, James Kempf, Martin Reisslein, Stefan Schmid, and Nadir Shah. Guest editorial scalability issues and solutions for software defined networks. *IEEE J. Sel. Areas Commun.*, 36(12):2595–2602, 2018.

[121] Philip I. Holler, Magnus K. Jensen, Hannah Marie K. Lockey, and Michele Albano. Strategising robocup in real time with uppaal stratego. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence (ICAART 2021)*, 2020.

[122] Matthias Hollick, Cristina Nita-Rotaru, Panagiotis Papadimitratos, Adrian Perrig, and Stefan Schmid. Toward a taxonomy and attacker model for secure routing protocols. *Comput. Commun. Rev.*, 47(1):43–48, 2017.

[123] Dmitry Ivanov, Kim G. Larsen, Sibylle Schupp, and Jirí Srba. Analytical solution for long battery lifetime prediction in nonadaptive systems. In *QEST*, volume 11024 of *Lecture Notes in Computer Science*, pages 173–189. Springer, 2018.

[124] Manfred Jaeger, Giorgio Bacci, Giovanni Bacci, Kim Guldstrand Larsen, and Peter Gjøl Jensen. Approximating euclidean by imprecise markov decision processes. In *ISoLA*, volume 12476 of *Lecture Notes in Computer Science*, pages 275–289. Springer, 2020.

[125] Manfred Jaeger, Peter Gjøl Jensen, Kim Guldstrand Larsen, Axel Legay, Sean Sedwards, and Jakob Haahr Taankvist. Teaching stratego to play ball: Optimal synthesis for continuous space mdps. In *ATVA*, volume 11781 of *Lecture Notes in Computer Science*, pages 81–97. Springer, 2019.

[126] Manfred Jaeger, Kim G. Larsen, and Alessandro Tibo. From statistical model checking to run-time monitoring using a Bayesian network approach. In *International Conference on Runtime Verification (RV)*, pages 517–535, 2020.

[127] Manfred Jaeger, Marco Lippi, Giovanni Pellegrini, and Andrea Passerini. Counts-of-counts similarity for prediction and search in relational data. *Data Min. Knowl. Discov.*, 33(5):1254–1297, 2019.

[128] Manfred Jaeger and Oliver Schulte. Inference, learning, and population size: Projectivity for SRL models. *CoRR*, abs/1807.00564, 2018.

[129] Manfred Jaeger and Oliver Schulte. A complete characterization of projectivity for statistical relational models. In *IJCAI*, pages 4283–4290. ijcai.org, 2020.

[130] Amirhossein Jangdoost, Reza Keypour, and Hessam Golmohamadi. Optimization of distribution network reconfiguration by a novel rca integrated with genetic algorithm. *Energy Systems*, July 2020.

[131] Erkki Jantunen, Giovanni Di Orio, Csaba Hegedus, Pal Varga, Istvan Moldovan, Felix Larrinaga, Martin Becker, Michele Albano, and Pedro Maló. Maintenance 4.0 world of integrated information. In K. Popplewell, KD. Thoben, T. Knothe, and R. Poler, editors, *Enterprise Interoperability VIII. Proceedings of the I-ESA Conferences*, pages 67–78. Springer, Cham, Cham, Switzerland, 2019.

[132] Cyrille Jégourel, Kim G. Larsen, Axel Legay, Marius Mikučionis, Danny Bøgsted Poulsen, and Sean Sedwards. Importance sampling for stochastic timed automata. In *SETTA*, volume 9984 of *Lecture Notes in Computer Science*, pages 163–178, 2016.

[133] Finn Verner Jensen and Thomas Dyhre Nielsen. *Bayesian Graphical Models*, pages 1–9. Wiley StatsRef: Statistics Reference Online. Wiley, 2016.

[134] Jesper Stenbjerg Jensen, Troels Beck Krøgh, Jonas Sand Madsen, Stefan Schmid, Jirí Srba, and Marc Tom Thorgersen. P-rex: fast verification of MPLS networks with multiple link failures. In *CoNEXT*, pages 217–227. ACM, 2018.

[135] Jonas Finnemann Jensen, Kim Guldstrand Larsen, Jirí Srba, and Lars Kaerlund Oestergaard. Efficient model-checking of weighted CTL with upper-bound constraints. *Int. J. Softw. Tools Technol. Transf.*, 18(4):409–426, 2016.

[136] Lasse S. Jensen, Isabella Kaufmann, Kim G. Larsen, Søren M. Nielsen, and Jirí Srba. Model checking and synthesis for branching multi-weighted logics. *J. Log. Algebraic Methods Program.*, 105:28–46, 2019.

[137] Mathias Claus Jensen, Kim Guldstrand Larsen, and Radu Mardare. Weighted branching systems: Behavioural equivalence, behavioural distance, and their logical characterisations. In *FORMATS*, volume 11022 of *Lecture Notes in Computer Science*, pages 145–161. Springer, 2018.

[138] Mathias Claus Jensen, Anders Mariegaard, and Kim Guldstrand Larsen. Symbolic model checking of weighted PCTL using dependency graphs. In *NFM*, volume 11460 of *Lecture Notes in Computer Science*, pages 298–315. Springer, 2019.

[139] Peter Gjøl Jensen, Kenneth Yrke Jørgensen, Kim G. Larsen, Marius Mikucionis, Marco Muñiz, and Danny Bøgsted Poulsen. Fluid model-checking in UPPAAL for covid-19. In *ISoLA (1)*, volume 12476 of *Lecture Notes in Computer Science*, pages 385–403. Springer, 2020.

[140] Peter Gjøl Jensen, Kim Guldstrand Larsen, and Jirí Srba. Real-time strategy synthesis for timed-arc petri net games via discretization. In *SPIN*, volume 9641 of *Lecture Notes in Computer Science*, pages 129–146. Springer, 2016.

[141] Peter Gjøl Jensen, Kim Guldstrand Larsen, and Jirí Srba. Ptrie: Data structure for compressing and storing sets via prefix sharing. In *ICTAC*, volume 10580 of *Lecture Notes in Computer Science*, pages 248–265. Springer, 2017.

[142] Peter Gjøl Jensen, Kim Guldstrand Larsen, and Jirí Srba. Discrete and continuous strategies for timed-arc petri net games. *Int. J. Softw. Tools Technol. Transf.*, 20(5):529–546, 2018.

[143] P.G. Jensen, D. Kristiansen, S. Schmid, M.K. Schou, B.C. Schrenk, and J. Srba. A fast and quantitative what-if analysis tool for mpls networks. In *CoNEXT 2020 - International Conference on emerging Networking EXperiments and Technologies*. ACM, 2020. To appear.

[144] Tobias Skovgaard Jepsen, Christian S. Jensen, and Thomas Dyhre Nielsen. Graph convolutional networks for road networks. In *SIGSPATIAL/GIS*, pages 460–463. ACM, 2019.

[145] Tobias Skovgaard Jepsen, Christian S. Jensen, and Thomas Dyhre Nielsen. Relational fusion networks: Graph convolutional networks for road networks. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–12, 2020. to appear.

[146] Tobias Skovgaard Jepsen, Christian S. Jensen, Thomas Dyhre Nielsen, and Kristian Torp. On network embedding for machine learning on road networks: A case study on the danish road network. In *BigData*, pages 3422–3431. IEEE, 2018.

[147] Patrick Kalmbach, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. Poster abstract: Themis: A data-driven approach to bot detection. In *INFOCOM Workshops*, pages 1–2. IEEE, 2018.

[148] Patrick Kalmbach, Lion Gleiter, Johannes Zerwas, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. Modeling ip-to-ip communication using the weighted stochastic block model. In *SIGCOMM Posters and Demos*, pages 48–50. ACM, 2018.

[149] Patrick Kalmbach, Johannes Zerwas, Péter Babarczi, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. Empowering self-driving networks. In *SelfDN@SIGCOMM*, pages 8–14. ACM, 2018.

[150] Shyam Lal Karra, Kim Guldstrand Larsen, Florian Lorber, and Jirí Srba. Safe and time-optimal control for railway games. In *RSSRail*, volume 11495 of *Lecture Notes in Computer Science*, pages 106–122. Springer, 2019.

[151] Shyam Lal Karra, Kim Guldstrand Larsen, Marco Muñiz, and Jirí Srba. On-the-fly synthesis for strictly alternating games. In *Petri Nets*, volume 12152 of *Lecture Notes in Computer Science*, pages 109–128. Springer, 2020.

[152] Isabella Kaufmann, Kim Guldstrand Larsen, and Jirí Srba. Synthesis for multi-weighted games with branching-time winning conditions. In *Petri Nets*, volume 12152 of *Lecture Notes in Computer Science*, pages 46–66. Springer, 2020.

[153] Wolfgang Kellerer, Raouf Boutaba, Prosper Chemouil, Rafael Pasquini, Giovanni Schembra, Stefan Schmid, Sandra Scott-Hayward, and Kohei Shiomoto. Guest editors' introduction: Special section on novel techniques for managing softwarized networks. *IEEE Trans. Netw. Serv. Manag.*, 15(4):1192–1196, 2018.

[154] Jin Hyun Kim, Axel Legay, Louis-Marie Traonouez, Abdeldjalil Boudjadar, Ulrik Nyman, Kim G. Larsen, Insup Lee, and Jin-Young Choi. Optimizing the resource requirements of hierarchical scheduling systems. *SIGBED Rev.*, 13(3):41–48, 2016.

[155] Andrej Kiviriga, Kim Guldstrand Larsen, and Ulrik Nyman. Randomized refinement checking of timed I/O automata. In Jun Pang and Lijun Zhang, editors, *Dependable Software Engineering. Theories, Tools, and Applications - 6th International Symposium, SETTA 2020, Guangzhou, China, November 24-27, 2020, Proceedings*, volume 12153 of *Lecture Notes in Computer Science*, pages 70–88. Springer, 2020.

[156] Arne Köhn, Julia Wichlacz, Álvaro Torralba, Daniel Höller, Jörg Hoffmann, and Alexander Koller. Generating instructions at different levels of abstraction. In *Proceedings of the 28th International Conference on Computational Linguistics (Coling'20)*, 2020. To appear.

[157] Robert Krösche, Kashyap Thimmaraju, Liron Schiff, and Stefan Schmid. I DPID it my way! A covert timing channel in software-defined networks. In *Networking*, pages 2017–225. IFIP, 2018.

[158] Mitja Kulczynski, Florin Manea, Dirk Nowotka, and Danny Bøgsted Poulsen. The power of string solving: Simplicity of comparison. In *AST@ICSE*, pages 85–88. ACM, 2020.

[159] Kim G. Larsen. Synthesis of safe, optimal and compact strategies for stochastic hybrid games (invited paper). In *CONCUR*, volume 140 of *LIPIcs*, pages 2:1–2:5. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[160] Kim G. Larsen, Simon Laursen, and Martin Zimmermann. Limit your consumption! finding bounds in average-energy games. In *QAPL*, volume 227 of *EPTCS*, pages 1–14, 2016.

[161] Kim G. Larsen and Axel Legay. On the power of statistical model checking. In *ISoLA (2)*, volume 9953 of *Lecture Notes in Computer Science*, pages 843–862, 2016.

[162] Kim G. Larsen and Axel Legay. Statistical model checking: Past, present, and future. In *ISoLA (1)*, volume 9952 of *Lecture Notes in Computer Science*, pages 3–15, 2016.

[163] Kim G. Larsen and Axel Legay. 30 years of statistical model checking. In *ISoLA (1)*, volume 12476 of *Lecture Notes in Computer Science*, pages 325–330. Springer, 2020.

[164] Kim G. Larsen, Axel Legay, Marius Mikučionis, Brian Nielsen, and Ulrik Nyman. Compositional testing of real-time systems. In *ModelEd, TestEd, TrustEd*, volume 10500 of *Lecture Notes in Computer Science*, pages 107–124. Springer, 2017.

[165] Kim G. Larsen, Florian Lorber, and Brian Nielsen. 20 years of UPPAAL enabled industrial model-based validation and beyond. In *ISoLA (4)*, volume 11247 of *Lecture Notes in Computer Science*, pages 212–229. Springer, 2018.

[166] Kim G. Larsen, Florian Lorber, Brian Nielsen, and Ulrik Nyman. Mutation-based test-case generation with ecdar. In *ICST Workshops*, pages 319–328. IEEE Computer Society, 2017.

[167] Kim G. Larsen, Radu Mardare, and Bingtian Xue. Probabilistic mu-calculus: Decidability and complete axiomatization. In *FSTTCS*, volume 65 of *LIPIcs*, pages 25:1–25:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[168] Kim G. Larsen, Radu Mardare, and Bingtian Xue. On decidability of recursive weighted logics. *Soft Comput.*, 22(4):1085–1102, 2018.

[169] Kim G. Larsen, Marius Mikucionis, Marco Muñiz, and Jirí Srba. Urgent partial order reduction for extended timed automata. In Dang Van Hung and

Oleg Sokolsky, editors, *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*, volume 12302 of *Lecture Notes in Computer Science*, pages 179–195. Springer, 2020.

[170] Kim G. Larsen, Marius Mikučionis, Marco Muñiz, Jirí Srba, and Jakob Haahr Taankvist. Online and compositional learning of controllers with application to floor heating. In *TACAS*, volume 9636 of *Lecture Notes in Computer Science*, pages 244–259. Springer, 2016.

[171] Kim G. Larsen, Doron Peled, and Sean Sedwards. Memory-efficient tactics for randomized LTL model checking. In *VSTTE*, volume 10712 of *Lecture Notes in Computer Science*, pages 152–169. Springer, 2017.

[172] Kim Guldstrand Larsen. Dependable and optimal cyber-physical systems. In *SOFSEM*, volume 10139 of *Lecture Notes in Computer Science*, pages 3–10. Springer, 2017.

[173] Kim Guldstrand Larsen. Timed and untimed energy games. In *CIAA*, volume 10329 of *Lecture Notes in Computer Science*, pages xiii–xvi. Springer, 2017.

[174] Kim Guldstrand Larsen. Validation, synthesis and optimization for cyber-physical systems. In *TACAS (1)*, volume 10205 of *Lecture Notes in Computer Science*, pages 3–20, 2017.

[175] Kim Guldstrand Larsen, Adrien Le Coënt, Marius Mikučionis, and Jakob Haahr Taankvist. Guaranteed control synthesis for continuous systems in uppaal tiga. In *CyPhy/WESE*, volume 11615 of *Lecture Notes in Computer Science*, pages 113–133. Springer, 2018.

[176] Kim Guldstrand Larsen and Axel Legay. Statistical model checking the 2018 edition! In *ISoLA (2)*, volume 11245 of *Lecture Notes in Computer Science*, pages 261–270. Springer, 2018.

[177] Kim Guldstrand Larsen, Florian Lorber, and Brian Nielsen. 20 years of real real time model validation. In *FM*, volume 10951 of *Lecture Notes in Computer Science*, pages 22–36. Springer, 2018.

[178] Thi Thieu Hoa Le, Roberto Passerone, Uli Fahrenberg, and Axel Legay. Contract-based requirement modularization via synthesis of correct decompositions. *ACM Trans. Embed. Comput. Syst.*, 15(2):33:1–33:26, 2016.

[179] Thi Thieu Hoa Le, Roberto Passerone, Ulrich Fahrenberg, and Axel Legay. A tag contract framework for modeling heterogeneous systems. *Sci. Comput. Program.*, 115-116:225–246, 2016.

[180] Axel Legay, Dirk Nowotka, and Danny Bøgsted Poulsen. Automatic verification of LLVM code. *CoRR*, abs/2006.02670, 2020.

[181] Florian Lorber, Kim G. Larsen, and Brian Nielsen. Model-based mutation testing of real-time systems via model checking. In *ICST Workshops*, pages 59–68. IEEE Computer Society, 2018.

[182] Kasper Søe Luckow, Bent Thomsen, and Stephan Korsholm. Hvmtp: A time predictable and portable java virtual machine for hard real-time embedded systems. *Concurrency and Computation: Practice & Experience*, 29(22), 2017.

[183] Arne Ludwig, Szymon Dudycz, Matthias Rost, and Stefan Schmid. Transiently secure network updates. In *SIGMETRICS*, pages 273–284. ACM, 2016.

[184] Arne Ludwig, Szymon Dudycz, Matthias Rost, and Stefan Schmid. Transiently policy-compliant network updates. *IEEE/ACM Trans. Netw.*, 26(6):2569–2582, 2018.

[185] Tamás Lukovszki, Matthias Rost, and Stefan Schmid. It's a match!: Near-optimal and incremental middlebox deployment. *Comput. Commun. Rev.*, 46(1):30–36, 2016.

[186] Tamás Lukovszki, Matthias Rost, and Stefan Schmid. Approximate and incremental network function placement. *J. Parallel Distributed Comput.*, 120:159–169, 2018.

[187] Manuel Luque, Thomas D. Nielsen, and Finn Verner Jensen. Anytime decision making based on unconstrained influence diagrams. *Int. J. Intell. Syst.*, 31(4):379–398, 2016.

[188] Anders L. Madsen, Frank Jensen, Antonio Salmerón, Helge Langseth, and Thomas D. Nielsen. A parallel algorithm for bayesian network structure learning from large data sets. *Knowl. Based Syst.*, 117:46–55, 2017.

[189] Anders L. Madsen, Kristian G. Olesen, Heidi Lynge Løvschall, Nicolaj Søndberg-Jeppesen, Frank Jensen, Morten Lindblad, Mads Lause Mogensen,

and Trine Søby Christensen. Prediction of high risk of deviations in home care deliveries. In *10th Int. Conf. on Probabilistic Graphical Models (PGM)*, 2020.

[190] Anders L. Madsen, Kristian G. Olesen, Jørn Munkhof Møller, Nicolaj Søndberg-Jeppesen, Frank Jensen, Thomas Mulvad Larsen, Per Henriksen, Morten Lindblad, and Trine Søby Christensen. A software system for predicting patient flow at the emergency department of aalborg university hospital. In *10th Int. Conf. on Probabilistic Graphical Models (PGM)*, 2020.

[191] Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. Quantitative algebraic reasoning. In *LICS*, pages 700–709. ACM, 2016.

[192] Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. On the axiomatizability of quantitative algebras. In *LICS*, pages 1–12. IEEE Computer Society, 2017.

[193] Radu Mardare, Prakash Panangaden, and Gordon D. Plotkin. Free complete wasserstein algebras. *Log. Methods Comput. Sci.*, 14(3), 2018.

[194] Anders Mariegaard and Kim Guldstrand Larsen. Symbolic dependency graphs for $$\text {PCTL}^{>}_{\le }$$ model-checking. In *FORMATS*, volume 10419 of *Lecture Notes in Computer Science*, pages 153–169. Springer, 2017.

[195] Andrés R. Masegosa, Ana M. Martínez, Helge Langseth, Thomas D. Nielsen, Antonio Salmerón, Darío Ramos-López, and Anders L. Madsen. d-vmp: Distributed variational message passing. In *Probabilistic Graphical Models*, volume 52 of *JMLR Workshop and Conference Proceedings*, pages 321–332. JMLR.org, 2016.

[196] Andrés R. Masegosa, Ana M. Martínez, Helge Langseth, Thomas D. Nielsen, Antonio Salmerón, Darío Ramos-López, and Anders L. Madsen. Scaling up bayesian variational inference using distributed computing clusters. *Int. J. Approx. Reason.*, 88:435–451, 2017.

[197] Andrés R. Masegosa, Ana M. Martínez, Darío Ramos-López, Rafael Cabañas, Antonio Salmerón, Helge Langseth, Thomas D. Nielsen, and Anders L. Madsen. AMIDST: A java toolbox for scalable probabilistic machine learning. *Knowl. Based Syst.*, 163:595–597, 2019.

[198] Andrés R. Masegosa, Ana M. Martínez, Darío Ramos-López, Helge Langseth, Thomas D. Nielsen, and Antonio Salmerón. Analyzing concept drift: A case study in the financial sector. *Intell. Data Anal.*, 24(3):665–688, 2020.

[199] Andrés R. Masegosa, Thomas D. Nielsen, Helge Langseth, Darío Ramos-López, Antonio Salmerón, and Anders L. Madsen. Bayesian models of data streams with hierarchical power priors. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 2334–2343. PMLR, 2017.

[200] Andres R. Masegosa, Darío Ramos-López, Antonio Salmerón Cerdán, Helge Langseth, and Thomas Dyhre Nielsen. Variational inference over nonstationary data streams for exponential family models. *Mathematics*, 8(11), 2020.

[201] Matteo Mio, Robert Furber, and Radu Mardare. Riesz modal logic for markov processes. In *LICS*, pages 1–12. IEEE Computer Society, 2017.

[202] Balázs Németh, Mark Szalay, Janos Doka, Matthias Rost, Stefan Schmid, László Toka, and Balázs Sonkoly. Fast and efficient network service embedding method with adaptive offloading to the edge. In *INFOCOM Workshops*, pages 178–183. IEEE, 2018.

[203] Brian Nielsen, Michele Albano, Thomas Pedersen, Anders Hødholdt, and Christian Østergaard Laursen. Battery lifetime estimation and optimization for an urban water system wireless sensor network. 2020. Accepted for IWA 2020 poster presentation - to be presented IWA 2021 (Corona postponed); null ; Conference date: 18-10-2020 Through 23-10-2020.

[204] Flemming Nielson, René Rydhof Hansen, and Hanne Riis Nielson. Adaptive security policies. In *Proceedings of the 9th International Symposium on Leveraging Applications of Formal Methods (ISoLA 2020)*, volume 12477 of *Lecture Notes in Computer Science*, pages 280–294. Springer, 2020.

[205] Gilles Nies, Marvin Stenger, Jan Krčál, Holger Hermanns, Morten Bisgaard, David Gerhardt, Boudewijn Haverkort, Marijn Jongerden, Kim G. Larsen, and Erik R. Wognsen. Mastering operational limitations of leo satellites – the gomx-3 approach. *Acta Astronautica*, 151:726–735, 2018.

[206] Emil J. Njor, Florian Lorber, Nikolaj I. Schmidt, and Sebastian R. Petersen. Conformance testing in UPPAAL: A diabolic approach. In *ICST Workshops*, pages 35–42. IEEE, 2020.

[207] Mateja Novak, Tomislav Dragicevic, Frede Blaabjerg, and Ulrik Mathias Nyman. Analytical performance verification of fcs-mpc applied to power electronic converters: A model checking approach. In *Proceedings of 2017 IEEE*

18th Workshop on Control and Modeling for Power Electronics (COMPEL), IEEE Workshop on Control and Modeling for Power Electronics (COMPEL). IEEE Press, July 2017. 18th IEEE Workshop on Control and Modeling for Power Electronics, COMPEL 2017 ; Conference date: 09-07-2017 Through 12-07-2017.

[208] Mateja Novak, Ulrik Nyman, Tomislav Dragicevic, and Frede Blaabjerg. Analytical design and performance validation of finite set MPC regulated power converters. *IEEE Trans. Ind. Electron.*, 66(3):2004–2014, 2019.

[209] Mateja Novak, Ulrik Nyman, Tomislav Dragicevic, and Frede Blaabjerg. Statistical model checking for finite-set model predictive control converters: A tutorial on modeling and performance verification. *I E E E Industrial Electronics Magazine*, 13(3):6–15, September 2019.

[210] Mateja Novak, Ulrik Mathias Nyman, Tomislav Dragicevic, and Frede Blaabjerg. Statistical performance verification of fcs-mpc applied to three level neutral point clamped converter. In *Proceedings of 2018 20th European Conference on Power Electronics and Applications (EPE'18 ECCE Europe)*, pages 1–10. IEEE Press, September 2018. 20th European Conference on Power Electronics and Applications, EPE 2018 ECCE Europe ; Conference date: 17-09-2018 Through 21-09-2018.

[211] Adrian Ogierman, Andréa W. Richa, Christian Scheideler, Stefan Schmid, and Jin Zhang. Sade: competitive MAC under adversarial SINR. *Distributed Comput.*, 31(3):241–254, 2018.

[212] Yoann Padioleau, Julia L. Lawall, René Rydhof Hansen, and Gilles Muller. Documenting and automating collateral evolutions in Linux device drivers. In Joseph S. Sventek and Steven Hand, editors, *Proceedings of the 2008 European Conference on Computer Systems (EuroSys 2008)*, pages 247–260, 2008.

[213] Thomas Kobber Panum, Kaspar Hageman, René Rydhof Hansen, and Jens Myrup Pedersen. Towards adversarial phishing detection. In *CSET @ USENIX Security Symposium*. USENIX Association, 2020.

[214] Thomas Kobber Panum, Kaspar Hageman, Jens Myrup Pedersen, and René Rydhof Hansen. Haaukins: A highly accessible and automated virtualization platform for security education. In *ICALT*, pages 236–238. IEEE, 2019.

[215] Thomas Kobber Panum, René Rydhof Hansen, and Jens Myrup Pedersen. Kraaler: A user-perspective web crawler. In *TMA*, pages 153–160. IEEE, 2019.

[216] Roberto Passerone, Daniela Cancila, Michele Albano, Sebti Mouelhi, Sándor Plósz, Erkki Jantunen, Anna Ryabokon, Mohamed-Emine Laarouchi, Csaba Hegedüs, and Pál Varga. A methodology for the design of safety-compliant and secure communication of autonomous vehicles. *IEEE Access*, 7:125022–125037, 2019.

[217] Debasish Pattanayak, Kaushik Mondal, Partha Sarathi Mandal, and Stefan Schmid. Convergence of even simpler robots without position information. In *NETYS*, volume 10299 of *Lecture Notes in Computer Science*, pages 69–85, 2017.

[218] Debasish Pattanayak, H. Ramesh, Partha Sarathi Mandal, and Stefan Schmid. Evacuating two robots from two unknown exits on the perimeter of a disk with wireless communication. In *ICDCN*, pages 20:1–20:4. ACM, 2018.

[219] Mathias Ruggaard Pedersen, Giorgio Bacci, Kim Guldstrand Larsen, and Radu Mardare. A hemimetric extension of simulation for semi-markov decision processes. In *QEST*, volume 11024 of *Lecture Notes in Computer Science*, pages 339–355. Springer, 2018.

[220] Mathias Ruggaard Pedersen, Nathanaël Fijalkow, Giorgio Bacci, Kim G. Larsen, and Radu Mardare. Timed comparisons of semi-markov processes. In *LATA*, volume 10792 of *Lecture Notes in Computer Science*, pages 271–283. Springer, 2018.

[221] Thomas Pedersen, Michele Albano, and Brian Nielsen. Reengineering the lifecycle of arrowhead applications: from skeletons to the client library. In *IECON*, pages 5519–5524. IEEE, 2019.

[222] Bruna Soares Peres, Olga Goussevskaia, Stefan Schmid, and Chen Avin. Brief announcement: Distributed splaynets. In *DISC*, volume 91 of *LIPIcs*, pages 58:1–58:3. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[223] Yvonne Anne Pignolet, Stefan Schmid, and Gilles Trédan. Load-optimal local fast rerouting for resilient networks. In *DSN*, pages 345–356. IEEE Computer Society, 2017.

[224] Yvonne Anne Pignolet, Stefan Schmid, and Gilles Trédan. Tomographic node placement strategies and the impact of the routing model. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(2):42:1–42:23, 2017.

[225] Yvonne Anne Pignolet, Stefan Schmid, and Gilles Trédan. Tomographic node placement strategies and the impact of the routing model. In *SIGMETRICS (Abstracts)*, pages 75–77. ACM, 2018.

[226] Christian W. Probst, Florian Kammüller, and René Rydhof Hansen. Formal modelling and analysis of socio-technical systems. In *Semantics, Logics, and Calculi*, volume 9560 of *Lecture Notes in Computer Science*, pages 54–73. Springer, 2016.

[227] Darío Ramos-López, Andrés R. Masegosa, Ana M. Martínez, Antonio Salmerón, Thomas D. Nielsen, Helge Langseth, and Anders L. Madsen. MAP inference in dynamic hybrid bayesian networks. *Prog. Artif. Intell.*, 6(2):133–144, 2017.

[228] Darío Ramos-López, Andrés R. Masegosa, Antonio Salmerón, Rafael Rumí, Helge Langseth, Thomas D. Nielsen, and Anders L. Madsen. Scalable importance sampling estimation of gaussian mixture posteriors in bayesian networks. *Int. J. Approx. Reason.*, 100:115–134, 2018.

[229] Darío Ramos-López, Antonio Salmerón, Rafael Rumí, Ana M. Martínez, Thomas D. Nielsen, Andrés R. Masegosa, Helge Langseth, and Anders L. Madsen. Scalable MAP inference in bayesian networks based on a map-reduce approach. In *Probabilistic Graphical Models*, volume 52 of *JMLR Workshop and Conference Proceedings*, pages 415–425. JMLR.org, 2016.

[230] Anders P. Ravn, Bent Thomsen, Kasper Søe Luckow, Lone Leth Thomsen, and Thomas Bøgholm. Firm deadline checking of safety-critical java applications with statistical model checking. In Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfsdóttir, Axel Legay, and Radu Mardare, editors, *Models, Algorithms, Logics and Tools*, volume 10460 of *Lecture Notes in Computer Science*, pages 269–288, Germany, 2017. Springer. Essays Dedicated to Kim Guldstrand Larsen¡br/¿on the Occasion of His 60th Birthday ; Conference date: 25-07-2017.

[231] Matthias Rost and Stefan Schmid. Charting the complexity landscape of virtual network embeddings. In *Networking*, pages 55–63. IFIP, 2018.

[232] Matthias Rost and Stefan Schmid. Virtual network embedding approximations: Leveraging randomized rounding. In *Networking*, pages 370–378. IFIP, 2018.

[233] Antonio Salmerón, Anders L. Madsen, Frank Jensen, Helge Langseth, Thomas D. Nielsen, Darío Ramos-López, Ana M. Martínez, and Andrés R. Masegosa. Parallel filter-based feature selection based on balanced incomplete block designs. In *ECAI*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 743–750. IOS Press, 2016.

[234] Antonio Salmerón, Rafael Rumí, Helge Langseth, Thomas D. Nielsen, and Anders L. Madsen. A review of inference algorithms for hybrid bayesian networks. *J. Artif. Intell. Res.*, 62:799–828, 2018.

[235] Liron Schiff and Stefan Schmid. PRI: privacy preserving inspection of encrypted network traffic. In *IEEE Symposium on Security and Privacy Workshops*, pages 296–303. IEEE Computer Society, 2016.

[236] Liron Schiff and Stefan Schmid. Study the past if you would define the future: Implementing secure multi-party SDN updates. In *SwSTE*, pages 111–116. IEEE, 2016.

[237] Liron Schiff, Stefan Schmid, and Petr Kuznetsov. In-band synchronization for distributed SDN control planes. *Comput. Commun. Rev.*, 46(1):37–43, 2016.

[238] Liron Schiff, Ofri Ziv, Manfred Jaeger, and Stefan Schmid. Netslicer: Automated and traffic-pattern based application clustering in datacenters. In *Big-DAMA@SIGCOMM*, pages 21–26. ACM, 2018.

[239] Henrik Schiøler, Luminita Totu, Jan Dimon, Kim Guldstrand Larsen, and Jakob Haahr Taankvist. Time optimal robust fleet management of micro UAV through timed games formulation. In *CCTA*, pages 146–152. IEEE, 2018.

[240] S. Schmid and J. Srba. A method for fast verification of prefix-manipulating computer networks, 2017. Patent application PA 201700672. Filed by Aalborg University, November 24th, 2017.

[241] Stefan Schmid. From self-stabilization to self-optimization: Principles of distributed network design. In *PODC*, page 493. ACM, 2018.

[242] Stefan Schmid and Jirí Srba. Polynomial-time what-if analysis for prefix-manipulating MPLS networks. In *INFOCOM*, pages 1799–1807. IEEE, 2018.

[243] Martin Schoeberl, Andreas Engelbredt Dalsgaard, René Rydhof Hansen, Stephan Erbs Korsholm, Anders P. Ravn, Juan Ricardo Rios Rivas, Tórur Biskopstø Strøm, Hans Søndergaard, Andy J. Wellings, and Shuai Zhao. Safety-critical java for embedded systems. *Concurr. Comput. Pract. Exp.*, 29(22), 2017.

[244] Julius Schulz-Zander, Carlos Mayer, Bogdan Ciobotaru, Raphael Lisicki, Stefan Schmid, and Anja Feldmann. Unified programmability of virtualized network functions and software-defined wireless networks. *IEEE Trans. Netw. Serv. Manag.*, 14(4):1046–1060, 2017.

[245] Roshan Sedar, Michael Borokhovich, Marco Chiesa, Gianni Antichi, and Stefan Schmid. Supporting emerging applications with low-latency failover in P4. In *NEAT@SIGCOMM*, pages 52–57. ACM, 2018.

[246] Bhargava Shastry, Markus Leutner, Tobias Fiebig, Kashyap Thimmaraju, Fabian Yamaguchi, Konrad Rieck, Stefan Schmid, Jean-Pierre Seifert, and Anja Feldmann. Static program analysis as a fuzzing aid. In *RAID*, volume 10453 of *Lecture Notes in Computer Science*, pages 26–47. Springer, 2017.

[247] Gayathri Sugumar, Rajasekar Selvamuthukumaran, Tomislav Dragicevic, Ulrik Nyman, Kim G. Larsen, and Frede Blaabjerg. Formal validation of supervisory energy management systems for microgrids. In *IECON*, pages 1154–1159. IEEE, 2017.

[248] Martin Tappler, Bernhard K. Aichernig, Giovanni Bacci, Maria Eichlseder, and Kim G. Larsen. L^*-based learning of markov decision processes. In *FM*, volume 11800 of *Lecture Notes in Computer Science*, pages 651–669. Springer, 2019.

[249] Kashyap Thimmaraju, Gábor Rétvári, and Stefan Schmid. Virtual network isolation: Are we there yet? In *SecSoN@SIGCOMM*, pages 1–7. ACM, 2018.

[250] Kashyap Thimmaraju, Liron Schiff, and Stefan Schmid. Outsmarting network security with SDN teleportation. In *EuroS&P*, pages 563–578. IEEE, 2017.

[251] Kashyap Thimmaraju, Bhargava Shastry, Tobias Fiebig, Felicitas Hetzelt, Jean-Pierre Seifert, Anja Feldmann, and Stefan Schmid. The vamp attack: Taking control of cloud systems via the unified packet parser. In *CCSW@CCS*, pages 11–15. ACM, 2017.

[252] Kashyap Thimmaraju, Bhargava Shastry, Tobias Fiebig, Felicitas Hetzelt, Jean-Pierre Seifert, Anja Feldmann, and Stefan Schmid. Taking control of sdn-based cloud systems via the data plane. In *SOSR*, pages 1:1–1:15. ACM, 2018.

[253] Alessandro Tibo, Paolo Frasconi, and Manfred Jaeger. A network architecture for multi-multi-instance learning. In *ECML/PKDD (1)*, volume 10534 of *Lecture Notes in Computer Science*, pages 737–752. Springer, 2017.

[254] Alessandro Tibo, Manfred Jaeger, and Paolo Frasconi. Learning and interpreting multi-multi-instance learning networks. *Journal of Machine Learning Research*, 21:1–60, 2018.

[255] Alessandro Tibo, Thomas Dyhre Nielsen, Manfred Jaeger, Malthe Ahm, and Peter Rasch. Automatic anomaly detection for sewage network sensors. 2020.

[256] Nicolas Tran, Patrick Speicher, Robert Künnemann, Michael Backes, Álvaro Torralba, and Jörg Hoffmann. Planning in the browser. In *System Demonstration at the 30th International Conference on Automated Planning and Scheduling (ICAPS'20)*, 2020. To appear.

[257] Erik Ramsgaard Wognsen, René Rydhof Hansen, Kim Guldstrand Larsen, and Peter Koch. Energy-aware scheduling of FIR filter structures using a timed automata model. In *DDECS*, pages 163–168. IEEE, 2016.

[258] Johannes Zerwas, Patrick Kalmbach, Carlo Fuerst, Arne Ludwig, Andreas Blenk, Wolfgang Kellerer, and Stefan Schmid. AHAB: data-driven virtual cluster hunting. In *Networking*, pages 379–387. IFIP, 2018.

[259] Zhengkui Zhang, Brian Nielsen, and Kim G. Larsen. Distributed algorithms for time optimal reachability analysis. In *FORMATS*, volume 9884 of *Lecture Notes in Computer Science*, pages 157–173. Springer, 2016.

[260] Zhengkui Zhang, Brian Nielsen, and Kim G. Larsen. Time optimal reachability analysis using swarm verification. In *SAC*, pages 1634–1640. ACM, 2016.

[261] Zhengkui Zhang, Brian Nielsen, Kim Guldstrand Larsen, Gilles Nies, Marvin Stenger, and Holger Hermanns. Pareto optimal reachability analysis for simple priced timed automata. In *ICFEM*, volume 10610 of *Lecture Notes in Computer Science*, pages 481–495. Springer, 2017.

[262] Jiaqi Zheng, Guihai Chen, Stefan Schmid, Haipeng Dai, and Jie Wu. Chronus: Consistent data plane updates in timed sdns. In *ICDCS*, pages 319–327. IEEE Computer Society, 2017.

[263] Jiaqi Zheng, Guihai Chen, Stefan Schmid, Haipeng Dai, Jie Wu, and Qiang Ni. Scheduling congestion- and loop-free network update in timed sdns. *IEEE J. Sel. Areas Commun.*, 35(11):2542–2552, 2017.

[264] Jiaqi Zheng, Bo Li, Chen Tian, Klaus-Tycho Foerster, Stefan Schmid, Guihai Chen, and Jie Wux. Scheduling congestion-free updates of multiple flows with chronicle in timed sdns. In *ICDCS*, pages 12–21. IEEE Computer Society, 2018.

5 Human-Centered Computing

5.1 Executive Summary

The research group Human-Centred Computing focuses on design, development, and application of interactive computing technologies. The research is empirical and solution-oriented and pursued in relation to the international research communities of systems development and human-computer interaction.

The group has successfully maintained its quality and impact of publications and improved its internationalization. Further, the group has worked on funding strategies and has attracted grants and become parts of consortia. The group has reorganized and changed name during the past period, and group has grown significantly over the past few years both for academic staff as well as PhD students.

5.2 Profile

The research group Human-Centred Computing (HCC) focuses on design, development, and application of interactive computing technologies. The research group has a long tradition for empirical and solution-oriented research with a focus on processes, products, and theory. The process-oriented interest contributes to methodological knowledge of agile development, design, evaluation, and innovation of interactive systems. The product-oriented interest contributes to knowledge on the specifics of interactive systems, their applications in work, leisure, organizations, teams as well as other social arrangements. The theoretical knowledge interest is on innovation, design, evaluation, and the application of interactive systems in practice. This focus is pursued in relation to the international research communities of human-computer interaction, information systems, and their intersections.

The HCC group has a strong element of empirical research involving humans, whether they are individual users or developers of software applications, groups

of users and teams of developers, or the use and development in a larger social or organisational context. One part of the empirical research is experimental in the form of design, construction, and study of interactive prototypes to explore and solve specific problems. Another part is action-oriented and directed at digitalisation and software development, and it is often in close collaboration with companies in improving their practices.

The HCC group conducts research activities within three main areas: interaction design, user studies, and design and development.

5.2.1 Interaction Design

The interaction design research area is primarily design-oriented and is concerned with the design of probes, prototypes, and products that can be both digital and physical, and specifically with designing interactive products to support the way people communicate and carry out activities in their everyday and working lives. This research is highly focused on interactive technology or elements of such technologies. The group's research into interaction design is driven by the design and implementation of concrete functional prototype systems and technology probes. Prototype systems are designed and built to fit specific use contexts, explore new technologies, or interaction techniques.

During the period from 2016 to 2020, the group's interaction design research activities can be divided into four streams:

Interaction Techniques The research on interaction techniques is about exploring and evaluating new and innovative ways for people to interact with digital technology. This research is driven by the continuing emergence and improvement of new input and output technologies for human-computer interaction, such as computer vision, multi-touch and context sensors, which enable radically new interaction designs.

Mobility and Mobile Technologies The research stream on mobility and mobile technologies is about designing technology that is either mobile or intended for mobile use. While mobile phones or smartphones and related apps are natural candidates for such research, this research also involves different kinds of mobile technologies like cars or robots.

Smart Homes The research stream on smart homes is about designing and understanding digital technology for the home or domestic setting. This research is driven by the growing diffusion of technologies for the home, in particular Internet-connected devices that enable remote monitoring or controlling, but also technolo-

gies that support people in living more sustainable or making sustainable choices. Finally, this research includes the integration of smart or intelligent technologies or sensors into buildings.

Ubiquitous Computing and Ecologies The research stream on ubiquitous computing and ecologies is about designing new interactive systems for people's everyday personal and working lives, enabling them to do things that they could not do before. The research is driven by the opportunities offered by the increasing prevalence and maturity of ubiquitous technologies in society, such as smartwatches or Internet of Things, allowing for fundamentally new applications of interactive systems.

5.2.2 User Studies

The user studies research area is highly empirical and is concerned with the study, evaluation, and deployment of interactive technologies. Such studies are a natural part of human-centred computing, where the aim is to understand how technologies should be designed or deployed, and the studies can be both normative and descriptive. Empirical studies are carried out through longitudinal real-world deployments, or controlled experiments (in a usability laboratory), depending on the research question and aim. Collected data in such user application studies can either be qualitative or quantitative or may follow a mixed methods approach. This kind of research has a long history in the group, dating back to its work on usability evaluations and in particular comparison of field and lab testing during the late 1990s and early 2000s, and has been very successful in terms of quality, impact, and visibility.

During the period from 2016 to 2020, the group's interaction design research activities can be divided into two streams:

Field and Longitudinal The research stream on field and longitudinal studies is about conducting studies that are typically conducted outside the lab entailing natural and non-scripted interaction and could also include studies that are long-term and longitudinal in nature, where the focus is on changes to interaction over time. As such, a primary focus is for this kind of research is on addressing or increasing ecological validity.

User Experience The research in this area entails methods, techniques, and metrics for conducting user experience studies. In particular, the activity has focused on how new metrics for evaluating or understanding interaction with technology or the perceived value of technology. The research in this area is a continuation of the work in the group on usability testing and engineering.

5.2.3 Design and Development

The research in design and development seeks to understand, design and improve development processes. This includes organizational practices pertaining to agility or UX integration. It further includes research that involves designs that aim in challenging social norms, assumptions, and practices through technology. The research involves empirical investigations of organizations that either develop or use software-based systems and applications. Research efforts engaged with practice in organizations are based on qualitative and proactive research methods such as case studies and action research. Closely related to the empirical work, the research also focuses on research methodology, such as developing new methods and techniques for studying interactive systems.

During the period from 2016 to 2020, the group's research activities on design and development can be divided into three streams:

Digitalization The focus is on topics of addressing how digital technologies are developed and used in companies, public organisations, cities, and society at large. It also includes research on empirical, qualitative research approaches employed when studying digitalization. Examples are research on business cases, big data project and benefits, values of digitalization, data governance, and smart cities.

Agility The focus on agility attends to the rapid changes in the environment of most software development projects and how this has motivated the use of agile development methods in software companies. The purpose is to empirically understand and improve the emerging practices and competences from combining agile methods with safety-critical development, open-source development and user experience work.

Critical Computing This research stream deals with research that aims to make people critically reflect on their behaviours, practices, or beliefs in relation to the social, cultural, or political implications of the technologies they use. Examples of work in this area include challenging energy consumption practices through provocative artifacts, designing technologies for non-use, and research on algorithmic fairness.

5.3 Staff

5.3.1 Current Staff

Professors: Peter Axel Nielsen, Jesper Kjeldskov (20%), Mikael B. Skov

Associate professors: Ivan Aaen, John Persson, Anders Bruun, Timothy Merritt, Dimitrios Raptis, Florian Echtler

Assistant professors: Eleftherios Papachristos, Enrique Encinas, Niels van Berkel, Daniel Russo, Rikke H. Jensen, EunJeong Cheon, Nicolai Brodersen Hansen

PostDocs: Michael Kvist Svangren

PhD students: Stine S. Johansen, Maria Kjærup, Eike Schneiders, Maria Hoffmann Jensen, Alisa Ananjeva, Rune M. Jacobsen, Martin Lindrup, Kasper F. Skov, Ashna M. Zada, Christopher Getschmann, Sujay Shalawadi

Research Assistants: Julie Torpegaard, Anne Ellegaard Christensen

5.3.2 Staff Development

The group has had a significant development in terms of scientific staff over the past five years (grown from 13 members in 2015 to 30 members in 2020), and this is graphically illustrated in Figure 5.1. Peter Axel Nielsen and Jesper Kjeldskov have been professors throughout the period (Jesper was appointed Head of Department during spring 2019 and is currently 20% research-active in the group). During the period we have had several promotions on all levels of the academic ladder. Mikael B. Skov was promoted from associate to full professor. Anders Bruun and Dimitrios Raptis went from assistant to associate professors, while Rikke H. Jensen continued in a Postdoc (in the DPW group) after completing her PhD, and later returned to HCC as an assistant professor, and finally Michael Kvist Svangren went from PhD student to Postdoc. Several staff members went from research assistant to PhD students including Stine S. Johansen, Maria Kjærup, Eike Schneiders, Alisa Ananjeva, and Rune M. Jacobsen. Additionally, the group has been joined by staff from other institutions including Timothy Merritt (Aarhus University) and Florian Echtler (Bauhaus-Universität Weimar) as associate professors, and Enrique Encinas (Northumbria University), Niels van Berkel (University College London), Daniel Russo (University of Limerick), EunJeong Cheon (Indiana University), Nicolai Brodersen Hansen (TU Eindhoven) as assistant professors.

Members of staff have also left the research group during the period where Jan Stage left the group, and Jeni Paay is currently on leave. Jane Billestrup left the group after successfully completing her Ph.D. studies, while Nis Bornoe left to work in industry. Several research assistants including Heidi Nielsen, Sara Nielsen, Mette Elsborg, Shagen Djanian left the group to pursue PhD studies at other departments or working in industry.

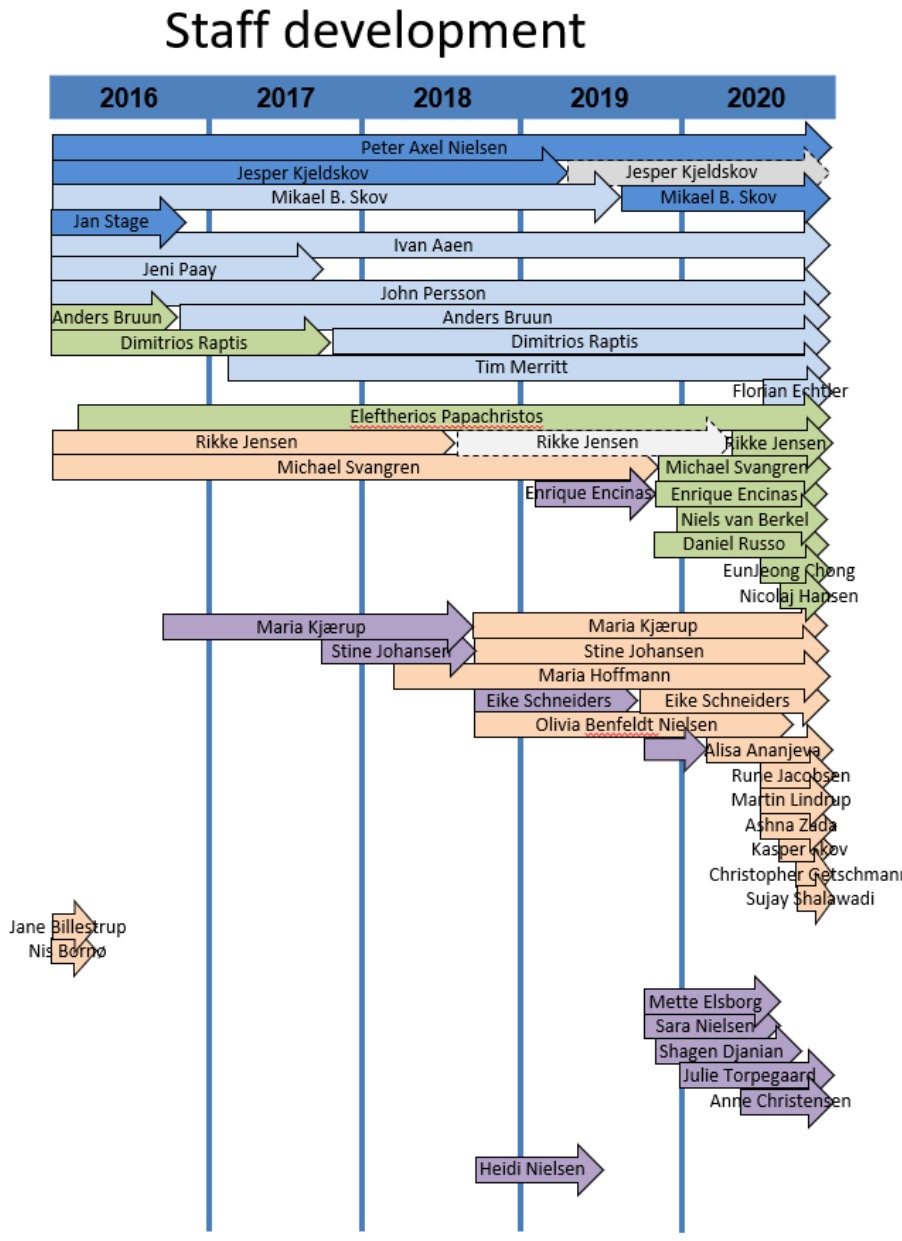

Figure 5.1: Staff members in the HCC group from 2016 to 2020.

5.4 Goals 2016-2020

The following goals for the period 2016-2020 were set as part of the previous research evaluation (January 2016).

Focused publication strategy The goal was to continue to pursue A conferences within the HCI research community and A journals within the IS research community with vigilance, and further to perceive spin-off workshops at A and B conferences as training grounds for PhD students. Finally, it was a goal to increase joint publications across the group.

Research funding The goal was to pursue a two-legged approach for attracting for research funding namely through participation in larger research consortia on societal concerns and grand challenges (e.g., Innovation Foundation and Horizon 2020) and through smaller funding sources often in collaboration with companies (e.g., industrial PhD grants).

Recruitment The goal was to increase research staff numbers based on our strong involvement in the department's education portfolio and in particular the interaction design education and the iDA master's education. The goal was to turn the revenue from these educations into scientific positions and to the right competences to these positions by searching actively and internationally. Further, it was a stated goal to increase the number of PhD students and research assistants.

Industry collaboration The goal was to maintain and expand the close collaboration with software companies but also companies from other industries. The idea was that collaboration would not only confront our research on relevancy but could also provide testing grounds for joint funding applications.

Maintain international and national research collaborations The goal was to maintain collaboration with existing international and national research institutions to increase joint publications and funding. The goal was further to build on existing partnerships with local Danish companies to get their commitment of resources toward joint research projects.

Research themes for 2016-2020 The group identified several themes that seemed very relevant and promising for the coming period:

- **From Usability to User Experience:** The goal was to move on from measuring usability to measuring user experience, and more specifically to study how UX can be measured through psychophysiological measurements.
- **Interaction design for data-intensive cyber-physical systems:** The goal was to do research on interaction design for digital ecosystems, where several

systems, computerized as well as physical, are interconnected and used in concert as an extension of our previous initiatives on ubiquitous and mobile computing.

- **Value-driven software development:** The goal was to conduct research on more intangible or wicked development processes of new software systems and to study them through business cases and business models at the intersection between the software company and its customers and users.
- **Smart city and smart society:** The goal was to pursue research that lies between the smart city technologies and the necessary understanding of users, interactions, benefits, and social processes.

5.5 Activities and Results

5.5.1 Interaction Design

Interaction Techniques The group's research on interaction techniques has primarily been focused on activities related to mobile computing and ubiquitous computing, and we have had a primary focus on techniques for mobile technologies. Often, we are concerned with the design, development, and evaluation of techniques for interaction on such mobile or ubiquitous technologies.

A stream of research during the last five years has been to develop new techniques for cross-device interaction and study them in experiments. These efforts were initiated in the previous period and further elaborated during this period. We have had a particular focus on how mobile devices can be part of digital ecologies, and how interaction techniques can support cross-device interaction, for example how to transfer information or data from one device to another device. As an illustrative example, we designed different techniques (Pinch, Swipe, Throw, and Tilt) and studied how they could be used to move information from a mobile device to a large display.

A different stream of research under interaction techniques has focused on interacting with in-vehicle systems. Cars or different vehicle types often require different interaction techniques than those used in traditional desktop or mobile technologies as drivers need to focus on the driving activity. During the period, we have developed interaction techniques to study how to detect safe and concentrated driving utilizing human activity recognition. Also, we had a research collaboration with Volkswagen Corporate Research to develop ideas and techniques for interacting with future autonomous cars.

In the period 2016-20, this research activity has produced 18 publications of which 4 are A-level conference papers.

Mobile and Mobility The research group has focused for many years on understanding user experience while "on the go." This focus has expanded and has moved beyond mobile devices to include research on people in the mobility context as well, for example to understand the challenges and opportunities for technology with autonomous and shared vehicles.

Research on mobile technologies continued with an expanded focus on user experience and the application of technologies to support people facing the challenges of daily life. This has involved exploring the motivations and behaviours of players in mobile games that involve physical activity such as the popular game, Pokémon Go. Also, personal wellbeing and health supported by mobile technologies has been an important focus area in this period. Outcomes have included the development and theoretical contributions for technologies that help those trying to quit smoking and technologies that help young adults manage low self-esteem.

During this period much focus has been placed specifically on the mobility and transportation context, which has experienced an increase in the use of digital technologies and the rise of electric vehicles. This has involved a PhD research project that examined shared mobility and passenger trip planning for ride-sharing services as well as studies on how to align electric vehicle charging with household electricity production. Also, embracing the emerging smartwatch technologies, research has explored the potential to understand user experience and activity, for example using the inertial sensors to recognize and distinguish drivers from passengers in the automotive context, while another project focused on utilizing sensors in the cabin to assist drivers to reflect on how they hold the steering wheel.

Industrial collaborations have been established and ongoing with Swedish and German mobility companies. One area of research explored how technologies support highly mobile individuals across the journey from door to door (Lufthansa). The work resulted in a characterization on technology use when planning and carrying out trips. Another area of focus has been on the design and evaluation of future interfaces to explore the challenges in the future of mobility. This has involved various master thesis projects in collaboration with (Volkswagen Corporate Research). These projects centred on the development of technologies to address challenges of autonomous driving, interactions in public transportation, and interactions with intelligent agents.

In the period 2016-20, this research activity has produced 16 publications of which 7 are A-level conference papers. One Ph.D. dissertation was produced.

Smart Homes The group's research on smart homes has partly grown out of the group's activities and collaborations on domestic computing. Our previous work on smart homes has primarily focused on how household inhabitants could stay in touch while separated, for example how technology could support connectedness while parents and children were separated due to work or travelling.

During this period, a significant part of our work and activities within smart homes have focused on sustainable behaviour and energy consumption in private households. Particularly we have worked with electricity consumption as part of the DiCyPS Research Centre and more recently the H2020 FEVER project. As part of this, we developed an interface and prototype called HeatDial, which controls private household heat pumps and allows inhabitants to exhibit flexibility on when electricity is consumed. Through collaboration with a local energy company, HeatDial was used in several local households for extended periods of time allowing us to study how people perceived and negotiated electricity consumption and flexibility. With this work, we have attempted to contribute to some Sustainable Development Goals, including "Affordable and clean energy" and the subgoal (7.a) of investing in energy infrastructure and clean energy technology, and "Responsible consumption and production" and it's subgoal (12.2) of achieving sustainable management and efficient use of natural resources. This has primarily been done through understanding of domestic electricity consumption, but also by providing interface for managing consumption. Finally, and as part of energy consumption, we have established collaboration with Yolande Strengers, Monash University and colleagues on sustainable HCI and household consumption. This collaboration has led to a research visits and several co-authored publications.

Another part of the smart home research stream has continued and extended our work on home monitoring of cardiac patients while out of hospital. This work has primarily been conducted in collaboration with researchers and nurses at Aalborg University Hospital and has focused on how patients can become collaborative agents, where they actively provide information and data about their own health and situation while at home. This extends our previous work with home monitoring and how patients perceived more automatic logging and registering of data.

In the period 2016-20, this research activity has produced 13 publications of which 4 are A-level journal articles, and 3 are A-level conference papers. One PhD student has graduated from this activity.

Ubiquitous Computing and Ecologies An important direction of research over this period has been to look beyond the traditional, screen-based user interfaces to embrace the increasingly diverse and complex interconnectivity of digital systems and modalities. These interconnected systems and artifacts are slowly merging into

everyday objects and environments and have become ubiquitous. Various activities have explored the design and user experience supported by such embedded and networked computing systems and have included contributions to theoretical understanding of ubiquitous media and ecologies. We also participated in a Dagstuhl seminar on ubiquitous computing called "Ubiquitous Computing Education: Why, What, and How". In terms of modalities and formal characteristics, the group has explored tabletop interactions, sound zones, pervasive games, embedded interfaces in the body, living materials, and child-computer interaction.

During the period, we started to include robots (and drones) in our work with ubiquitous computing and ecologies. We joined a cross-department research initiative on HRI at the TECH faculty at AAU and a PhD project has been established that explores human-robot interaction between multiple robots and people. An overlapping area of research has involved the materials of interaction including shape-changing interfaces. This research has focused on expanding theoretical understanding and mapping the main challenges and aspirations related to this field of research. Research collaborations with the Human Media Lab at Queen's University has focused on human-drone interaction, which has involved interactions with drone swarms as an interactive 3D display. This collaboration resulted in various public demonstrations including a public expo and televised demonstrations.

The ISOBEL project was started during this period, which focuses on the development and interactions with sound zones in the home and in the medical context. Sound zones enable sound to be selectively delivered to various co-located people in a room to support the needs of everyone whether they want to listen to their own unique audio track for a movie or to enable patients in a busy hospital to find comfort without the disturbing sounds of diagnostic equipment or other distractions.

Through collaboration, we extended our work on ubiquitous computing to utilizing living materials for interactive systems. Together with experts in the field, we led the review of Living Media Interfaces and mapped out the open research challenges for interaction with biological materials in HCI and took part in the microfluidics workshop at Atlas Institute in Colorado. Empirical research was conducted to investigate the personal meaning and experiences with technologies that are embedded into the body (Underneath the Skin) In this intersection between robotics and biology, the group was part of a successfully funded Erasmus+ project ABRA (Artificial Biology, Robotics and Art), which is a joint project to develop a joint master's program to explore the transdisciplinary intersection of biology, robotics, and art.

A continuing area of focus for the group has been on child-computer interaction.

This has involved participation in a special interest group (ACM SIG), research collaborations with industry partners, public demonstrations and user studies with children. Together with the LEGO Foundation, master thesis students have conducted research on Braille bricks for teaching computational thinking for blind children and another group developed a simplified cardboard robot for supporting physical and social play. This led to the development of 4 highly refined prototypes called "Robert Robot" that were deployed and studied at the LEGO World Expo in 2019 with children in an open play area.

In the period 2016-20, this research activity has produced 22 publications of which 2 are A-level journal articles, and 7 are A-level conference papers.

5.5.2 User Studies

Field and Longitudinal The group's research on field and longitudinal user studies has partly grown out of the previous activities in the group on usability testing in the field, but recently we have moved beyond focusing on usability metrics to focusing on broader aspects of use and utility of technology. This shift in focus can partly be explained in the general focus of user studies within HCI research, but it is also a result of the role that user studies play in our research activities. Currently, the group's research often has a focus on providing empirical evidence or accounts of the role and use of off-the-shelf interactive applications or developed prototypes with emphasis on natural and non-scripted interaction and on ecological validity.

Field studies involving users play a significant role in our community, and we have continued to conduct field studies as part of our research activities. In the DiCyPS research centre, we work with digital technologies for mobility and energy consumption, and both domains have included field studies of developed prototypes or commercial products. As an example, we have conducted studies of electric car owners and their use of technology to support mobility. Data collection has included on-site visits and interviews, technology tours in private homes, and co-driving sessions. During the period, we attracted funding for two pilot studies of different forms of mobility. In the YILU project, we collaborated with a German travel company to understand frequent travellers needs and behaviour in door-to-door travel, while the PlusTur project focused on understanding digital mobility services in rural regions of Denmark.

During the period 2016-2020, we have initiated activities on longitudinal studies as a response to address issues like learning or novelty-effects. Thus, our aim with these activities is to study and understand change and stability over time

for human interaction with interactive technology. Longitudinal studies can be difficult to conduct in many research projects (e.g., lack of user commitment, time costs, timing), but our involvement and activities in the DiCyPS research centre provided an opportunity to conduct a longitudinal study with eight households over 18 months controlling heat pumps using our developed prototype. We were able to observe how our participants interacted with the technology over time, and we found that participants kept using functionality in the interface to specify comfort levels in their houses and although they may not have regulated these as often as it is necessary to take full advantage of the developed prototype. While these kinds of studies require resources and time to mature, the results have helped us to illustrate how interaction, user behaviours and consumption change over time.

As part of working with longitudinal studies, we have established new collaborations with prominent longitudinal study researchers namely Prof. Jens Gerken (Westphalian University of Applied Sciences) and Asso. Prof. Evangelos Karapanos (Cyprus University of Technology). Both have been active in conducting longitudinal studies and influencing our research fields, and during 2020, we have (with Gerken and Karapanos) been co-editing a book in Springer Human-Computer Interaction Series on "Advances in Longitudinal HCI Research" with a public call-for-papers, and the book is expected to be published in 2021. Finally, we have introduced a method from ecology to the HCI community, in which field and longitudinal data is analysed through the lens of complex dynamic systems. This has resulted in the organisation of a workshop at CHI 2021 with both upcoming and established researchers in the field (including Prof. Alan Dix, Swansea University).

In the period 2016-20, this research activity has produced 12 publications of which 3 are A-level journal articles, and 3 are A-level conference papers.

User Experience The research in this area entails the interplay between usability engineering and software development practices in the industrial and public sector, identifying users' design preferences of smartphones and Human-AI interaction. The most prevalent work within this activity revolves around methods for data collection to assess user experience, which includes real-time physiological data, questionnaire development and validation, as well as crowdsourcing. We have developed methods and tools to evaluate dimensions of user experience not typically assessed through classical usability testing. In particular, the activity has focused on evaluating and understanding interaction with technology as well as the perceived value of technology.

We developed the "Cool Questionnaire" as a tool to operationalize the notion of coolness. This can be used by researchers and designers to assess how users perceive the desirability, rebelliousness, usability, classic aesthetics and hedonic quality

of a product. These five factors were derived through exploratory and confirmatory factor analysis and enable researchers and designers to target properties of a design and improve its perceived coolness. For instance, a design may be perceived as aesthetically pleasing, but not sufficiently rebellious for the product to be perceived as cool. The tool is offered for free online and it started to have an impact as it is being used from other reseachers and practitioners. In 2020, the International Journal of Human-Computer Studies published our special issue 'Human accuracy in mobile data collection', in which we asked the research community to reflect on the quality of participant data collected in the wild. The special issue built on established and new research relationships, including Prof. Katarzyna Wac (University of Copenhagen, University of Geneva) and Prof. Anna L. Cox (University College London).

Another stream of user experience research has dealt with developing and validating a method to conduct real time assessments of emotional reactions towards interaction designs. This work has taken its offset in the cued-recall debriefing method (CRD) proposed in the early 90's that provides cues to study participants to better recall past events. CRD essentially combines the advantages known from the Day Reconstruction Method (DRM) and Experience Sampling Method (ESM) for collecting data on participant experiences in natural settings. We revisited CRD in the context of evaluating user experiences of interaction designs. We did this by using physiological sensors to identify events during interactions with technology that leads to increased emotional arousal. Physiological data is captured in real time and used as cues for study participants to reflect upon specific events (as indicated by the GSR sensors) at a later and more convenient time of the day. This makes CRD less burdensome on participants compared to ESM yet providing real time data to support recollection through cues, which is not present when using DRM. We have developed and studied the use of CRD with physiological data as cues in controlled experimental settings to make initial validations of this approach. Findings show that study participants have comparable physiological reactions between actual interaction with systems and during retrospective CRD sessions. This indicates that CRD enables re-immersion into past experiences. Also, we found that non-expert participants were able to interpret the physiological data to an extent that goes well beyond random guessing. More recently, we have also deployed this method in natural settings in a context of using physiological data for the purpose of self-reflection. This is also known as "Life-Logging" where participants reflect on daily events that were particularly emotionally charged in order to support mental well-being. This work has established collaboration with Professor Effie Lai-Chong Law, Dept. of Computer Science, Leicester University, UK.

In the period 2016-20, this research activity has produced 24 publications of which 5 are A-level journal articles, and 7 are A-level conference papers. One Ph.D. dissertation was produced as part of this activity.

5.5.3 Design and Development

The HCC groups research in design and development falls in three streams: digitalization, agility, and critical computing.

Digitalization Research in this stream has all in all been conducted in PhD projects and in collaborative action research with companies and public organisations.

Systems development has moved away from being about developing software and developing it from scratch to activities that acknowledge that all development is re-development and that data and in particular big data have become increasingly important and systems do not only involve interaction functionality. Data governance and big data development projects have been in focus in two PhD projects in collaboration with external partners. A PhD project on data governance in public organisations has been conducted and completed during this period. The results address how to govern access and utilisation of closed and open data about citizens and companies in a municipality. Another PhD project has started researching on the utilisation of big data analytics in the global wind turbine company Vestas Wind Systems. The results concern how to improve on the benefits realization from analytics projects. These projects link well with ongoing research on the values co-created through digitalisation, and a key result has been published as an approach to formulating business cases for digitalisation.

Research on digitalisation of energy systems with a particular focus on consumers of district heating has started and is still ongoing. The research sees collaboration with two other research groups at the university in energy planning and architecture district heating providers, and a company developing consumer apps for monitoring energy consumption. The research has led to stronger ties between the partners and the partners have engaged in large efforts in applying for external funding. The research has so far reached results on digital infrastructures and how (little) effect consumer data has on consumers in reducing energy consumption.

Empirical research approaches have been extended by contributing to problematization as part of action and design research. This has led to results in how to empirically support the problematization process as to push in the direction that action and design research not only leads to solution, but also that the qualities of these solutions can be explained vis a vis their problem setting.

A recent project has started collaboration in a large externally funded project, DREAMS, where the purpose is to study how environmental assessment for large building constructions can be digitalised. The project employs action research and involves several partners from building contractors, ministerial environmental offices, engineering companies and universities (environmental assessment, linked databases, natural language processing). The project has started late 2020 and funds a PhD student and a postdoc.

In the period 2016-20, this research activity has produced 19 publications of which 6 are A-level journal articles, and 3 are A-level conference papers. One PhD student has graduated from this activity.

Agility The HCC group's research on agility in software development has been ongoing since the 1990's, which at the time was under the heading of software process improvement. Today, many software development organizations are committing to agile methods and practices, which has been the area of concern for several case studies and action research efforts in the HCC group. In the last 5 years, several publications in good journals and conferences show a consolidation of this long-lasting research on agility in relation to topics such as safety-critical software development, intrafirm knowledge transfer, small software organizations, and user experience.

Safety-critical software development may be under the constraint that parts of the product development (hardware, mechanics, and clinical) in e.g., a large pharmaceutical company, requires plan-driven elements while still committing to agile development practices. This stream of research has trough case studies and systematic literature reviews resulted in models and practical insights on how intrafirm knowledge transfer and meshing is possible in agile and safety-critical software development.

Agility in small software organizations is an area of concern that compared to a large pharmaceutical company is much more dynamic and employs significantly fewer formal processes. This stream of research is based on case research and action research, which has resulted in insights on how to improve project management agility, dynamic capabilities, and on integrating user experience work with agile development through user stories.

The work on user experience in agile software development has been the focus for small-scale research projects that has enabled a new collaboration across the HCC group. The research stream has focused on effective and efficient integration of user experience analysis and design with agile software development in small software organizations. This collaboration, furthermore, allowed joint learning and

development of case study and action research across the different research traditions in human-computer interaction and information systems. From this research collaboration, a research assistant was recruited for a PhD position for the digitalization stream to also further bridge these research areas. We have collaborated with Marta Larusdottir from Reykjavik University, Iceland and Lene Nielsen from the IT University, Denmark.

The HCC group's long tradition for action research has also attracted a research visit and collaboration on this research method's adaptation to the open-source communities. As part of further strengthening the research on agility, the group has recently hired a new assistant professor that already is researching open source, software quality, and personality traits of software engineers. This also adds to the use of quantitative research and multimethod research experience aligns well with action research and its underlying pragmatic philosophy that focus on problems and practical relevance.

In the period 2016-20, this research activity has produced 16 publications of which 5 are A-level journal articles, and 2 are A-level conference papers.

Critical Computing In the period 2016-20, a new stream of research emerged within the HCC group. This stream aims to move beyond user-centred design and the more traditional focus on the utility, usability, and user experience of technology. Positioned within the 3rd wave of HCI, we call this stream critical computing. The appearance of the critical computing stream was motivated by research efforts within both the fields of human-computer interaction and design. Rather than studying how to produce technologies that make users' interactions more effective and efficient, the stream's research purpose is to make people critically reflect on their behaviours, practices, or beliefs in relation to the social/cultural/political implications of the technologies they use. In most of these research efforts, research is conducted through a physical, digital, or fictional design (Research through Design - RtD) that can be characterized as critical, reflective, provocative, etc., often through the use of specific theoretical underpinnings (such as social practice theory, or critical theory).

Currently, the group's research has focused on challenging energy consumption practices, non-use of technology in specific contexts, embedding reflective and critical thinking in university teaching activities, and the perception of fairness in future artificial intelligence systems. Within the context of energy consumption, the group has produced provocative prototypes that challenged their users and urged them to reconsider their energy consumption practices, expanded on the theoretical understandings of designing provocative technologies, and produced knowledge on how to challenge resource consumption through alternative design

visions such as hygge. All these activities contributed to new understandings on how to reduce the energy consumption of households. Within non-use, the group's research work primarily focused on family settings and explored how technology can be used to make people reflect on their technology use in specific contexts (such as family dinners). Within teaching critical computing, the group has produced a book on how to utilize fiction as a research tool, and a number of publications that summarize different teaching perspectives and approaches within university settings. Finally, on the topic of fairness in AI-based systems we have studied governmental AI policy documents to uncover cultural similarities and differences.

As part of working within critical computing, the group has established collaboration with prominent researchers such as Prof. Mark P. Blythe. Besides co-authoring publications, our activities with Prof. Mark P. Blythe extended in organizing a PhD course on Critical Theory and Interaction Design in Aalborg. Finally, the group was also invited to participate as a guest on the ERASMUS+ project speculativeEDU.

In the period 2018-20, the critical computing research stream has produced 11 publications of which 1 is an A-level journal article, and 8 are A-level conference papers.

5.6 Own Evaluation

We evaluate our research in the last five years against the plans we made five years ago and provide an overall evaluation of our research efforts in the period. We also briefly address the UN Sustainability Development Goals (SDGs), and finally we explicitly address the recommendations made by the evaluation panel in 2016.

5.6.1 Publication Strategy

As a continued effort, we set out to consolidate our publication strategy by pursuing A conferences within the HCI research community and A journals within the IS research community. By far the most significant part of our own evaluation is on the topic of quality and impact of publications. This was also part of the plan for the group almost ten years ago (to focus on A journals and A conferences), and in 2016 we decided to consolidate this strategy. Within IS research, there is a strong focus on journals, e.g., the AIS Bucket-of-8, while HCI research has a strong focus on conference publications including among others the ACM CHI conference. Quality of research publications can be measured in terms of how high ranking our outlets are and impact can be measured in terms of how well cited

our publications are. We first analyse the number of publications at A-, B-, and C-levels and then we analyse h-indexes for staff members in the group.

Production Analysis The overall publication production was relatively stable in the period compared to the previous period going from a total of 138 publications (during 2011-2015) to 141 publications (during 2016-2020). This number covers an increase in numbers of A level publications (going from 60 to 70), and minor decreases in B and C level publications (see Table 5.1). Calibrating for the increase in staff numbers, the total number of publications are slightly lower than previous period, but it is important to understand the group's production based on some observations. First, while the effective number of senior staff has been stable during the period (fewer full professors versus more associate professors), junior staff has increased significantly in numbers. Especially, the number of assistant professors, PhD students, and research assistants has gone significantly up, and PhD students and research assistant productivity are generally (and expectedly) lower than that of senior staff. This can partly explain the slight lower production. Also, the consolidation (and even increase) of A level publications can partly explain this slightly lower production.

In order to assess if the publication strategy for 2016-20 has been effective, it is necessary to analyse the numbers further. A-level journals have gone up from 20 (14%) to 26 (18%), whereas A-level conferences have gone up from 40 (29%) to 44 (32%). It should be noted that during the previous period, the number of A-level conferences went from 21 to 40 implying that the group has maintained and consolidated a high focus on top level conferences. On the other hand, the number of workshop publications (level C) has gone slightly down from 27 (20%) to 22 (16%), and perhaps play a minor role as training grounds for PhD students than during previous periods. Several of the PhD students in the group have been successful in publishing their work in A- and B-level outlets. This has been a deliberate ambition within the group to motivate and train the PhD students to publish at the highest level.

The increase in staff numbers is also visible in production of publications with total numbers over the years namely 2016 (N=20), 2017 (N=18), 2018 (N=33), 2019 (N=19), 2020 (N=57).

H-Index Analysis Like the previous research evaluation, we have conducted a simple citation analysis based on h-index of the individual group members in order to understand the impact of the HCC group's research. All senior staff members of the HCC group have a Google Scholar profile page (as part of the department policy) and we used these pages in our h-index analysis.

Table 5.1: Number of publications ranked by outlet level for journals and conferences.

	2011-2015	**2016-2020**
A Outlets	60 (43%)	70 (50%)
A-Journals	20 (14%)	26 (18%)
A-Conferences	40 (29%)	44 (32%)
B Outlets	49 (36%)	46 (32%)
B-Journals	15 (11%)	9 (6%)
B-Conferences	34 (25%)	37 (26%)
C Outlets	29 (21%)	25 (18%)
C-Journals	2 (1%)	3 (2%)
C-Conferences	27 (20%)	22 (16%)
Sum	138	141

Like previous periods, we notice that the HCC group is characterized by having a relatively even distribution in the h-index for the individual members (see Figure 5.2). Here we notice that four members of the group have an h-index above 20 (also N=4 in the previous period) and that the group covers a spectrum in h-index from 5 to 35. Further, we can also calculate (not in the figure) that the total number of citations for the group is currently 20.305 where 10.212 of these citations is from the current research evaluation period (since 2015). However, the group size has grown since 2015 (primarily in 2019 and 2020), but at the same time, two senior members with citations above 2000 have left the group in the period.

We can further observe that all three full professors in the group now have a total number above 3000 citations as opposed to the previous evaluation where none of the full professors had 3000 citations or above.

5.6.2 Research Funding

The goal was to pursue a two-legged approach for attracting for research funding namely through participation in larger research consortia on societal concerns and grand challenges (e.g., Innovation Foundation Denmark and Horizon 2020) and through smaller funding sources often in collaboration with companies (e.g., industrial PhD grants). In the period, the group has pursued attempts and ap-

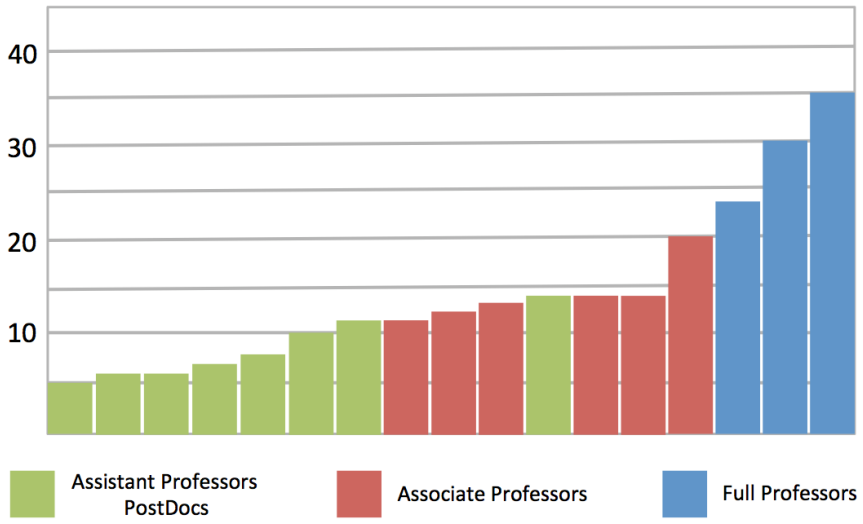

Figure 5.2: Citation illustration for h-index (from Google Scholar). The columns signify individual members of the HCC group including full professors, associate professors, and assistant professors/postdocs.

proaches to attract funding from the Innovation Foundation Denmark and Horizon 2020, but also other funding schemes.

One major externally funded research project, ISOBEL, has started during in 2019. The project is funded by Innovation Foundation Denmark (IFD) and has a total budget of 18 Million DKK (HCC: 5.5 Million DKK). The work for attracting this project was initiated by the group and was started in 2016, shortly after the previous research evaluation and was granted in 2019. ISOBEL builds on the group's activities within smart homes, interaction techniques, and user studies. The HCC group is also part of another IFD project namely DREAMS on digitalisation of environmental assessments, which was granted in 2020. The project has a total budget of 18.2 Million DKK (HCC: 0.8 Million DKK). In 2020 the group got funding from the Industrial Foundation for a project on security in software processes. The project has a total budget of 10.5 Million DKK (HCC: 2 Million DKK). Also, the group got funding for an industrial PhD project on value creation from big data analytics at Vestas Wind Systems A/S from IFD (Total: 2.5 Million; DKK; HCC: 0.3 Million DKK) for an industrial PhD student. In late 2020, the ongoing DFG-funded project DeUCoE with two PhD students was moved from Bauhaus-Universität Weimar in

Germany to the HCC group, with a remaining budget of approximately 1.7 Million DKK. Finally, the group attracted a smaller project PanTra (0.15 Million DKK) from the Carlsberg Foundation on effects of the COVID-19 pandemic lockdown on software professionals.

During this period (2016-20), the group has actively worked on getting involved in EU/Horizon2020 consortia and projects, and this is new compared to the previous period. The group has a small part of the H2020 project FEVER (in collaboration with the other AAU CS research groups) on user studies in smart homes (HCC: 0.3 Million DKK). The FEVER project was started in 2020. But besides this, the group has worked on different approaches and strategies to become part of H2020 applications utilizing professional networks around Europe. For example, we have participated in consortia developing meetings within the network organisation ERTICO, and the group was part of unsuccessful application in 2019 on mobility. The group has also actively worked with collaborators in primarily Greece to become part of H2020 applications, while some proposals have been rejected, the group continues to take part in developing these for possible resubmission. The group is also a part of an Erasmus+ grant ABRA (Artificial Biology Robotics and Art), which is a three-year project to develop a joint master program with Aalto University, University of Trento, and Institute of Advanced Design Studies, with a total budget of 2.4 Million DKK (HCC: 0.1 Million DKK)

The group has been quite successful in attracting strategic funds from Aalborg University during the period. The TECH Faculty and Aalborg University has strategic funds to support and seed research projects and ideas, and during the period the group has attracted funding for new initiatives on energy and district heating digitalization (1.4 Million DKK in 2018), interaction with robots (1.5 Million DKK in 2019), and interaction with sound (1.5 Million DKK in 2018). This funding is typically part of creating consortia and collaboration within AAU with the purpose to create new external funding applications. As an example, the funding for the interaction with sound was part of the ISOBEL application process, while the district heating digitalization is part of a budget-wise similarly sized application also towards the Danish Innovation Foundation. Additionally, the group has got funding for several minor research projects from various sources including public and private companies, e.g., PlusTur (0.25 Million DKK in 2018), Door2Door (0.15 Million DKK in 2019), BUS-Support (0.17 Million DKK in 2020).

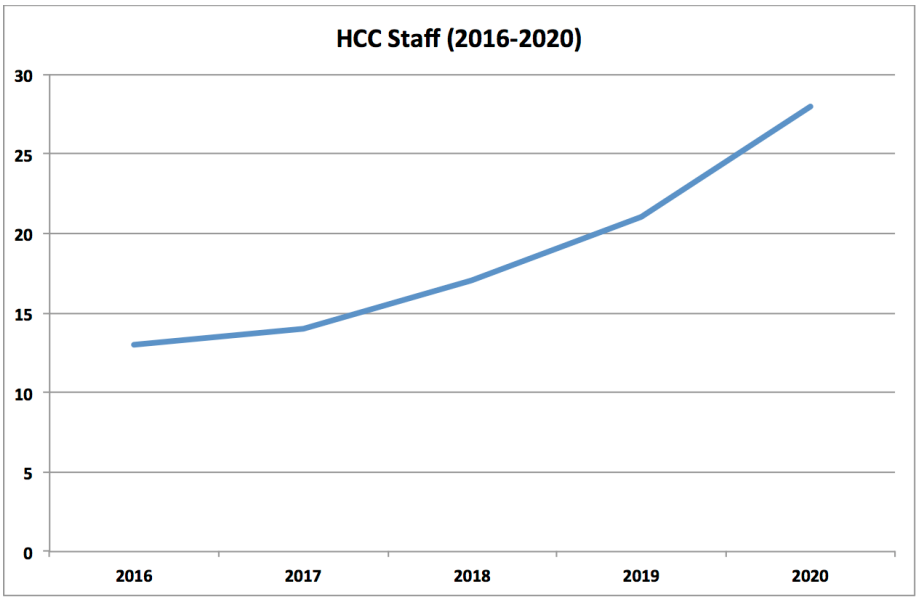

Figure 5.3: Total number of staff in the HCC group including all position types.

5.6.3 Recruitment

The goal was to increase research staff numbers in the HCC group mainly based on the group's strong involvement in the education portfolio of the department. Recruitment was also mentioned by the evaluation panel in the previous research evaluation, who argued that recruitment is one of the most important strategic considerations. The panel suggested to apply different strategies on recruitment to attract high excellence staff for the academic positions (after PhD), but also recommended to increase the number of PhD students in the group.

During the period the group has successfully recruited scientific staff and have had a significant growth in staff number from a total of 13 staff members in 2015 to 30 staff members in 2020. Figure 5.3 presents the total number of HCC staff members during the period, and this confirms a yearly increase in staff members, but also a growing trend with 2020 going from 22 to 30 people.

The faculty in the group (thus excluding postdocs, PhD students, and research assistants) has gone from 9 to 17 staff members. The growth is primarily a result of more assistant professors with three assistant professors in 2015 to eight assistant professors in 2020, while tenured staff (full professors and associate professors) has only had a minor growth from 7 to 9 staff members. Also, the group has still three

full professors, while Jesper Kjeldskov is partly research active. Also, the number of PhD students have increased notably from two to eleven students (this will be further analysed and discussed in section x.6.7). The growth in numbers of staff is partly due to the group's considerable involvement in the department's education portfolio, but also partly due to how economic funds are allocated in the university. Thus, the strong involvement in the department's teaching portfolios has had a significant and positive impact on faculty staff recruitment during this period.

Part of the recruitment strategy has been to attract excellent candidates from international institutions as well as attracting own AAU candidates. As part of this strategy, we have more systematically advertised our open positions through professional network and various social media platforms. Also, we have also recruited while participating at international conferences where we have had informal talks and interviews with potential candidates. These recruitment initiatives have been highly successful, and the group has recruited candidates from several international – and well-known – institutions including University College London, University College Cork, Lero - the Irish Research Centre for Software, Aarhus University, TU Eindhoven, Northumbria University, Indiana University, Bauhaus-Universität Weimar. As an example, one employed candidate, Daniel Russo, was invited to be part of the ACM Future of Computing Academy after a very competitive selection process while applying for an assistant professorship at AAU.

While the group has previously recruited internationally, this period has been different due to the high numbers of new positions, but also the relative high number of international applicants. Currently, six of the seven assistant professors in the group have obtained their PhD degree from another university than AAU, and five of them from a university outside Denmark. Part of this could be attributed to the fact, that the CS department during the period started to offer tenure-track assistant professor positions (three assistant professors in the group are currently tenure-track).

Finally, the group has significantly increased the number of PhD students in the period and currently 11 PhD students are employed in the group compared to the 3-7 PhD students at any time during the previous period. As a result, the permanent staff/PhD student ratio is currently 1.22 PhD students per permanent staff compared to 0.43 PhD students per permanent staff at the end of the previous period.

5.6.4 Industry Collaboration

The group has continued to work closely with industry and other external partners, such as municipalities and semi-public companies. The goal was to maintain and expand the close collaboration of the group with software companies but also companies from other industries. Unlike previous periods where the group had close collaboration with classical software development companies, more of the collaboration partners during this period are not software developers, but typically represent the domain of interest. As examples, we have established collaboration (funded) with Bang & Olufsen, SoundFocus and Wavecare on interactive sound zones through the ISOBEL project. Funded collaboration has also been established with Nordjyllands Trafikselskab (the public transportation authority of North Jutland) on digital opportunities for mobility and Aalborg Municipality on district heating as well as Yilu (Lufthansa) on mobility. Such new types of collaboration still have the same purposes and aims (as previous types of collaboration), namely stressing relevance and importance of the research projects and activities.

5.6.5 International and National Research Collaboration

International and national research collaboration continues to be important for the group and several new collaborations has been established - in particularly with international researchers. This has been a direct consequence of the focus on H2020 applications, but also our focus on specific research topics as well as the changed nature of industry collaboration (as illustrated above).

During the period, the group has developed and expanded the established collaborations with researchers from various Scandinavian universities (e.g., University of Agder, and Copenhagen Business School) and Australian universities (e.g., University of Melbourne, Queensland University of Technology, Swinburne University of Technology, Monash University). These collaborations have integrated PhD student research stays and have primarily expanded the group's activities smart homes and sustainable resource consumption and mobility and has led to several co-authored publications. New collaboration have been established with KTH and Nottingham University in connection with PhD students research stays.

Finally, the recruitment and hiring of external staff members coming from international universities in Europe and US have naturally extended our network and collaboration to involve several new partners. As an example, in 2020 we organized a PhD course on critical theory and interaction design involving a lecturer from Northumbria University, UK as a direct result of hiring an assistant professor from

same institution.

5.6.6 Research Themes for 2016-2020

The group identified several themes that seemed very relevant for the coming period:

- **From Usability to User Experience:** The goal was to move on from measuring usability to measuring user experience, and more specifically to study how UX can be measured through psychophysiological measurements.
- **Interaction design for data-intensive cyber-physical systems:** The goal was to do research on interaction design for digital ecosystems, where several systems, computerized as well as physical, are interconnected and used in concert as an extension of our previous initiatives on ubiquitous and mobile computing.
- **Value-driven software development:** The goal was to conduct research on more intangible or wicked development processes of new software systems and to study them through business cases and business models at the intersection between the software company and its customers and users.
- **Smart city and smart society:** The goal was to pursue research that lies between the smart city technologies and the necessary understanding of users, interactions, benefits, and social processes.

The research funding landscape has highly focused on grand challenges in the past five years, and this has influenced the group's work and themes of research.

Related to the first theme, the group has moved away from usability towards user experience. This is directly reflected in the profile and activities of the group, where usability no longer plays a significant role. This move is also partly reflected in the research community, and the group has worked on different aspects of user experiences, e.g., psychophysiological measurements but also developed a questionnaire for measuring coolness of interactive products.

Related to the second theme, the group implicitly work with aspects of data-intensive cyber-physical systems, but we are not actively using the term cyber-physical systems. This term is part of the research centre DiCyPS and our work on this has primarily to do with interaction design and the four sub-themes of interaction techniques, mobility and mobile technologies, smart homes, and ubiquitous computing and ecologies.

Related to the third theme, the group has worked with the value of new software-based systems. This work includes the development of methods to support value creation in government organizations and for big data analytics projects in a large private organization. This research and its further development take place within the digitalization theme.

Related to the fourth theme, the group is highly involved in several activities around smart cities and smart societies and have worked on challenges of optimising energy consumption (for example electricity and district heating), and the group has generally worked on various aspects of contemporary digitalization.

5.6.7 Panel Suggestions and Recommendations

The evaluation panel provided and listed ten specific recommendations for the group in the previous evaluation. While some of them have been addressed in the self-evaluation above, these panel recommendations and suggestions are discussed individually in the following to explicitly address them.

Renaming the group The group is now called Human-Centred Computing (HCC). The group changed its name after the previous research evaluation as a direct consequence of the recommendation by the evaluation panel. The group was formerly known as the Information Systems group, and as pointed out by the panel, this was partly due to historical reasons. As an area or discipline, HCC generally represents a research field that aims at bridging the existing gaps between the various disciplines involved with the design and implementation of computing systems that support people's activities. In this respect, HCC fits the focus and activities of the group quite well and embraces the different perspectives within the group. Also, similar groups with the HCC name exist around the world, for example the Department of Computer Science, University of Copenhagen, Denmark and at the School of Interactive Computing, Georgia Tech, USA.

Vision and mission The group has a long tradition of having multiple staff members that are equal or peers and clearly this can lead to multiple directions for research and related activities, but on the other hand, it also provides flexibility for the individual researchers who can pursue different goals and adjust to societal changes. During the period, we have worked with the vision of the group – also as part of renaming the group. Also, at the time of changing name, the group and particularly the research leaders discussed alternatives to group names including composite names with two or more disciplines, but HCC was chosen to aim for a bridging name that would span the different disciplines in the group. Internally, this has had a positive effect in the group on perceived group coherence, and we

currently see a growing number of co-authored publications between new author constellations within the group. As part of the work on vision and mission, it is also worth noticing that the group has changed leadership during the period, where one research leader has left the group, one is currently head of department (since early 2019), another is still research leader, and one internal member was promoted to research leader during the period (in 2019). This has impacted the group in terms of research direction.

Impact: As noted in the previous research evaluation, the group has worked over the past decade on scientific impact for research and publications by focusing on outlet quality and is now publishing research in high quality conferences and journals. During the previous period (2016-20), the group has primarily worked with impact through becoming partners in externally or internally funded projects with significant societal impact. First, the research projects in section x.6.2 deal with various societal challenges, e.g., the ISOBEL project on health issues of noise, and almost all our funded projects address significant challenges today. Secondly, the group has actively sought to take part in several new AAU initiatives like the Human-Robot Interaction Lab at AAU or strategic research projects (e.g., district heating) where impact is on working with relevant society problems and often in close collaboration with external partners.

Through the strong and continuous involvement with several of the department's educations, the group has during the last period worked actively to involve student assignments and projects in societal challenges. This has among other things included strategic collaboration with external partners like Aalborg Municipality on contemporary issues within sustainability – for example waste collection and sorting in Aalborg City. The strategic work in this regard has been to elevate the collaboration beyond individual projects and instead involve whole classes of students to increase the impact.

Improving funding applications Funding of research projects was discussed in section x.6.2 on the specific successful projects during this period. We have gradually been working towards having HCC internal review and quality assurance as part of the application process, for assisting younger members of the group in getting relevant feedback on their applications. However, this has not been formalized and has mostly been conducted on an ad-hoc basis - partly explainable to high workloads during the first part of the period, and extensive onboarding of new staff during the last part of the period. Also, in 2020 the department hired a full-time fundraiser, that the group has primarily used for identifying relevant funding schemes, but also for application review and practical matters in setting up project consortia.

Strategic recruitment Recruitment was discussed in detail in section x.6.3. In summary, we have worked extensively with recruitment during the last five years and have strategically initiated and worked with approaches to attract and recruit candidates with high excellence. Among others this has included formal and informal job interviews at conferences, advertising of open job positions at conferences and various online media, inviting candidates to give seminars in the group, using our international network to announce open positions, and offering tenure-track positions since 2019. As a result, we have been able to identify and encourage strong candidates to apply for our positions. We see these strategic initiatives as successful as the group has been able to attract a high number of very qualified people both internal and external and therefore, the group har been able to grow. Also, we have a record high number of open positions for the last couple of years and has been able to fill the positions with qualified candidates.

PhD students As part of recruitment and project funding, the group has increased the number of PhD students significantly. This increase is described in section x.6.3. This increase has been the result of increased funding activity both externally and internally funding schemes, but also thorough the hiring of permanent staff with existing and transferable funding and PhD students.

Multidisciplinary collaboration The group has expanded collaboration activities within the university over the last five years – primarily as part of new research initiatives. As mentioned above, the Human-Robot Interaction Lab initiative is an example of such collaboration, and this has led to collaborations with new people, for example PhD student co-supervision with researchers from Department of Electronic Systems and Department of Communication and Psychology. But several of the funded research projects involve multidisciplinary collaboration with researchers from other disciplines within AAU, e.g., sound people in the ISOBEL project, mobility and transport researchers in DiCyPS and Plustur projects, and energy planning researchers in the district heating project.

Reworking strategy This recommendation had to do with the changing society and societal needs and how this would affect the group and the work in the group. This has been addressed above in "Vision and mission" and "Impact".

Increased ambition and self-esteem A primary ambition of the group over the past five years has been to grow (number of staff) as direct consequence of the group's high and increased activities within the department's education portfolio. As noted elsewhere, this has been achieved, and the group now has a size that is more comparable to the two other groups in the department.

We further see a natural and strong need for human-centred approaches in both

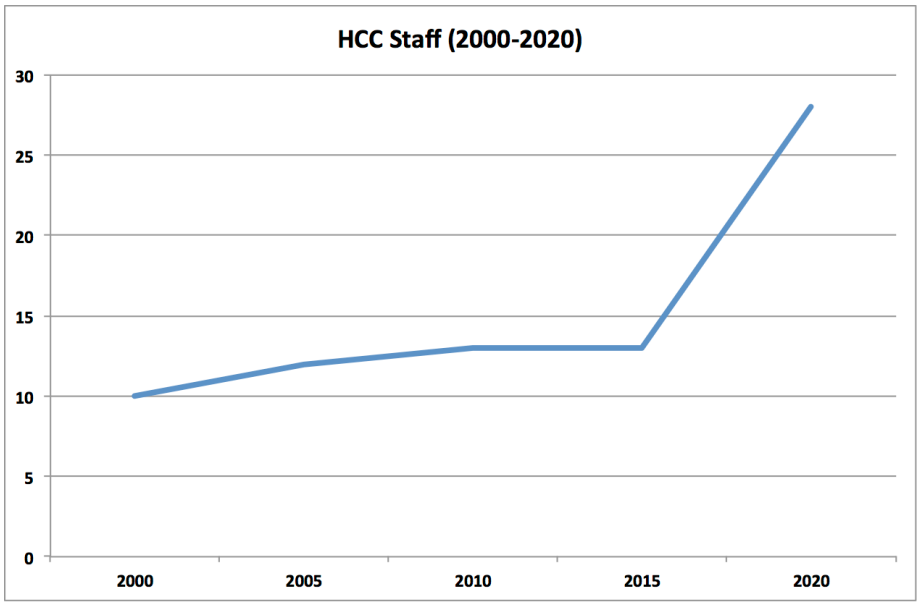

Figure 5.4: Total number of staff in the HCC group since 2000.

computing activities and projects, but also in several societal challenges where digitalization and human involvement is strongly needed.

The strategic importance of education for research development This has been addressed in several parts of this evaluation, e.g., "Impact".

5.7 Plans for 2021-2025

Based on the activities and results from the past five years as well as the self-evaluation, the group has the following plan for the coming period. This includes strategic important areas to focus on as well as specific research themes.

Group Organization and Strategy The significant growth of the group over the past years and the potential to grow even further in the coming years calls for activities and initiatives that aims to orchestrate and guide the group and further develop the strategy of the group. The group has changed in size and thus also organization and structure compared to how the group has been for the last 15-20 years (see Figure 5.4).

The renaming of the group to HCC has helped shaping a group vision and more internal coherence, but also created a clearer profile to the world both internally

and externally. However, at the same time, the increased group size has resulted in a higher number of individual researchers (at different levels) with own goals and specific research areas, and this calls for leadership of the group. Furthermore, the group can expect to grow even further over the next years due to the still increasing activities within teaching, but also the fact that the department has started the software engineering education in Copenhagen where the group is also involved. Finally, the group has changed leadership during the previous period as one research leader left the group, another became head of department, and one was promoted.

Based on the above, the group and the group leaders will in the coming period work with group organization and particularly have focus on re-shaping the strategy and vision, but also actively and explicitly work with group structure(s) and how to organize first order research activities (e.g., paper writing, funding application writing) with second order research activities (e.g., onboarding, culture). To facilitate this process, the group will engage with an external professional research management consultancy to help organize activities, efforts, and direction. This collaboration has already been initiated.

Increased Activity With the increased size, the group plan to increase activities in the coming five years in terms of number of publications, number of submitted applications, and a higher number of funded projects.

For publications, we expect to see a higher yearly number of publications compared to the previous period and while maintaining our focused publication strategy (focus on A level outlets), we plan to have increased activity for lower-level publications through secondary submission categories like poster presentations, demos, workshops, etc. While lower-level publications have less prestige and perhaps direct impact, such publications can provide a good and supplementary platform for dissemination of research results. Also, workshop participation enables interaction and collaboration with researchers that sometimes lead to new activities, e.g., joint publications, funding applications, research stays. With several new colleagues from seven national and international institutions in the group, we expect to increase collaboration with external partners at other institutions. Such collaboration could be co-organized workshop activities, joint publications, funding applications and projects, and research stays. Finally, we seek to play a significant role and apply for funding as part of the newly established research centre DIREC.

Funding Application Process As part of the increased activity with funding applications, we will work towards more systematic and structured processes for writing applications including specifying research problems, identifying relevant funding schemes, and providing feedback on application drafts. While such guidance and

feedback are already happening ad-hoc today, we will work towards making this more explicit and structured. Also, with the increased size and partly diverse group of staff members, we see a need to structure this further.

Recruitment and Staffing Depending on university and department economy, we plan to grow even more during the next years. Thus, recruitment is still relevant for group, and we need to strategically focus on what competences are needed in the group – both from a research point of view but also from a teaching point of view. Like the previous period, we plan to continue utilizing various methods for attracting excellent candidates, e.g., inviting them for seminars and interviews, using network to identify candidates.

We further plan to maintain the number of PhD students relatively to the number of permanent staff. But the group would like to increase the number of postdocs, which is currently low compared to the number of other staff categories. With the increase of large consortia with industrial partners, we see a need to involve more postdocs that are more mature as researchers and can work independently on complex problems and challenges. Also, having a larger group, we plan to work on attracting more guest researchers who can stay with our group for shorter or longer periods at all levels of the academic ladder.

Research themes for 2021-2025 A few research themes seem very relevant for the coming period though some are very likely to change as the research areas evolve:

- **Sustainability/SDGs:** We will continue to further align our activities and project portfolio to sustainability and the Sustainability Development Goals (SDGs). Sustainability has been of focus for several years in the group, e.g., the work energy consumption in private households, and it provides relevance to the research activities and projects, e.g., community usages of dietary carbon footprint Also, the growing focus on SDGs in various funding bodies and in AAU's strategy impede this focus for the group.

- **Human-AI:** Artificial intelligence (AI) is rapidly growing in importance these years both as part of the general society and in research, and AI is also relevant for human-centred computing. The group will expand our work on human involvement in the design and evaluation of technologies integrating AI or ML components, and while this has already started, we expect this to grow significantly in the next period. Also, the group will take part in the cross-group AI team in the department. Potential focus areas include collaborative AI, fairness of AI algorithms, or explainable AI.

- **Digitalization:** The group will continue and further extend the work on digitalization. The group will proactively engage with public and private organizations on how digitalization can transform their products, services, and organizations. The engaged scholarship with practitioners should create both research insights and practical benefits for sustainable development.
- **Empirical Software Engineering:** This theme will focus on the empirical observation of software artifacts and software engineering professionals to develop and empirically validate software engineering theories and assumptions.
- **Smart homes:** Driven by the ISOBEL project, we will continue activities around smart homes and particularly how technologies can support domestic living. Our competences lie within tools and methods for understanding domestic routines and how people interact with technology at home.
- **Mobility:** We have had a strong focus on mobile technologies over the past decade, and we plan to continue work on mobility – how can mobile technologies support sustainable transportation in urban and rural areas.
- **Tangible interfaces:** We have a growing interest on tangible and physical interfaces – both on how to design such interfaces and how to study these either in the lab or in the field. Tangible interfaces play an important role in more and more interfaces, and our newly established facilities in the design studies.
- **Critical Computing:** We plan to continue working with interaction designs that aim in helping and making people reflect on their behaviours and practices within a variety of contexts (e.g., collaborative robots on the factory line, gender issues), challenging the status quo in general, and studying their impact in depth typically through field studies.

5.8 Sustainable Development Goals

During this evaluation period (2016-2020), the UN Sustainable Development Goals (SDGs) have become increasingly important both in society in general, but also in research activities on focus and funding. This has also been the case for the HCC group. Several of our projects and activities deal with challenges that are directly or indirectly related to the SDGs, and more funding bodies have increased their focus on these goals. The focus on SDGs and sustainability more generally works

well for the HCC group for many of our activities and provides focus and relevance for our work. Examples of our growing focus and work within sustainability are the activities on the following:

- **SDG 3: Good Health and Well-Being.** During the last part of this period, we have initiated work on sound zones and unwanted noise (which have impact on stress and well-being). We have particularly focused on how users can interact with these sound zones and how such interaction can be studied in different contexts (private homes and hospitals). This work will continue in the coming period.
- **SDG 11: Sustainable Cities and Communities.** Our activities have studied how digital technologies can support more green and sustainable transport and mobility. This has particularly focused on car and ride sharing services, and how they are used and applied, but also how digital technology can support transportation in rural parts of Denmark.
- **SDG 12: Responsible Consumption and Production.** Our activities have primarily focused on consumption of energy resources in domestic homes and particularly electricity consumption and district heating. Here the group's activities address among others the sub-goal (12.2) of achieving sustainable management and efficient use of natural resources with research on how households can or should consume electricity and heating. We plan to pursue further work with SDGs for the next years (further discussed in section 5.7).

5.9 AI and Machine Learning

The HCC group has during the period started to work on design and use of technologies that involve AI or machine learning technology. Within the human-centred computing research disciplines, interacting with AI technology has been growing significantly over the past years and is sometimes referred to as Human-AI. While we anticipate more focus on Human-AI in the coming period, we already have activities on AI or machine learning, and the HCC group has also joined the cross-group AI-initiative at the department. Here we provide a short list of AI/ML-related activities from the previous period.

- **AI Policy and Algorithm Fairness.** We have started to investigate how different regions and countries implement and describe AI policies as part of regulations and governing. Related fairness of AI-based systems, we have

studied the effects of visualisation techniques on perceived fairness of predictors.

- **Mobility and Human Activity Recognition.** In line with our work on mobility and transportation, we have started to investigate how people can interact with technologies that utilizes human activity recognition. Particularly, out focus has been on how ML technology can detect safe car driving.

- **Interaction with Domestic Digital Assistants.** As digital assistants increasingly inhabit domestic homes, we have during the period started to study how people interact with such assistants and which roles the assistants play in domestic settings.

5.10 Committee Evaluation

The report authored by HCC provides an extensive description of the journey from where the group was in 2015, what measures have been taken over the last five years and an outline of potential research themes for the upcoming five-year period. In addition, it provides a more in-depth analysis of the various research themes currently being conducted, a description of the staff, with people leaving and new recruitments over the period, an analysis over the publication strategy and an analysis of citations and h-index based on Google scholar. In addition, interviews were conducted with the head of department separately, a meeting with the management team of the department, a meeting with assistant professors and postdocs as well as with the PhD students.

First of all, we want to praise the tremendous development that has taken place over the last five years, with many new hirings, substantial growth and development in the size of the organization. The Human Centered Computing group is doing well in terms of growth in staff, education (basic and advanced) and societal impact. The field of HCC has an important contribution to make to the future of CS in education and research. It is important that HCC issues are considered, and the user perspective and societal concerns are addressed, at the university as well as in society. It has an important contribution to make due to its multidisciplinarity. Just to take on example the profession of User Experience (UX) is one of the fastest growing disciplines in the IT sector. The HCC group shows great industrial collaboration and relevance through its industrial collaboration. Additionally, Aalborg university has a great reputation for innovation and problem-based learning (PBL) which makes up an exciting potential for collaboration.

Research areas – lacking a coherent strategy and vision for the future Let us first talk about the HCC research focus. In 2´the report the research themes for 2016-2021 was presented as:

- From Usability to User Experience
- Interaction design for data-intensive cyber-physical systems
- Value-driven software development
- Smart city and smart society

But in the presentation of the research the following emerging themes are discussed:

Interaction design
- Novel technology-oriented interaction techniques,
- Mobile technologies from phones to autonomous systems,
- Smart homes supporting sustainable living,
- Ubiquitous technologies, new applications for work and leisure.

User studies
- Long-term longitudinal field studies,
- Methods, techniques and metrics for UX.

Design and development
- Digitalization for business, public and society,
- Agility in combination with safety-critical, open source and UX,
- Critical computing; social, cultural and political.

This development is a natural consequence of the changing situation in the surrounding society, of the pursuing of the interests of the staff being recruited and as natural consequences of the received funding and opportunities that occur. However, the presentation gives at hand an impression of a strategy meaning that the HCC research group "lets a thousand flowers bloom" and, consequently, may risk losing coherence, convergence, and opportunities to building critical mass and strength in a few selected areas. It is about making priorities.

As a result, the vision of the future research themes for the upcoming 5-year period contains the following themes:

- Sustainability Development Goals (SDG)
- Human-centred AI
- Digitalization
- Empirical software engineering

- Smart homes
- Mobility
- Tangible interfaces
- Critical computing

Some of the themes are building upon earlier research and some are new emerging topics that may entail opportunities for funding, collaboration and societal challenges or missions. Pursuing issues such as human-centered AI, digitalization and SDG is timely and smart and looks promising for the future. But when asked about the future themes the impression was that this requires more elaborate work.

Additional observations The HCC department emphasizes its important profile addressing societal impact – addressing societal needs and industrial collaboration – which is an important part of the HCC profile and that should be maintained. But there is of course always the risk that there is an unbalance between industrial research needs and basic research opportunities.

When it comes to publications, the HCC group adopted an improved strategy as a consequence of the earlier evaluation, which has given satisfactory results, but there is still time for further development. But when it comes to citations, much more is to be expected. The citations have developed from 10.000 citations in 2015 to 20.000 citations in total now which is good, but many of the citations are based on the earlier work. One can question whether the citations are a consequence of planned action or if it happened by chance.

Tremendous recruitment has happened in the past 5-year period, particularly recruitment of junior faculty on assistant professor level. The assistant professors are meant to be building the future of the group. The group currently has only one postdoc, which is under-critical for a growing field, but on the other hand the group has really developed its PhD education recruitment, which has been a great and necessary development.

Finally, even if this is a research evaluation, it is important to mention education, which is an important output of high quality and an important solid source of income for the group.

On an organizational level the following observations have been made:

There is a stark difference in both quality and size between the different groups within the CS department, even if some improvements have been made since the

last evaluation, all the groups have grown substantially. The size difference can, in itself, pose many different problems.

The way that the department implemented the recruitment of assistant professors making a difference between tenured track assistant professors with a clear path towards promotion to associate and full professors and non-tenured assistant professors which is a time limited position quite similar to a post doc, but with a bigger teaching percentage. Tenure track versus non-tenure track assistant professors versus postdocs may be a source of many future problems due to the similarity in work task and title, but the hugely different prospects for the future.

The HCC group has good national and international collaboration but what about internal collaboration at the university and, above all, at the department? There is an immense potential for HCC to engage with the research at the other groups at the department where HCC has the potential of making a great contribution.

It also seems like the department in a terrific way has been able to work on improving the culture and collaboration which is promising for the future.

One of the major challenges for the HCC group recently has been the fact that they lost one of the major research leaders of the group to the management position as head of department, which both is a strength in contributing to the important leadership of the department, but that has had an impact on the strategy, production, and the power of the HCC group.

Recommendations – on Research

Vision: There is a need to work on and clarify the future directions of the research. HCC carefully need to choose and package the future research areas and themes. There are clear important and emerging predominant topics in sustainability, AI and ML, digitalization that could constitute a good base for the future research themes, but it needs to be easily communicable, and needs prioritization and strategic investments.

Applied and problem-oriented research: Clearly the applied and problem-oriented research approach is an asset, but there is still a need to focus on quality, novelty, contribution, and innovation. There is no room to hide behind the fact that we are talking about mostly applied research.

Publications: The goal for this last five-year period was to work according to a focused publication strategy. However, the total number of publications during the last five-year period in comparison to the previous five-year period is almost the same, even if the number of staff has gone up (from 9 to 17 faculty). The share of A-level conferences and journals has gone up

slightly and B-level is almost the same. The recommendation is to put an increased focus on, and discussion about publications to help further increase the quality and quantity to be able to approach the quality of the other groups.

Citations: When it comes to citations there is definitely room for improvement, and citations do not happen by accident. The recommendation is to conduct more of an analysis of citations also on an individual level and set targets that may influence the choice of publications, venues, topics, and collaborators. The purpose is to increase awareness of how you gather citations to increase visibility, reputation and subsequently the impact of the research. For example, going for open science publications is one step towards increasing the citations.

Recommendations – on an organizational level

Funding strategy: Highest priority must be given to securing more external funding, either by recruiting somebody with the potential of pulling in funding or by taking in scientific support or coaching to crack the code of securing funding

Compensating for loss of top researcher: The fact that Jesper Kjeldskov became HoD means that HCC lost one of its top researchers with the highest potential of securing substantial funding. This could have been compensated for by recruiting a top-level researcher to provide that type of leadership at HCC.

Leadership support: Provide coaching and leadership education to the identified future leaders of your research.

Recruitment strategy: Recruiting on assistant professorship level is great. Working with search committees to headhunt potential great candidates that are willing to apply for jobs at AAU, invite them to see the environment and give talks, supply appealing startup packages to be able to recruit at a higher level. There is an immense potential in developing AAU's recruitment strategy and processes to ease meeting the ambitious goals set up.

Diversity and inclusion: Increasing gender equity on all levels, from students to faculty, but students need role models, so you need to start with faculty. Great internationalization – 6(7) Ass. Prof did their PhD outside of AAU.

Digitalization of the university: the most crucial factor, transforming the ways all work is conducted, digitalized education, digitalization of research data, transforming the university business to meet the needs of the future

Education: Recognize education as an important asset for your research, make the education more research based, and consider researching your education

DEPARTMENT OF
COMPUTER SCIENCE

AALBORG
UNIVERSITY

SELMA LAGERLØFS VEJ 300
DK-9220 AALBORG EAST · DENMARK

ISBN 978-87-991383-3-3