SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Dagstuhl News

January - December 2009

Volume 12
2011

Leibniz Gemeinschaft

Dagstuhl Publications
Dagstuhl News January-December 2009
http://drops.dagstuhl.de/opus/volltexte/2011/3106

ISSN 1438-7581

Schloss Dagstuhl, the Leibniz Center for Informatics is operated by a non-profit organization. Its objective is to promote world-class research in computer science and to host research seminars which enable new ideas to be showcased, problems to be discussed and the course to be set for future development in this field.

Period:	January - December 2009
Frequency:	One per year
Online version:	http://drops.dagstuhl.de/portals/dagstuhl_news/
Associates:	Gesellschaft für Informatik e.V. (GI), Bonn Technical University of Darmstadt University of Frankfurt Technical University of Kaiserslautern Karlsruhe Institute of Technology (KIT) University of Stuttgart University of Trier Saarland University Max Planck Society e.V. (MPG) French National Institute for Research in Informatics and Automatic Control (INRIA) Dutch National Research Institute for Mathematics and Informatics (CWI)
Membership:	The Center is a member of the Leibniz Association.
Funding:	The Center receives federal and state funding
Information:	Schloss Dagstuhl Office Saarland University Campus 66123 Saarbrücken, Germany Phone: +49-681-302-4396 E-mail: service@dagstuhl.de http://www.dagstuhl.de/

Welcome

Here are the Dagstuhl News for 2009, the 12th edition of the "Dagstuhl News", a publication for the members of the foundation "Informatikzentrum Schloss Dagstuhl", the *Dagstuhl Foundation* for short.

The main part of this volume consists of collected summaries from the 2009 Dagstuhl Seminars reports and manifestos from the Dagstuhl Perspectives Workshops.

We hope that you will find this information valuable for your own work or informative as to what colleagues in other research areas of Computer Science are doing. The full reports for 2009 are available at Dagstuhl's webpage.[1] You may be irritated that you receive the Dagstuhl News 2009 in 2011. Well, not all organizers supply their result digests within the requested time.

Our online-publication service, started to publish online proceedings of our Dagstuhl Seminars, is catching on as a service to the Computer Science community. The Dagstuhl Research Online Publication Server (DROPS) (`http://www.dagstuhl.de/drops/`) hosts the proceedings of a few external workshop and conference series. The Leibniz International Proceedings in Informatics, LIPIcs, is slowly taking on momentum.

```
http://www.dagstuhl.de/lipics/
http://drops.dagstuhl.de/opus/institut_lipics.php?fakultaet=04
```

Please read the announcement to learn more about LIPIcs.

The extension building with 7 more rooms is under construction. It will allow us to run two Seminars in parallel once the building is finished.

Recently, two copies of Dagstuhl have opened their doors, the Shonan Village Center[2], a Dagstuhl in Japan, and one at the Infosys campus in Mysore, India [3].

There are also attempts to establish copies of Dagstuhl in Korea and in China. We seem to have something right.

Thanks
I would like to thank you for supporting Schloss Dagstuhl through your membership in the *Dagstuhl Foundation*. Thanks go to Fritz Müller for editing the summaries collected in this volume.

Reinhard Wilhelm (Scientific Director)

Saarbrücken, December 2010

[1] `http://drops.dagstuhl.de/opus/institut.php?fakultaet=01&year=09`
[2] `http://www.nii.ac.jp/shonan/`
[3] `http://albcom.lsi.upc.edu/ojs/index.php/beatcs/article/view/27/26`

Contents

Chapter 1

Data Structures, Algorithms, Complexity

1.1 Hybrid and Robust Approaches to Multiobjective Optimization

Seminar No. **09041** Date **18.01.–23.01.2009**
Organizers: Salvatore Greco, Kalyanmoy Deb, Kaisa Miettinen, Eckart Zitzler

The seminar "Hybrid and Robust Approaches to Multiobjective Optimization" was a sequel to two previous Dagstuhl seminars (04461 in 2004 and 06501 in 2006). The main idea of this seminar series has been to bring together two contemporary fields related to multiobjective optimization – Evolutionary Multiobjective Optimization (EMO) and Multiple Criteria Decision Making (MCDM) – to discuss critical research and application issues for bringing the entire field further and for fostering future collaboration.

This particular seminar was participated by 53 researchers actively working in multiobjective optimization. The purpose of the seminar was to discuss two fundamental research topics related to multiobjective optimization: interactive methods requiring optimization and decision making aspects to be integrated for a practical implementation and robust multiobjective methodologies dealing with uncertainties in problem parameters, objectives, constraints and algorithms. The seminar was structured to have more emphasis on working group discussions, rather than individual presentations, so that the open and free environment and facilities of Schloss Dagstuhl could be fully utilized.

Overall, the seminar provided a free atmosphere for everyone to speak and discuss freely about her or his research interests and ideas for considering robust and interactive methods for multiobjective optimization. Several future collaborative research strategies were planned involving researchers from both EMO and MCDM fields. It is hoped that in the next Dagstuhl seminar on the topic some of these collaborative research efforts will be presented.

1.2 Adaptive, Output Sensitive, Online and Parameterized Algorithms

Seminar No. **09171** Date **19.04.–24.04.2009**
Organizers: Jérémy Barbay, Alejandro Lopez-Ortiz, Rolf Niedermeier

Traditionally the analysis of algorithms measures the complexity of a problem or algorithm in terms of the worst-case behavior over all inputs of a given size. However, in certain cases an improved algorithm can be obtained by considering a finer partition of the input space. For instance, it has been observed that in certain applications, sequences to be sorted are nearly in sorted order. In this setting one would expect that such sequences should be sorted in less time than a perfectly shuffled sequence. An adaptive sorting algorithm takes advantage of existing order in the input, with its running time being a function of the disorder in the input.

The workshop was organized into a serie of tutorials and "bridging" talks in the first two days, followed by three days of more regular talks grouped by pairs of themes, with a large amount of time left for interaction in the afternoon, and two "exchange sessions" on Tuesday and Wednesday evenings.

The workshop succeeded in attracting many young students, and a proportion of female participants larger than usual in computer science. The survey attests in particular that the workshop suggested new directions of research (22 participants rated the sentence "The seminar identified new research directions." on average of 4.05 out of 5), but that participants would prefer to receive the schedule of the workshop earlier.

During the exchange sessions, many participants mentionned that they enjoyed from hearing about proof techniques and open problems in areas they were not familiar with before. After the session, several participants, both young and more experienced, contacted the organizers separately to express their satisfaction with the social aspect of the seminar.

1.3 Search Methodologies

Seminar No. **09281** Date **05.07.–10.07.2009**
Organizers: Rudolf Ahlswede, Ferdinando Cicalese, Ugo Vaccaro

The main purpose of this seminar was to provide a common forum for researchers interested in the mathematical, algorithmic, and practical aspects of the problem of *efficient searching*, as seen in its polymorphic incarnation in the areas of computer science, communication, bioinformatics, information theory, and related fields of the applied sciences. We believe that only the on site collaboration of a variety of established and young researchers engaged in different aspects of *search theory* might provide the necessary humus for the identification of the basic search problems at the conceptual underpinnings of the new scientific issues in the above mentioned areas. We aim at uncovering common themes and

structures among these problems, by analyzing them through interdisciplinary lens, and tools from a variety of areas, ranging from Algorithmics to Computational Complexity, from Information Theory to Combinatorics. The more recent challenges provided by the areas of Communications and Molecular Biology call for more attention at the application side of the problems. Therefore, together with the conceptual understanding and the efficient algorithmic solutions, we shall focus also on the studies of new heuristics and experimental methods as well as the theoretical understanding of the well established ones.

We carefully chose a group of outstanding researchers, of different expertise but nonetheless fluent in diverse languages of sciences. They brought their different views of the themes of the original proposal of this seminar. Through the several discussions and the two open problem sessions, we aimed at laying the basis for new perspectives, and solutions to arise.

We shall now briefly describe some of the main areas of research and the problems addressed in the talks and in the common discussions.

The ubiquitous nature of group testing makes it a gold mine for investigators in Search Theory. Group testing has been proved to find applications in a surprising variety of situations, including quality control in product testing searching for files in storage systems, screening for experimental variables, data compression, computation of statistics in the data stream model, and testing for concentration of chemical and pathogenic contaminants. Group testing has been recently applied to Computational Molecular Biology, where it is used for screening library of clones with hybridization probes, and sequencing by hybridization. The contributions by P. Damaschke, G.O.H. Katona, A.J. Macula, and E. Triesch reported on some recent development in this area. In the presentation by A. Zhigljavsky, the case when tests can be affected by noise is also considered. Fault-tolerant search strategies were also considered in C. Deppe's talk. He reported on the equivalence between combinatorial channels with feedback and combinatorial search with adaptive strategies, giving new constructive bounds, when the error is proportional to the blocklength/the number of tests.

The study of gene expression, protein structure, and cell differentiation has produced huge databases which are heterogeneous, distributed, and semi-structured. We are interested in the problem of processing queries that involve specialized approximate pattern matching and complex geometric relations. See, e.g., the contribution by E. Porat for application of group testing to problems of approximate pattern matching.

In multi-access communication one has to coordinate the access of a set of stations to a shared communication medium. It is known that this problem and probabilistic group testing are strongly tied. We focussed on the fascinating relations among the combinatorial structures that are at the conceptual bottom of deterministic multi-access communication and non-adaptive group testing, namely superimposed codes and their many variants. The importance of these structures, that appear in an astonishing variety of problems cannot be overestimated. C. Colbourn, H. Aydinian, E. Porat and G. Wiener, presented some new combinatorial constructions for selection by intersection and superimposed encoding.

A new area of research where group testing techniques are finding fertile ground for new developments is the one of compressed sensing. The presentation by C. Colbourn and O. Milenkovic focussed on some aspects of this new fascinating area of investigation.

Evaluating a function by probing the smallest possible set of variables is at the core of studying the decision tree model for Boolean functions. Function evaluation algorithms play also a central role in automatic diagnosis and more generally in computer aided decision making systems. Relevant to this area of research was the presentation by M. Milanič who reported on game tree evaluation in the priced information model. Game tree search was also dealt with in I. Althofer presentation which focussed on Monte Carlo techniques.

Data compression is another area of investigation, which is tightly connected to search. An easy example of such connection is given by the Huffman trees which provide optimal prefix free compression and, equivalently, search strategy with optimal average number of questions. Variants of the Huffman coding problem are also important in problems of information transmission and storing. M. Golin presented new dynamic programming based approach for variants of the Huffman coding problem. T. Gagie reported on constructing minimax trees.

The presentation by E. Kranakis reported on a different model of search, the one of *randezvous* problems. Here, several agents living in a common domain, want to find each other at a common place and time. The question is what strategies they should choose to maximize their probability of meeting. Such problems have applications in the fields of synchronization, operating system design, operations research, and even search and rescue operations planning.

There were also some contributions that extended beyond the set of main topics: K. Kobayashi reported on new results on the capacity formula of finite state channels, and S. Riis, presentation introduced the (private) entropy of a directed graph in a new network coding sense, and related it to the concepts of the guessing number of a graph.

1.4 Algorithms and Complexity for Continuous Problems

Seminar No. **09391** Date **20.09.–25.09.2009**
Organizers: Thomas Müller-Gronbach, Leszek Plaskota, Joseph F. Traub

This was already the 10th Dagstuhl Seminar on Algorithms and Complexity for Continuous Problems over a period of 18 years. It brings together researchers from different communities working on computational aspects of continuous problems, including computer scientists, numerical analysts, applied and pure mathematicians, and statisticians. Although the Seminar title has remained the same many of the topics and participants change with each Seminar. Each seminar in this series is of a very interdisciplinary nature.

Continuous problems arise in diverse areas of science and engineering. Examples include multivariate and path integration, approximation, optimization, operator equations, (stochastic) ordinary as well as (stochastic) partial differential equations. Typically, only partial and/or noisy information is available, and the aim is to solve the problem with a given error tolerance using the minimal amount of computational resources. For example,

in multivariate numerical integration one wants to compute an ε-approximation to the integral with the minimal number of function evaluations.

Still growing need of efficiently solving more and more complicated computational problems makes this branch of science both important and challenging.

The current seminar attracted 58 participants from 11 different countries all over the world. About 30% of them were young researchers including PhD students. There were 53 presentations covering in particular the following topics:

- tractability of high dimensional problems
- computational stochastic processes
- numerical analysis of operator equations
- inverse and ill-posed problems
- applications in computer graphics and finance

The work of the attendants was supported by a variety of funding agencies. This includes the Deutsche Forschungsgemeinschaft, the National Science Foundation and the Defense Advanced Research Projects Agency (USA), and the Australian Research Council. Many of the attendants from Germany were supported within the DFG priority program SPP 1324 on "Extraction of Quantifiable Information from Complex Systems", which is strongly connected to the topics of the seminar.

As always, the excellent working conditions and friendly atmosphere provided by the Dagstuhl team have led to a rich exchange of ideas as well as a number of new collaborations.

Selected papers related to this seminar will be published in a special issue of the Journal of Complexity.

1.5 Algebraic Methods in Computational Complexity

Seminar No. **09421** Date **11.10.–16.10.2009**
Organizers: Manindra Agrawal, Lance Fortnow, Thomas Thierauf, Chris Umans

The seminar brought together more than 50 researchers covering a wide spectrum of complexity theory. The focus on algebraic methods showed once again the great importance of algebraic techniques for theoretical computer science. We had almost 30 talks, most of them about 40 minutes leaving ample room for discussions. We also had a much appreciated open problem session. In the following we describe the major topics in more detail.

Scott Aaronson gave the opening talk on the relationship between problems that are efficiently solvable by quantum algorithms, captured by the class BQP, and the classical

polynomial time hierarchy, PH. This addresses a problem which is open since the earliest days of quantum computing. Scott presented new evidence that quantum computers can solve problems outside PH, and related the question to frontier topics in Fourier analysis, pseudorandomness, and circuit complexity. Valentine Kabanets talked on algebrization, a notion introduced by Scott Aaronson and Avi Wigderson which extends the old notion of relativization considerably. Since the 1970s we know that we need non-relativizing techniques to separate complexity classes like P and NP. Since then, a few techniques have been developed that indeed don't relativize. However, Scott and Avi showed that all these techniques algebrize, but that we need nonalgebrizing techniques to separate P from NP. Hence they have established a new barrier. Valentine proposed an axiomatic approach to algebrization, which complements and clarifies the approach of Scott and Avi. He presented logical theories formalizing the notion of algebrizing techniques so that most algebrizing results are provable within these theories and separations requiring non-algebrizing techniques are independent of them. Algorithms that use only small amount of space draw much attention these days. Meena Mahajan proposed an algebraic variant of deterministic log-space which is motivated by Valiant's algebraic model of computation. A great result was presented by Fabian Wagner: the graph isomorphism problem (GI) for planar graphs can be solved in logspace. This has to be contrasted with the fact that for general GI, we don't even have a polynomial time algorithm. He also showed that the result can be extended to K_5-free and $K_{3,3}$-free graphs.

We had a number of talks on coding theory and PCPs. Eli Ben-Sasson talked on linear codes that are affine-invariant and locally testable. Eli argued that such codes must have a low rate. Sergey Yekhanin considered the Nearest Codeword Problem (NCP) which is known to be NP-complete. Sergey considerably improved the deterministic approximation algorithms known for NCP. Atri Rudra talked on the error detection problem for codes in the streaming model. Many participants were excited to hear a brand-new result of Anna Gal giving lower bounds on the rate of certain locally decodable codes, a class of codes introduced by Katz and Trevisan in 2000. For these codes it suffices to read a constant number of bits of the word received to retrieve one bit of the original input with high probability.

In an impressive talk, Dana Moshkovitz gave a very elegant algebraic proof for the low error PCP Theorem. Since she had to skip many details in the morning talk, she presented a full proof in a special evening session.

Ilan Newman talked on geometric embeddings of finite metric spaces into spaces of small dimension. The celebrated Johnson-Lindenstrauss Theorem states such an embedding for the Euclidian metric. Ilan pointed out that the situation for the ℓ_1-metric is far less understood. He defined a notion related to the dimension, the cut-dimension, and showed an embedding for ℓ_1 into a space of small cut-dimension.

In a one hour lecture, Nitin Saxena gave a very interesting survey-type talk on polynomial identity testing (PIT), with a focus on his own exciting results. Nitin considers polynomials described by depth-3 circuit of the form $\Sigma\Pi\Sigma$, where the top addition gate has fan-in k and the second level multiplication gates have fan-in d. Hence d is the degree of the polynomial. The circuit is associated with a matrix defined from the coefficients of the polynomial defined by the circuit. The rank of the circuit is defined as the rank of this

matrix. If the circuit computes the zero-polynomial, then its rank is bounded. Previously, the best rank bound known was $2^{O(k^2)}(\log d)^{k-2}$ by Dvir and Shpilka (STOC 2005). This bound is exponential in k. Nitin improved this bound dramatically to $O(k^3 \log d)$. This is no longer exponential in k and is close to the optimal bound because there is a $\Omega(k \log d)$ lower bound.

Ronen Shaltiel introduced the notion of typically-correct derandomization of a randomized algorithm A, which is a deterministic algorithm B (preferably of the same complexity as A) that agrees with A on most inputs. The standard notion of derandomization requires B to agree with A on all inputs. Ronen demonstrated that the relaxed goal sometimes allows better derandomization than is known for the standard notion. For example, it is possible to unconditionally simulate a randomized AC^0-algorithm by a deterministic AC^0-algorithm that succeeds on most inputs. It also allows polynomial time deterministic simulation of BPP under assumptions that are incomparable to those used in the hardness-versus-randomness tradeoffs as for example by Impagliazzo and Wigderson.

We had a series of talks on circuit complexity. Arkadev Chattopadhyay considered solution sets of systems of generalized linear equations modulo a composite integer m that is a product of two distinct primes. The main result is that such solution sets have exponentially small correlation with the boolean function MOD_q, when m and q are relatively prime. This bound is independent of the number of linear equations. As a consequence, Arkadev derives the first exponential lower bound on the size of depth-3 circuits of type MAJ of AND of MOD_m computing the function MOD_q. This solves a long standing open problem.

V. Arvind defined the Hadamard product of multivariate polynomials which is motivated by the Hadamard product of matrices. He studied the arithmetic circuit and branching program complexity of the product, showed several applications, and established connections to polynomial identity testing.

Michal Koucký presented a surprising upper bound for polynomial size constant depth circuits built from modular counting gates, CC^0-circuits: the AND function can be computed by uniform probabilistic CC^0-circuits that use only $O(\log n)$ random bits. This has to be contrasted with a conjecture by Barrington, Straubing and Thrien (1990) that the Boolean AND function cannot be computed (deterministic) CC^0-circuits.

Ryan Williams presented a new method for exactly solving certain NP-hard search problems. The high-level idea is to encode a subset of potential solutions of a search problem with a multivariate polynomial that can be efficiently evaluated. This polynomial is then evaluated on carefully chosen points over a group algebra that will "cancel out" all non-solutions and preserve some solutions with decent probability. This basic method has led to new randomized algorithms for several fundamental problems, most notably the longest path problem.

In cryptography, steganography is the art of encoding secret messages into unsuspicious covertexts such that an adversary cannot distinguish the resulting stegotexts from original covertexts. Rüdiger Reischuk pointed out that the commonly used definition of security of a stegosystem has certain pitfalls. Therefore he proposed a different notion of security which is called undetectability.

As is evident from the list above, the talks ranged over a broad assortment of subjects with the underlying theme of using algebraic techniques. It was very fruitful and has hopefully initiated new directions in research. Several participants specifically mentioned that they appreciated the particular focus on a common class of techniques (rather than end results) as a unifying theme of the workshop. We look forward to our next meeting!

1.6 The Constraint Satisfaction Problem: Complexity and Approximability

Seminar No. **09441** Date **25.10.–30.10.2009**
Organizers: Andrei A. Bulatov, Martin Grohe, Phokion Kolaitis, Andrei Krokhin

The constraint satisfaction problem, or CSP in short, provides a unifying framework in which it is possible to express, in a natural way, a wide variety of algorithmic problems, including propositional satisfiability, graph colorability, and systems of equations. This framework has been extensively used in theoretical computer science, both as a mathematical object with rich structure that deserves investigation in its own right and as a versatile vehicle for algorithmic techniques. The constraint satisfaction problem was studied in the 1970s by researchers in artificial intelligence working on computer vision. From the 1980s on, it has been studied in database theory under the guise of the conjunctive query containment problem, as well as in combinatorics and finite model theory under the name of the homomorphism problem for graphs and for arbitrary relational structures. Only in the last decade, however, it was realized that all these problems are different faces of the same fundamental problem. Consequently, it is important to analyze and pinpoint the computational complexity of certain algorithmic tasks related to constraint satisfaction.

Constraint satisfaction has been ubiquitous in computational complexity theory from its early beginnings. For example, as mentioned earlier, propositional satisfiability and graph colorability, two of the very first problems shown to be NP-complete, are particular cases of CSP. Since the constraint satisfaction problem is computationally hard in its full generality, researchers have toiled for the past thirty years to discover tractable cases of CSP and have strived, and continue to strive, to delineate the boundary between tractability and intractability for this problem. During the past two decades, an impressive array of diverse methods from several different mathematical fields, including universal algebra, logic, and graph theory, have been used to analyze both the computational complexity of algorithmic tasks related to the constraint satisfaction problem and the applicability/limitations of algorithmic techniques. Although significant progress has been made on several fronts, some of the central questions remain open to date. The most prominent among them is the Dichotomy Conjecture for the complexity of the decision version of CSP posed by Feder and Vardi in 1993.

The seminar brought together forty researchers from different highly advanced areas of constraint satisfaction and with complementary expertise (logical, algebraic, combinatorial, probabilistic aspects). The list of participants contained both senior and junior researchers and a small number of advanced graduate students.

The seminar included two substantial tutorials: one on the classification of the complexity of constraint languages via methods of logic and universal algebra (given by A. Krokhin from Durham U, UK), and the other on the approximability of CSP (given by V. Guruswami from Carnegie Mellon U, US). The recent breakthroughs on the topic of the seminar were presented by their respective authors in one-hour lectures, as follows:

1. P. Austrin (New York U, US), Approximation Resistance
2. A. Bulatov (Simon Fraser U, CA), Counting CSP
3. M. Kozik (Jagiellonian U, PL), CSPs of Bounded Width
4. D. Marx (Tel Aviv U, IL), Structural Complexity of CSPs: The Role of Treewidth and Its Generalisations
5. P. Raghavendra (U Washington, US), Complexity of Approximating CSPs.

Other participants presented, in 19 further 30-minute talks, their re- cent results on a number of important questions concerning the topic of the seminar.

The seminar was essentially the first meeting of researchers from both the constraint satisfaction community and the complexity of approximation community. The general consensus was that both communities have learned a lot about the goals, methods and techniques of one another, and that there is a potential of a systematic interaction that may cross-fertilize the areas and open new research directions, so it is worthwhile to maintain communication and to arrange further joint meetings.

1.7 Graph Search Engineering

Seminar No. **09491** Date **29.11.–04.12.2009**
Organizers: Lubos Brim, Stefan Edelkamp, Eric Hansen, Peter Sanders

Graph Search algorithms and their variants play an important role in many branches of computer science. All use duplicate detection in order to recognize when the same node is reached via alternative paths in a graph. This traditionally involves storing already-explored nodes in random-access memory (RAM) and checking newly-generated nodes against the stored nodes. However, the limited size of RAM creates a memory bottleneck that severely limits the range of problems that can be solved with this approach. Although many clever techniques have been developed for searching with limited RAM, all eventually are limited in terms of scalability, and many practical graph-search problems are too large to be solved using any of these techniques.

Over the past few years, several researchers have shown that the scalability of graph-search algorithms can be dramatically improved by using external memory, such as disk, to store generated nodes for use in duplicate detection. However, this requires very different search strategies to overcome the six orders-of-magnitude difference in random-access speed between RAM and disk. We discussed recent work on external-memory graph search, including duplicate-detection strategies (delayed, hash-based, and structured); integration of

these strategies in external-memory versions of breadth-first search, breadth-first heuristic, and frontier search; the inclusion of a perfect hash function, as well as combining parallel and disk-based search; external-memory pattern-database heuristics; and applications of external-memory search to AI planning, automated verification, and other search problems. Implicit graph search that is in the scope of the seminar included deterministic and non-deterministic models, as well as game-theoretical and probabilistic models.

Moreover, the seminar was specifically concerned with algorithm designs for implicit graph search on modern personal computer architectures, e.g. subject to several processing units on the graphic card, and hierarchical memory including solid-state disks. Applications areas for new algorithm designs that exploit modern hardware were found in the model checking community, but also in AI planning and game playing.

The industrial impact was located mostly in the area of software validation, and to some extent in the area of hardware verification. We saw large social network analyses and efficient route planning that were close to industrial application. Investment of parallel and distributed hardware led to new and scalable solutions. In some cases, advances in PC hardware like SSD and GPU already could make a difference. We have had one guest from Synopsis that indicated how influencing actual research is in validating chip design. Other participants from industry were discussing the news in the field.

The seminar mixture of graph theoreticians and application oriented researchers was fruitful. On the one hand, we had the algorithm engineers with theoretical background in hierarchical memory algorithm designs, then the AI researchers concerned with solving their single- and multi-agent search challenges, and last but not least the formal method people, trying to certify or falsify hard- and software designs.

Overall, the seminar was a big success for seeding future research. Tutorials helped to bring the communities together. We thereby established that the gap between the fields was not big. Recently published results were mixed with new insights right from the research labs. Among many other important bricks of work there was the observation that different cache structures are very effective in detecting duplicates in RAM. To address new challenges we have had "open spaces" for discussing and attacking unresolved problems and current research trends.

1.8 Parameterized Complexity and Approximation Algorithms

Seminar No. **09511** Date **13.12.–17.12.2009**
Organizers: Erik Demaine, MohammadTaghi HajiAghayi, Daniel Marx

Many of the computational problems that arise in practice are optimization problems: the task is to find a solution where the cost, quality, size, profit, or some other measure is as large or small as possible. The NP-hardness of an optimization problem implies that, unless P = NP, there is no polynomial-time algorithm that finds the exact value of the optimum. Various approaches have been proposed in the literature to cope with NP-hard

problems. When designing approximation algorithms, we relax the requirement that the algorithm produces an optimum solution, and our aim is to devise a polynomial-time algorithm such that the solution it produces is not necessarily optimal, but there is some worst-case bound on the solution quality.

In parameterized complexity the running time is analyzed in finer detail: instead of expressing it as a function of the input size, one or more parameters of the input instance are defined, and we investigate the effect of these parameters on the running time. The goal is to design algorithms that work efficiently if the parameters of the input instance are small (even if the size of the input is large). More precisely, we say that a problem is fixed-parameter tractable (FPT) with parameter k if the problem can be solved in time $f(k)n^c$ for some function f and constant c. That is, our goal is to design algorithms that are polynomial in n and exponential only in the parameter k. The motivation behind this definition is that in practice we do not have to be able to solve the problem for any possible input: we might be able to define some parameter k that is typically small for the instances we encounter.

Until very recently, approximation algorithms and parameterized complexity have been considered to be two different approaches that have very little to do with each other. Indeed, the methodology of the two fields are very different: the design of approximation algorithms has its own distinctive set of tools such as linear programming, greedy algorithms, probabilistic methods, while parameterized complexity uses a different set of techniques: kernelization, bounded search trees, and extremal combinatorics. However, in the past few years, several connections between the two fields were identified that are worth investigating.

During the 4 days of the conference, 23 talks were given by the participants. Five of these talks were 60-minute surveys on various topics: Dániel Marx talked about several existing connections between approximation algorithms and fixed-parameter algorithms; Gregory Gutin talked about the rapidly growing area of parameterization above guaranteed values; Erik Demaine talked about the recent area of bidimensionality relevant to both approximation and fixed-parameter algorithms; Guy Kortsarz talked about relevant problems in wireless network design; and Daniel Lokshtanov talked about lower bounds on kernelization. As an additional highlight of the seminar, Holger Dell presented in detail his exciting new result about sparsification and its applications to kernel lower bounds.

It is becoming increasingly clear that kernelization—both upper bounds and lower bounds—are becoming a central focus of fixed-parameter algorithms. The talks of Lokshtanov and Dell have shown that kernelization lower bounds are a rich topic with much deep work to be done, and the talks of Demaine and Bodlaender have shown that "metakernelization" results are possible for wide ranges of problems. It is expected that in the next few years there will be substantial further progress on the topic of kernelization.

The seminar successfully brought together both experts and newcomers from the two fields of approximation algorithms and fixed-parameter algorithms, with many interesting interactions. The talks left plenty of time for discussion in the afternoon. An open problem session was held on Monday, and problems raised there were discussed by different groups throughout the seminar.

Chapter 2

Verification, Logic, Semantics

2.1 The Java Modeling Language (JML)

Seminar No. **09292** Date **12.07.–17.07.2009**
Organizers: Joseph Roland Kiniry, Gary T. Leavens, Robby, Peter H. Schmitt

Program verification has been a topic of research interest far into the history of computing science. Today, it is still a key research focus, see e.g., Hoare's Verified Compiler Grand Challenge and the Verified Software Initiative. A main facet in this effort is the ability to formally express properties that must be verified. Building on a long line of work in formal methods for reasoning about behavioral specifications of programs, several recent languages balance the desire for completeness and the pragmatics of checkability. In the context of the object-oriented programming paradigm, the Java Modeling Language (JML) is the most widely-adopted specification language in the Java formal methods research community.

The Java Modeling Language (JML) is a formal, behavioral specification language for Java. It describes detailed designs of Java classes and interfaces using pre- and postconditions, invariants, and several more advanced features. JML is used as a common language for many research projects and tools, including a runtime assertion checker (jmlc), tools to help unit testing (jmlunit), an extended static checker (ESC/Java), and several formal verification tools (e.g., LOOP, JACK, KRAKATOA, Jive, and KeY). JML is seeing some use in industry, particularly for financial applications on Java Smart cards and for verifying some security properties of a computer-based voting system.

Since JML is widely understood in the formal methods research community, it provides a shared notation for communicating and comparing many advances, both theoretical and practical, and it serves as a launching pad for research on advanced specification language features and tools. Researchers are using JML to study or express results for a wide variety of problems; these problems include verification logics, side effects (including frame axioms and modifies clauses), invariants, behavioral subtyping, null pointer dereferences, interfacing with theorem provers, information hiding, specifying call sequences in frameworks, multithreading, compilation, resource usage, and security. In addition to the tools

mentioned above, JML is also used to express, compare, or study tools for checking specifications, unit testing, and specification inference. JML is used to state research problems for formal specification languages and for general discussions of specification language design. JML has also inspired at least three other similar specification languages, Spec#, BML, and Pipa, and has influenced the design and tools for Eiffel. Representatives of these communities are included in the invitation list. JML tools are used in the implementation of at least two other specification languages: ConGu and Circus. At present, there are at least 19 research groups around the world that are cooperating on JML-related research. These groups, and others, have published well over 100 papers directly related to JML (see `www.jmlspecs.org/papers.shtml`).

The seminar will pull together and energize the broad community of JML researchers and developers. We plan to have seminar participants work together on JML's documentation, examples, pedagogical materials, and implementation infrastructure. The meeting will also provide a forum for considering changes to the language, for organizing community efforts, and for discussing recent work on formal methods relating to JML and its semantics. We plan to have fewer talks than an average Dagstuhl seminar and much more interaction and working sessions. We intend to involve the participants in writing documentation, examples, teaching materials, and library specifications. They will also discuss and debug software infrastructure and a novel semantics for JML. In addition, they will discuss and help organize the JML community.

2.2 Typing, Analysis and Verification of Heap-Manipulating Programs

Seminar No. **09301** Date **19.07.–24.07.2009**
Organizers: Peter O'Hearn, Arnd Poetzsch-Heffter, Mooly Sagiv

Most of today's software is written in procedural or object-oriented programming languages. Many of these programs make use of heap-allocated data. This is in particular true for object-oriented programs. Thus, analysis and verification techniques for heap-manipulating programs are crucial to avoid and find errors, to optimize implementations, and to verify properties in a huge class of modern software.

The heap has been a major obstacle to more widespread use of verification and analysis for real-world code. In the last ten years, though, research on analysis and verification for heap-manipulating programs has progressed significantly, in work mainly done by three research communities:

1. Ownership and region types for structuring object heaps, for alias control, and for encapsulation. The main idea is to restrict the way pointers are manipulated and/or restrict the shape of the heap.

2. Verification of heap manipulating programs. The main idea is to specify interesting properties of such programs and to develop formal methods for checking if the specifications are met by the program.

3. Static program analysis for heaps. The main idea is to automatically infer properties of programs. For example, many algorithms infer the shape of the heap at various program points.

The central purpose of this Dagstuhl seminar was to bring together top researchers from these three different communities and to investigate the synergies that can result from a combination of the techniques developed by these communities.

Participants and Organization

The seminar had 41 participants with a good distribution over the three research communities mentioned above. We were in particularly happy to have a good number of excellent young researchers as participants.

After the Monday morning sessions where each participant gave a short statement of his/her background and interest, we started with four overview talks covering the central topics and views of the different communities:

- Peter Müller: Ownership based types
- K. Rustan M. Leino: Comparing heap models: Ownership, dynamic frames, permissions
- Greta Yorsh: Shape analysis overview
- Viktor Kuncak: Theorem provers and decision procedures

The rest of the seminar was structured into research presentations (31 talks), presentation of challenge problems (three problems were presented and discussed), and discussions on how to exploit potential synergies of the different techniques.

Remarks on synergies

Ownership type information can be useful to static analyses and deductive verification. Analysis techniques can support type inference, allow generalizing type systems, and can automatically provide information for verification frameworks. Heap structuring techniques used in verification frameworks, like in separation logic, can be helpful to modularize static analyses. Besides combination of the techniques, another dimension of integration is given by the properties of interest such as, e.g., alias control, access modes, encapsulation, heap structure properties, and behavioral interface properties. Often these properties have to be analyzed together. E.g., certain heap analyses can only be applied in a modular way if the program satisfies some encapsulation restrictions. Also, programs that satisfy ownership requirements may be amenable to more efficient program analysis. A good witness of the close relation between functional and structural properties is separation logic.

2.3 Interaction versus Automation: The two Faces of Deduction

Seminar No. **09411** Date **04.10.–09.10.2009**
Organizers: Thomas Ball, Jürgen Giesl, Reiner Hähnle, Tobias Nipkow

Throughout the history of modern logic, there have been two strands of research: finding natural inference systems for a given problem domain and finding automatic procedures for solving specific logical problems. In computer science, these two strands became *interactive* and *automated* deduction. Powerful systems emerged in both camps (Coq, Isabelle, etc. versus Spass, Vampire, etc.), conferences were established, and separate communities developed.

However, none of the two kinds of systems were ideal for program verification. The interactive tools lacked the necessary automation and the automatic tools failed to cater for important aspects like arithmetic. And neither scaled well. Therefore a separate third, application-driven set of techniques and tools were developed. These are based on powerful automatic procedures for particular logical theories, ranging from propositional logic to arithmetic, and their combination, most notably in the form of SMT solvers. At the same time they were integrated with techniques from program analysis and automata theory. Again, a separate scientific community evolved.

Goals of the Seminar

There is clearly not just competition but also synergy among the three different approaches discussed in the previous section. For example, SMT solvers are successfully applied in program analysis and first-order provers are used in interactive systems. The KeY system is the result of combining an interactive approach to program verification with a high degree of automation. However, such combinations often raise questions and problems that require more interaction between the communities involved. These include

- exchange of formats for theories and proofs
- encoding of higher-order problems into first-order logic
- extension of automatic first-order provers with specific theories or abstraction techniques
- using automatic provers as servers that allow to incrementally add and delete formulas
- orchestration of interleaved automated and interactive inference
- rendering results of automated tools in human-readable form
- generation of proof certificates

- exploiting synergies between Abstract Interpretation and SMT solvers
- invariant inference, especially for quantified formulas
- exploiting program structure for efficient search
- test generation and support from SMT solvers
- programming language support for program analysis

The Dagstuhl seminar brought together the best researchers working on interactive and automatic deduction methods and tools, with a special emphasis on applications to program analysis and verification. In total we had 52 participants, mostly from Europe, but also from USA, Israel, and Australia. A good balance between more senior and junior participants was maintained. The program consisted of 39 relatively short talks, which gave ample time for discussion, both during and after the talks as well as during the meals and in the evenings. Altogether, we perceived the seminar as a very successful one, which allowed for cross-fertilization between research on interactive and on automated deduction. Moreover, it also helped to bridge gaps between foundational research on these topics and application-driven approaches; e.g., the transfer of new theoretical results into applications, or the discovery of new research problems motivated by applications.

2.4 Algorithms and Applications for Next Generation SAT Solvers

Seminar No. **09461** Date **08.11.–13.11.2009**
Organizers: Bernd Becker, Valeria Bertacco, Rolf Drechsler, Masahiro Fujita

In the last decade solvers for Boolean satisfiability (SAT solver) have successfully been applied in many different areas such as design automation, databases, artificial intelligence, etc. A major reason triggering this widespread adoption was the development of several sophisticated SAT techniques and as a result, today SAT solvers are the core solving engine behind many industrial and university tools as well.

However, common SAT solvers operate at the Boolean level and, in general, can only solve a satisfiability problem for formulas expressed in propositional logic. Due to the increasing complexity of the considered problems (e.g. exponential growth of the design sizes in circuit verification), in the last years several approaches have been studied which lift the solving engine to higher levels of abstractions and/or logics that have additional representational power, such as quantified Boolean logic or word level descriptions.

A new generation of SAT solvers - namely Quantified Boolean Formula (QBF) solvers, word-level solvers and SAT Modulo Theories (SMT) solvers - have been introduced. Furthermore, due to the development of multi-core processors, research in the area of (thread-) parallel SAT solving is growing and will be increasingly important in the near future.

The seminar brought together 36 experts from both 'worlds', i.e. researchers investigating new algorithms for solving SAT instances and researchers using SAT for solving problems in a range of application domains, with a particular focus in VLSI CAD (but not exclusively restricted to this area).

An intensive exchange during the seminar initiated discussions and cooperation among the participants and will hopefully lead to further improvements in the next generation SAT algorithms. Moreover, since most of the new techniques are not yet deployed in applications - even if they are often more competitive in contrast to traditional solving paradigms - the seminar provided an excellent forum to familiarize researchers in this area with the new techniques.

2.5 Computer-assisted proofs – tools, methods and applications

Seminar No. **09471** Date **15.11.–20.11.2009**
Organizers: B. Malcolm Brown, Erich Kaltofen, Shin'ichi Oishi, Siegfried M. Rump

Our seminars on computer-assisted proofs are intended to assemble a diverse group of scientists working on differing aspects of computer-assisted proofs and verification methods. The current one is the fifth initiated by Rump.

Computer-assisted proofs in general are characterized by the fact that part of a mathematical proof is assisted in an algorithmic way. This includes numerical calculations, taking account of all numerical errors, as well as symbolic computations.

This concept of computer-assisted proofs can be regarded as a special approach to constructive mathematics. In recent years, various mathematical problems have been solved by computer-assisted proofs, among then the Kepler conjecture (a 3 dimensional sphere packing problem), the existence of chaos, the existence of the Lorenz attractor, and more.

A major representative of computer-assisted proofs are so-called verification methods. These are algorithms verifying the correctness of the assumptions of mathematical theorems with rigor. These methods use solely floating-point arithmetic estimating all numerical errors. Therefore these methods are particularly fast. Besides the conference, a 163-page review article on verification methods by Rump was discussed which appeared in 2010 in Acta Numerica.

In our seminar various new and interesting verification methods were presented. For example, Tibor Csendes and his collaborators proved the chaotic behaviour of a double pendulum, a result which made it to large public media such as FAZ and TV programs. Moreover, a number of new and nontrivial problems related to existence, non-existence and behaviour of solutions of partial differential equations were presented. In particular the rigorous enclosure of sloshing frequencies (Behnke) and proof of photonic band gaps (Plum) attracted attention. Moreover, Nakao discussed the convergence speed of finite element smooth solutions on an L-shaped domain, Kobayashi presented a priori-error estimation for the approximate solution of a certain bi-harmonic equation to establish a

verification method for the driven-cavity problem, Nagatou discussed how to establish a theory for verifying the stability of traveling wave solutions for a certain PDE, and Wieners presented an abstract framework for verified constrained minimization. Traditional verification methods for specific problems like Frommer's square root of a matrix were presented as well. The computational speed can be improved by optimized BLAS routines with verified results as presented by Ozaki.

The discipline of symbolic computation is also well-suited to computer-assisted proofs. In particular the interplay between approximate and exact algebraic number arithmetic has recently lead to irrefutable computer proofs of real optimization problems that were unachieved before. Several researchers from symbolic computation presented approaches that either used exact methods or hybrid symbolic-numeric methods.

Exact linear algebra algorithms as they are found in the LinBox library can assist in proving theorems in graph theory, such as graph isomorphism problems (Clement Pernet's talk). Exact methods in polynomial algebra, in particular singularity removal from algebraic varieties are deployed in Mohab Safey El Din's (with Hoon Hong) Variant Quantifier Elimination (QE) algorithm. By relaxing the I/O specifications in Tarski's QE problem, instances that are notoriously difficult to tackle by software, for example from control theory, have become doable by VQE.

Another way of turning numerical computations into exact symbolic proofs is to prove real polynomial inequalities. At task is to consider a sum-of-squares proof (Artin's theorem) and first to proceed inexactly by a numeric semidefinite program solver, and second to convert the SOS expression into an exact polynomial identity with rational coefficients. Erich Kaltofen, Zhengfeng Yang, and Lihong Zhi with their students have successfully proved an instance of the monotone column permanent conjecture, given proofs for large degree for highly accurate factor coefficient bounds (known also as Siegfried Rump's model problem), and proved the positivity of polynomials arising in Voronoi diagram computations. H?rter used the sums-of-squares approach in a verified optimization algorithm in pure floating-point.

In computer algebra a common tool is exact arithmetic. Operations in the field of rational numbers or algebraic extensions are performed exactly rather than approximating them. Another way to get rid of errors in floating-point operations are so-called error-free transformations. Here the result of an operation between floating-point numbers is represented exactly by a pair of floating-point numbers. These approaches produce results at tremendous speed because those pairs are computed themselves with few floating-point operations. Among others Graillat presented how to compute dot products over finite fields using error-free transformations in pure floating-point, or Ogita extended an algorithm by Rump to solve extremely ill-conditioned problems with condition number way over 10^{100} in double precision floating-point arithmetic to calculate an LU- or Cholesky-decomposition.

Some of the basic algorithms using pure floating-point algorithms need very few operations, so that testing and comparing the quality of such algorithms becomes difficult. To overcome this, Langlois used detailed models of computer architectures for recent machines to model instruction-level parallelism. This technique explains why certain algorithms using error-free transformations are much faster than suggested by a pure flop count.

Finally, Arnold Neumaier allowed us to view the possible future of proofs in mathematics

with his FMathL. In this ambitious project plain TeX-files of mathematical proofs shall be syntactically and semantically analyzed and transformed into formalized and computer- and human-readable proofs. Although it sounds pretty futuristic, detailed plans suggest that it can be achieved some day.

The organizers refrained from presenting talks to give more space to the participants. As always, they and the 46 participants from 10 different countries of the seminar enjoyed the pleasant and stimulating atmosphere in Dagstuhl. Our own assessment is that computer-assisted proofs have several new exciting directions pursued by a number of established and young researchers, and we are already looking forward to the next seminar.

2.6 Software Synthesis

Seminar No. **09501** Date **06.12.–11.12.2009**
Organizers: Rastislav Bodik, Orna Kupferman, Doug Smith, Eran Yahav

Recent years have witnessed resurgence of interest in software synthesis, spurred by growing software complexity and enabled by advances in verification and decision procedures. This seminar brought together veterans of deductive synthesis as well as representatives of new synthesis efforts. Collectively, the seminar assembled expertise in diverse synthesis techniques and application areas.

The first half of the seminar focused on educating the participants in foundations and empirical results developed over the last three decades in the mostly isolated synthesis communities. The seminar started with tutorial talks on deductive synthesis, controller synthesis, inductive synthesis, and the use of decision procedures in program synthesis. The second half of the seminar led to a lot of discussion, boosted by talks on specific software synthesis problems.

The participants agreed that there are several reasons to actively explore synthesis now. First, software development, always non-trivial, is likely to become more complicated as a result of transition to multi-core processors. The hope is that we will synthesize at least the hard fragments of parallel programs. Second, deductive program verification and synthesis are intimately related; it seems promising to explore whether results in model checking and directed testing enable interesting synthesis. Third, by incorporating verification into synthesis we may be able to synthesize programs that are easier to verify than handwritten programs. Finally, the continuing Moore's Law may enable search powerful enough for synthesis of practical programs.

The seminar also led to identification of principles and open problems in benchmarking of software synthesis tools. In contrast to benchmarking of compilers and verifiers, experiments with synthesis must evaluate end-to-end benefits in programmer productivity; in particular, can the program be developed faster with the synthesizer than with a modern programming language? Short of performing a controlled user study, little can be said about the magnitude of these benefits. The situation is more favorable when comparing synthesis tools. The participants agreed that experiments reported in the literature must

identify the knowledge that the user had to formalize in the domain theory that made the synthesis possible. It was also deemed important to identify the formalism used in expressing the domain knowledge.

General Conclusions from the Seminar. The participants found the seminar to be educational and inspiring. We believe this was because of the unusual breadth of participants as well as the format, which revolved around tutorial-style talks that brought the participating communities together.

The participants believed that the talks should be shared with graduate students, who are usually exposed in their courses only to a fraction of synthesis techniques. This observation led to organization of summer school on synthesis, which will be held in Dagstuhl in summer 2011.

The need to create a collection of diverse synthesis results also led to a special issue of the STTT journal of software synthesis, which is under preparation.

2.7 Coalgebraic Logics

Seminar No. **09502** Date **06.12.–09.12.2009**
Organizers: Ernst-Erich Doberkat, Alexander Kurz

The seminar dealt with recent developments in the emerging area of coalgebraic logic and was the first Dagstuhl seminar on that topic. Coalgebraic logic is a branch of logic which studies coalgebras as models of systems and their logics. It can be seen as generalising and extending the classical theory of modal logic to more general models of systems than labelled transition systems. Traditionally, modal logics find their use when reasoning about behavioural and temporal properties of computation and communication, whereas coalgebras give a uniform account for a large class of different systems.

The seminar discussed foundational topics in a particular branch of logic, so problems which command a direct application in an industrial context were outside the seminar's scope. We expect, however, that specification methods related to coalgebraic logics will enter fields like model checking and other areas of industrial interest, once the mathematical foundations in this area are firmer and better understood.

Background

The following glossary puts coalgebraic logic in its larger context.

Modal logic is a field with roots in philosophical logic and mathematics. As applied to computer science it has become central in order to reason about the behavioural and temporal properties of computing and communicating systems, as well as to model properties of agents such as knowledge, obligations, and permissions. Two of the reasons for the success of modal logic are the following. First, many modal logics are —despite their remarkable expressive power— decidable and, therefore, amenable to automated reasoning

and verification. Second, Kripke's relational semantics of modal logic turned out to be amazingly flexible, both in terms of providing techniques to prove properties of modal logics and in terms of allowing the different applications of modal logic to artificial intelligence, software agents, etc.

Coalgebra is a more recent area. Following on from Aczel's seminal work on non-well founded set theory, coalgebra has been developed into a general theory of systems. The basic idea is that coalgebras are given with respect to a parameter F. Different choices of F yield, for example, the Kripke frames and models of modal logic, the labelled transition systems of process algebra, the deterministic automata of formal language theory, or the Markov chains used in statistics. Rutten showed that, in analogy with universal algebra, a theory of systems, called universal coalgebra, can be built uniformly in the parameter F, simultaneously covering the above and other examples. Crucial notions such as behavioural equivalence (observational equivalence, bisimilarity), final semantics and coinduction find their natural place here.

Coalgebraic logic combines coalgebra and modal logic to study logics of systems uniformly in the parameter F. Given the plethora of different transition systems and their ad hoc logics, such a uniform theory is clearly desirable. Uniformity means that results on, for example, completeness, expressivity, finite model property and complexity of satisfiability can be established at once for all functors (possibly satisfying some, usually mild, conditions). Additionally, there is also a concern for modularity: Typically, a parameter F is composed of basic features (such as input, output, non-determinism, probability). Modularity then means that the syntax/proof systems/algorithms for the logic of F are obtained compositionally from the syntax/proof systems/algorithms for the logics of the basic features.

Structuring the Seminar

When we planned the seminar, we envisaged six broad topics. We indicate which of the talks fall under which topic.

1. Category Theoretic Aspects of Coalgebraic Logic
2. Probabilistic Transition Systems
3. Stone Duality
4. Coalgebraic Logic, Automata Theory, Fixed Point Logics
5. Coalgebraic Logic for Structural Operational Semantics
6. Applied Coalgebraic Logic

Moss gave a presentation on new developments on the logic of recursion, which is one of the oldest topics in coalgebraic logic going back to the book Vicious Circles by Barwise and Moss (1996). New perspectives for coalgebraic logic where opened by the talks by Abramsky and Jacobs (quantum systems), and Pavlovic (security).

Chapter 3

Programming Languages, Compiler

3.1 SYNCHRON 2009

Seminar No. **09481** Date **22.11.–27.11.2009**
Organizers: Albert Benveniste, Stephen A. Edwards, Edward Lee, Klaus Schneider, Reinhard von Hanxleden

Synchronous languages have been designed to allow the unambiguous description of reactive, embedded real-time systems. The common foundation for these languages is the synchrony hypothesis, which treats computations as being logically instantaneous. This abstraction enables functionality and real-time characteristics to be treated separately, facilitating the design of complex embedded systems. Digital hardware has long been designed using the synchronous paradigm; our synchronous languages were devised largely independently and have placed the technique on a much firmer mathematical foundation.

Feedback from the user base and the continuously growing complexity of applications still pose new challenges, such as the sound integration of synchronous and asynchronous, event- and time-triggered, or discrete and continuous systems. This seminar aims to address these challenges, building on a strong and active community and expanding its scope into relevant related fields. This year's workshop includes researchers in model-based design, embedded real-time systems, mixed system modeling, models of computation, and distributed systems.

The seminar was successful in bringing together researchers and practitioners of synchronous programming, and furthermore in reaching out to relevant related areas. With a record participation in this year's SYNCHRON workshop of more than 50 participants and a broad range of topics discussed, the aims seem to have been well-met. The program of the seminar was composed of around 36 presentations, all of which included extensive technical discussions. The fields covered included synchronous semantics, modeling languages, verification, heterogeneous and distributed systems, hardware/software integration, reactive processing, timing analyses, application experience reports, and industrial requirements. The discussion identified and collected specific needs for future topics, in particular the integration of different models of computation.

The SYNCHRON workshop constitutes the only yearly meeting place for the researchers in this exciting field. The workshops on Synchronous Languages started in 1993 at Schloss Dagstuhl. Since then, the workshop has evolved significantly in its sixteen years of existence. One obvious change is the citizenship of its attendees, which has shifted from being largely French to being truly world-wide. But the biggest change is in its scope, which has grown to expand many languages and techniques that are not classically synchronous but have been substantially influenced by the synchronous languages' attention to timing, mathematical rigor, and parallelism. Also, while many of the most senior synchronous language researchers are still active, many younger researchers have also entered the fray and taken the field in new directions. We look forward to seeing where they take us next.

Chapter 4

Geometry, Image Processing, Graphics

4.1 Computational Geometry

Seminar No. **09111** Date **08.03.–13.03.2009**
Organizers: Pankaj Kumar Agarwal, Helmut Alt, Monique Teillaud

Computational Geometry Evolution

The field of computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric problems, which arise in a wide range of areas, including computer graphics, CAD, robotics computer vision, image processing, spatial databases, GIS, molecular biology, and sensor networks. Since the mid 1980s, computational geometry has arisen as an independent field, with its own international conferences and journals.

In the early years mostly theoretical foundations of geometric algorithms were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, researchers have started paying close attention to applications and implementations of geometric algorithms. Several software libraries for geometric computation (e.g. LEDA, CGAL, CORE) have been developed. Remarkably, these implementations emerged from the originally theoretically oriented computational geometry community itself, so that many researchers are concerned now with theoretical foundations as well as implementations.

Seminar Topics

The seminar focused on theoretical as well as practical issues in computational geometry. In the following, we list some of the currently most important topics in computational geometry, together with some of the leading researchers working in those areas whom were invited to this seminar:

- *Theoretical foundations* of computational geometry lie in combinatorial geometry and its algorithmic aspects. They are of an enduring relevance for the field, particularly the design and the analysis of efficient algorithms require deep theoretical insights. [Chazelle, Sharir, Welzl,...]

- Various *applications* such as robotics, GIS, or CAD lead to interesting variants of the *classical topics* originally investigated, including convex hulls, Voronoi diagrams and Delaunay triangulations, and geometric data structures. For example, pseudotriangulations, generalization of triangulations and developed in connection with visibility and shortest-path problems, have turned out to be useful for many other applications and are being investigated intensively. [van Kreveld, Mitchell, Streinu,...]

- Because of applications in molecular biology, computer vision, geometric databases, *shape analysis* has become an important topic. [Asano, Knauer]

- Another increasingly important application of computational geometry is *modeling and reconstruction of surfaces*. It brings about many interesting questions concerning fundamental structures like triangulations as well as new issues in *computational topology*. [Erickson,...]

- Massive geometric data sets are being generated by networks of sensors at unprecedented spatial and temporal scale. How to store, analyze, query, and visualize them has raised several algorithmic challenges. New computational models have been proposed to meet these challenges, e.g., streaming model, communication-efficient algorithms, and maintaining geometric summaries. [Arge, Efrat,...]

- *Implementation issues* have become an integral part of the research in computational geometry. Besides general software design questions especially *robustness* of geometric algorithms is important. Several methods have been suggested and investigated to make geometric algorithms numerically robust while keeping them efficient, which lead to interaction with the field of computer algebra, numerical analysis, and topology. [Everett, Lazard, Mehlhorn, Wolpert,...]

Participants

Dagstuhl seminars on computational geometry have been organized since 1990, lately in a two year rhythm, and always have been extremely successful and on a very high scientific level, possibly the highest of all meetings on computational geometry worldwide.

This year, 42 researchers from various countries and continents attended the meeting.

The feedback from participants was very positive.

Participants, especially junior researchers, appreciate the opportunity to meet leaders in the field and benefit from their expertise. Keeping the attendance small enough is a necessary condition for an easy communication and a good research atmosphere, but, having most leaders in the field still allows to invite some very promising younger people. This formula has been recognized as very successful for years.

4.2 Generalization of Spatial Information

Seminar No. **09161** Date **13.04.–17.04.2009**
Organizers: Sébastien Mustière, Monika Sester, Frank van Harmelen, Peter van Oosterom

From the early start of handling geo-information in digital environments, it has been attempted to automate the process of generalization of geographic information. Traditionally for the production of different map scale series, but more and more also in other contexts, such as the desktop/web/mobile use of geo-information, in order to allow to process, handle and understand possibly huge masses of data. Generalization is the process responsible for generating visualizations or geographic databases at coarser levels-of-detail than the original source database, while retaining essential characteristics of the underlying geographic information.

All current solutions for supporting different levels-of-detail are based on (static) copies that are (redundantly) stored at these levels. This makes dynamically adapting the map to new information and to the changing context of the user impossible. Besides the classic geo-information visualization requirement (supporting different scales), which has been solved only partly, there are also new requirements for generalization, making it even more difficult: it should be dynamic and suitable for progressive transfer. Furthermore, the objects to be visualized have expanded in dimension: the emerging 3D and temporal data. In order to make further progress in automated, machine generalization both the semantics of the spatial information and the user needs should be (further) formalized. Methods and techniques from the semantic web might be useful in this formalization and tools from knowledge engineering could be applied in the actual generalization based on reasoning. Interpretation of spatial constellations or situations is a process that is closely linked to human capabilities and can be formalized using formal semantics (OWL, ODM, etc.). Making implicit information explicit is needed not only for many spatial analysis problems, but also for aspects of information communication.

Spatial data also pose exciting questions for the algorithms and data structuring communities. It is vital that computational geometers meet with the spatial data community to exchange ideas, pose problems and offer solutions. Most algorithmic problems arising in that field are indeed geometric. In this context it must be noticed that the focus is more and more on 3D (and 3D plus time) geometric computations. In this respect, generalization operations and the resulting data have to be understood as processes, which will allow a broader and more flexible usage and re-generalization when changes in reality have occurred.

For a mass market (e.g. consumers of mobile maps) the human factors aspect is very important. The currently available (often mobile maps) solutions still have insufficient user-interfaces. Extremely important is the issue of context as the user gets ‘lost’ very easily on the small mobile displays when zooming and panning. Based on a selection of use cases (navigation, tourist support, etc.), User-Centered Design techniques should be applied to evaluate the interaction and the quality of the maps.

In this context, the main goal of the seminar was to bring together experts from digital cartography, knowledge engineering, computational geometry, computer graphics and cog-

nitive science, to lead to a fruitful exchange of different – but very close – disciplines and hopefully to the creation of new collaborations.

The seminar brought together experts from digital cartography, knowledge engineering, computational geometry and cognitive science, with the goal to lead to a fruitful exchange of different – but very close – disciplines and to the creation of new collaborations.

As a kind of conclusion: The idea of bringing together researches from various community, from semantic web specialists to geometry specialists has been widely thought of as fruitful. This led to many discussions on links between those fields of research. Some particular issues to be tackled appeared in many questions and discussions, and may lead to further research, like:

- How cognitive and semantic aspects could be integrated in the generalisation process that is (may be too much) geometry-oriented?
- How to handle the notions of fuzziness or uncertainty in semantic web techniques?
- Shall generalization be more task-oriented? How to do that?
- How to handle the notions of time and updates in both domains?

4.3 Scientific Visualization

Seminar No. **09251** Date **14.06.–19.06.2009**
Organizers: David S. Ebert, Eduard Gröller, Hans Hagen and Arie Kaufman

Resulting from a growth in data set size, complexity, and number of covered application areas, modern Scientific Visualization combines research from a wide variety of theoretical and practical fields such as mathematics, physics, biology and computer science. These research efforts yield a large number of different analysis, processing, and visualization techniques, allowing the efficient generation and presentation of visual results. This in turn directly contributes to the way domain experts are able to deduce knowledge from abstract data.

Emphasizing the heterogeneity of this research field, the Dagstuhl Scientific Visualization Seminar 2009 focused on a wide range of visualization topics such as "Knowledge Assisted Visualization", "Visual Exploration Environment", "Biomedical Visualization", and "Visualization of Vector- and Tensorfields". The seminar aimed to provide an open and international environment for the discussion of recent trends, breakthroughs and future directions of research in the area of visualization, fostering scientific exchange and collaboration among researchers of the Sci-Vis community and identifying new research directions.

In the course of the seminar, leading international scientists presented state-of-the-art summaries as well as novel research results and ideas. Among the discussed key topics were:

Interaction Techniques/Frameworks
To efficiently perform visual data analysis, end users and domain experts need not just be presented with visualization results, but have to be offered intuitive and efficient real-time interaction techniques and frameworks. User-centered approaches demonstrate, how human factors can influence the way data is processed and presented. Presentations and results from this seminar illustrated and devised methods for interactive data exploration and analysis.

Feature Definition and Extraction/Reconstruction
New data types and application fields require new types of features, novel extraction techniques and visualization algorithms. Work from a broad context of feature extraction and reconstruction in areas such as scalar-, vector- and tensorfield visualization was presented in the course of this seminar.

Visualization Metaphors
Complex abstract data properties can only be visualized with suitable visualization metaphors. The conception and design of such visualization metaphors is a key problem in data visualization. To this end, several results of this seminar aimed towards answering abstract questions about "system safety", "document history", and "algorithmic space".

Optimization Techniques
As existing work from the field of visualization is adapted to new application areas or visualization problems, an increase in size, structure or complexity of the given data necessarily leads to the development of optimized algorithms. This seminar identified algorithms and data structures for performance and accuracy improvement in key areas of scientific visualization such as (vector) field analysis.

Besides these topics, participants gave valuable presentations about conceptual, philosophical and psychological questions in visualization regarding the impact and benefit of user-centered approaches, research classification and other topics.

The productive setting at Dagstuhl made it possible, that a selection of ideas presented at this seminar as well as scientific results of this gathering are made available as Proceedings.

4.4 New Developments in the Visualization and Processing of Tensor Fields

Seminar No. **09302** Date **19.07.–24.07.2009**
Organizers: Bernhard Burgeth, David H. Laidlaw

This Dagstuhl Seminar was concerned with the visualization and processing of tensor fields, like its two predecessors: seminar 04172 organized by Hans Hagen and Joachim Weickert in April 2004, and the follow-up seminar 07022 in January 2007 with David Laidlaw and Joachim Weickert as organizers. Both earlier meetings were successful, resulting in well received books and riggering fruitful scientific interaction and exchange of experience

across interdisciplinary boundaries. We believe that the 2009 seminar will prove to have been equally successful.

Our main goal was to bring together researchers from rather different fields, ranging from visualization and image processing to applications in structural mechanics, fluid dynamics, elastography, and numerical mathematics. The scientific output will be collected in a post-proceedings volume currently being produced.

The Dagstuhl survey results suggest a success. All respondents said they would come to another seminar. All of the content-related responses were higher than those for the comparison group of seminars. In fact, no responses were lower than the comparison-group mean.

This third-in-a-series seminar was a great success. The attendees report high marks, and the expected book is starting off with strong involvement from the authors. We have hopes of proposing another in this series with two new organizers and one continuing one. The field continues to expand and mature, and Dagstuhl seminars continue to help make that process a robust one.

4.5 Geometric Networks, Metric Space Embeddings, Spatial Data Mining

Seminar No. **09451** Date **01.11.–06.11.2009**
Organizers: Gautam Das, Joachim Gudmundsson, Rolf Klein, Christian Knauer, Michiel Smid

This seminar has brought together scientists from three different communities (geometric networks, metric space embeddings, and spatial data mining) who have numerous interests in common, all of which are related to distance problems. The seminar was a continuation of the Dagstuhl seminar 06481 (Geometric Networks and Metric Space Embeddings) which was held in 2006. The main purpose of the current seminar was to continue, and intensify, the cooperation between the geometric network and the metric space communities. This time, we invited people from spatial data mining to provide the extra application stimulus.

A *geometric network* is a graph mapped into the Euclidean plane or a Euclidean space of low dimension. It is said to be a *spanner* if the network distance between any pair of nodes is bounded by a constant times their Euclidean distance. Geometric networks are the backbone of any model for the flow of goods, traffic or information. They also play an important role in telecommunication, VLSI design, motion planning (robotics), pattern matching, data compression, bio-informatics (gene analysis), and sensor networks. One is interested in spanners with other useful properties such as a linear number of edges, small total edge length, small node degree, few crossings, or small link diameter. Apart from these applications, geometric spanners have had great impact on the construction of approximation algorithms, e.g., for the traveling salesperson problem.

The similarity between individual objects of a given finite domain becomes apparent if the objects can be represented by points in the plane, or in 3-space, in such a way that the

Euclidean distance between two points corresponds to the similarity of the objects they are associated with. The question when such representations exist has led to the theory of embedding finite metric spaces into normed vector spaces. It is of particular interest for storage, visualization, and retrieval of high-dimensional data, e.g., in information retrieval.

Both problems (metric space embedding and spanner construction) have received a lot of attention during the last 10-15 years, within their respective communities. Indeed, the first monograph on spanners (co-authored by the 5th organizer) has been published meanwhile, and metric space embeddings have been addressed in several books and book chapters. In both cases, we are applying two different metrics to the point pairs of the same set, and we are looking for the maximum (or minimum) ratio of all values. In metric space embeddings we compare the measure of similarity of the objects against the (Euclidean) distance of their associated points, and the maximum ratio is called the distortion of the embedding. In a spanning geometric network, we compare the shortest path distance in the network against the Euclidean distance between two points; here, the extremal ratio is called dilation or stretch factor. In both areas many questions are open. For example, it is not known how to construct the triangulation of minimum dilation over a given point set, with or without extra Steiner points allowed.

Data mining can be seen as the science of extracting useful information from large data sets or databases. It is the principle of sorting through large amounts of data and finding relevant information. It is usually used by business intelligence organizations and financial analysts, but it is increasingly used in the sciences to extract information from the enormous data sets generated by modern experimental and observational methods. In many applications of data mining, the high dimensionality of the data restricts the choice of data processing methods. Such application areas include the analysis of market basket data, text documents, image data and so on; in these cases the dimensionality is large due to either a wealth of alternative products, a large vocabulary, or the use of large image windows, respectively. A common tool used to develop efficient algorithms is to reduce the number of dimensions. A statistically optimal way of dimensionality reduction is to project the data onto a lower-dimensional orthogonal subspace that captures as much of the variation of the data as possible. The best (in mean-square sense) and most widely used way to do this is principal component analysis (PCA); unfortunately, it is quite expensive to compute for high-dimensional data sets. A computationally simple method of dimensionality reduction that does not introduce a signifficant distortion in the data set is random projection. Here, the original high-dimensional data is projected onto a lower-dimensional subspace using a random matrix whose columns have unit lengths. Random projection has been found to be a computationally efficient, yet sufficiently accurate method for dimensionality reduction of high-dimensional data sets.

The seminar's aim was at crossfertilization between the three communities. The main results of the seminar can be summarized as follows.

- During the seminar, it became clear that methods developed in the theory of finite metric spaces help in analyzing geometric networks. Conversely, the algorithmic techniques developed in geometric network design are of interest to people working on embedding problems.

- There was a fruitful exchange of ideas which stimulated interesting discussions and future cooperations.
- The seminar will advance a comparative theory of distance measures.
- The knowledge gained during the seminar will help in reducing the complexity of high-dimensional data, as is important in data mining and related areas.

Chapter 5

Artificial Intelligence, Computer Linguistic

5.1 Normative Multi-Agent Systems

Seminar No. **09121** Date **15.03.–20.03.2009**
Organizers: Guido Boella, Pablo Noriega, Gabriella Pigozzi, Harko Verhagen

NorMAS-09 was the 4th NorMAS event, the second one at Dagstuhl, where the community working on normative multiagent systems had again the opportunity to meet and discuss.

The seminar was organized around short paper presentation sessions, from 9hrs to 15.30hrs and the remaining time devoted to group discussion. Each paper presentation was commented by a discussant —who had read the paper and sent comments to the authors beforehand– and then open to general debate. Priority was given to young and/or female researchers as presenters and discussants.

The papers presented were collected in DROPS proceedings before the meeting (to allow discussant to review them and to make them available to the other participants). A WIKI page was set up for participants to express the issues and technical concerns they find worth discussing in the seminar, and to give access to other seminar-related material.

Afternoon group discussion sessions were structured in two stages. The first one divided the participants into four groups and each group took the current concerns of the NorMAS community –previously gathered in the seminar WIKI– along with the actual contribution of the papers, and attempted to match concerns and contributions with the challenges of the field that had been proposed in the previous Dagstuhl seminar (NorMAS-07). The second stage consisted of a report of each group and a collective discussion. Slides of paper presentations, commentaries and of group discussions are being uploaded on the page containing the materials of the seminar.

Given that the Dagstuhl seminar format allowed more time and flexibility than a regular workshop, we were able to invite a participation of the European project COST Action "Agreement Technologies" (www.agreement-technologies.eu), that has Norms in Multi

Agent Systems as the topic of one of its five working groups, of which many of the NorMas-09 participants are members or supporters. The COST action organized a panel discussion on normative aspects of MAS in which the working group leader Cristiano Castelfranchi was accompanied by Thomas Agotnes, Guido Boella, Pablo Noriega and Timothy Norman. The COST action contributed further to the seminar by funding the participation of Bastin Savarimutu, a young researcher from University of Otago in New Zealand.

In addition to the presentations, group discussions and the Agreement Technologies panel, the seminar hosted an evening session about projects on Normative Multi Agent Systems where participants are currently involved. Although a similar initiative had taken place at NorMAS Dagstuhl Seminar 2007, the situation was then still too preliminary. In this session, the survey of projects showed that the area is in a healthy development. More than 15 projects are being funded at different levels: from international (ITA), European (ALIVE, COST Action Agreement technologies) as well as national and regional ones (ICT4LAW, Agreement technologies, Llibre Blanc de la Mediacio en Catalunya). A page of the WIKI website will include links to such projects. The session served also to discuss project proposals. From the discussion it became apparent that the ICT call of FP7 does not seem to hold a clearly defined spot for NorMAS topics. However, a new opportunity was identified: presenting an EUROCORE theme proposal. The goal of this initiative is to show to the Commission the existence of a large community working on norms: 44 researchers participated to the seminar, but more than 80 people had been invited, and many of them have large groups of people working also on norms.

Since the community has very different lines of research and to an extent belongs to different disciplines like sociology, law, computer science, economics, one of the goal of the seminar was to strengthen the roots of the community. For this reason, and in line with the COST Action Agreement technologies supporting the seminar, we agreed to identify the most significant papers for the NorMAS community with the likely intent to produce a seminal book around them. Thus, the participants have been invited to send a list of the papers which they believe to be most significant for normative multiagent systems and the organizers will sort out the results to get this initiative going.

As with the previous Dagstuhl seminar on NorMAS, the best revised papers will be published in specialized journals. In this occasion, a special issue of the Journal of Algorithms in Cognition, Informatics and Logic of Elsevier (ISI impact factor 1.2, A class in the Italian GRIN rating) and a special issue of the Artificial Intelligence and Law Journal will be produced after a new round of anonymous peer reviews.

The interest of the community present at Dagstuhl seems to have shifted from applications like Second Life to other domains like normative compliance, norm detection and norm enforcement. The amount of application papers was a bit higher than at NorMAS07 at Dagstuhl, partially due to new research groups joining the community to present their work and the interest of Web2.0 based applications in the area of norms.

New methodologies have been introduced like argumentation theory, while other more traditional ones like deontic logic have reduced their importance. Still some uncertainty exists about which is the right methodology: for example traditional modal approaches have shown some limitations. The meta-question being "is there a "right" methodology?" probably has the answer "no, it depends on the research question and focus". We are

happy to see an increasing spread in the type of questions and answers being put forward within the community, while at the same time the discussions suggest that the shared vocabulary and views do not suffer from this.

We are planning to continue on the NorMAS journey by organizing NorMAS2010 (probably at AISB2010) and NorMAS2011 (at Dagstuhl preferably) and setting up a NorMAS steering group.

5.2 Semantic Web, Reflections and Future Directions *(Dagstuhl Perspectives Workshop)*

Seminar No. **09271** Date **28.06.–03.07.2009**
Organizers: John Domingue, Dieter Fensel, James A. Hendler, Rudi Studer

With an ever increasing amount of data being stored and processed on computers, and the ubiquitous use of the Web for communication and dissemination of content, the world contains a vast amount of digital data that is growing ever faster. The available data is increasingly used to gain insights for science and research, to create commercial value, and to hold governments accountable. Semantic Web technologies for supporting machine-readable content aim at facilitating the processing and integration of data from the open Web environment where large portions of the publicly available data is being published. Since the first Dagstuhl seminar "Semantics for the Web" in 2000 the amount of machine-readable data on the Web has exploded, and Semantic Web technologies have matured and made their way from research labs and universities into commercial applications.

In this Perspectives Workshop participants from academia, industry, and government presented and discussed further directions for Semantic Web research. In general, the field of Semantic Web research has matured in the last decade. One indication for that, among others, is the discussion of advanced issues such as scalability. More data is becoming available and vocabulary management is being operationalised, and research on provenance tracking and technologies and methods addressing privacy concerns has commenced. How users can appropriately interact with the flood of data and leverage the data to satisfy information needs or gain insight is still an open question. Some prominent areas for applications of Semantic Web technologies discussed at the workshop are e-Science and mobile and sensor networks.

5.3 Cognition, Control and Learning for Robot Manipulation in Human Environments

Seminar No. **09341** Date **16.08.–21.08.2009**
Organizers: Michael Beetz, Oliver Brock, Gordon Cheng, Jan Peters

High performance robot arms are faster, more accurate, and stronger than humans. Yet many manipulation tasks that are easily performed by humans as part of their daily life are well beyond the capabilities of such robots. The main reason for this superiority is that humans can rely upon neural information processing and control mechanisms which are tailored for performing complex motor skills, adapting to uncertain environments and to not imposing a danger to surrounding humans. As we are working towards autonomous service robots operating and performing manipulation in the presence of humans and in human living and working environments, the robots must exhibit similar levels of flexibility, compliance, and adaptivity.

The goal of this Dagstuhl seminar is to make a big step towards pushing robot manipulation forward such that robot assisted living can become a concrete vision for the future. In order to achieve this goal, the computational aspects of everyday manipulation tasks need to be well-understood and the interaction of perceptual, learning, reasoning, planning, and control mechanisms thoroughly investigated. The challenges to be met include cooperation with humans, uncertainty in both task and environments, real-time action requirements, and the use of tools. The challenges cannot be met by merely improving the software engineering and programming techniques. Rather the systems need built-in capabilities to deal with these challenges. Looking at natural intelligent systems, the most promising approach for handling them is to equip the systems with more powerful cognitive mechanisms.

The potential impact of bringing cognition, control and learning methods together for robotic manipulation can be enormous. This urge for such concerted approaches is reflected by a large number of national and international research initiatives including the DARPA cognitive systems initiative of the Information Processing Technology Office, various integrated projects funded by the European Community, the British Foresight program for cognitive systems, huge Japanese research efforts, to name only a few.

As a result, many researchers all over the world engage in cognitive system research and there is need for and value in discussion. These discussions become particularly promising because of the growing readiness of researchers of different disciplines to talk to each other.

Early results of such interdisciplinary cross-fertilization can already be observed and we only intend to give a few examples: Cognitive psychologists have presented empirical evidence for the use of Bayesian estimation and discovered the cost functions possibly underlying human motor control. Neuroscientists have shown that reinforcement learning algorithms can be used to explain the role of dopamine in the human basal ganglia as well as the functioning of the bee brain. Computer scientists and engineers have shown that the understanding of brain mechanisms can result into reliable learning algorithms as well as control setups. Insights from artificial intelligence such as Bayesian networks and the associated reasoning and learning mechanisms have inspired research in cognitive psychology, in particular the formation of causal theory in young children.

These examples suggest that (1) successful computational mechanisms in artificial cognitive systems tend to have counterparts with similar functionality in natural cognitive systems; and (2) new consolidated findings about the structure and functional organization of perception and motion control in natural cognitive systems indicate in a number of cases much better ways of organizing and specifying computational tasks in artificial

cognitive systems.

5.4 Information Processing, Rational Belief Change and Social Interaction

Seminar No. **09351** Date **23.08.–27.08.2009**
Organizers: Giacomo Bonanno, James Delgrande, Hans Rott

The study of the formal aspects of information processing, belief formation and rational belief change is of central importance in a number of different fields. A new field of research, called Social Software, maintains that mathematical models developed to reason about the knowledge and beliefs of a group of agents can be used to deepen our understanding of social interaction and aid in the design of successful social institutions. Social Software is the formal study of social procedures focusing on three aspects: (1) the logical and algorithmic structure of social procedures (the main contributors to this area are computer scientists), (2) knowledge and information (the main contributors to this area are logicians and philosophers), and (3) incentives (the main contributors are game theorists and economists). Similarly, the most important question in Game Theory is how to rationally form a belief about other players' behavior and how to rationally revise those beliefs in light of observed actions. Traditionally Game Theory has relied mostly on probabilistic models of beliefs, although recent research has focused on qualitative aspects of belief change. A new branch of logic, called Dynamic Epistemic Logic, has emerged that investigates the epistemic foundations of game theory from the point of view of formal logic. There are various newly emerging links between the research areas mentioned above.

The purpose of the Workshop was to bring together researches from all these different areas and to promote an exchange of ideas and cross-fertilization between different fields. These researchers normally do not meet together.

Two very successful workshops with similar objectives took place at Schloss Dagstuhl in August 2005 and August 2007 (Seminars 05321 and 07351). Researchers from different fields (logicians, computer scientists, philosophers and economists) participated in these workshops and the anonymous surveys collected at the end gave enthusiastic evaluations of the events.

We saw the Dagstuhl Workshop as providing a forum where researchers in three broad areas (philosophy and logic, artificial intelligence and computer science, and economics and game theory) could address highly related (in some cases, the same) problems, in which work in one area could benefit research in another.

We found the Workshop successful, especially on the following two achievements: first, the seminar made participants aware of a commonality of interests across different disciplines; second, it suggested new directions for research that will probably be taken up by researchers in the next couple of years.

5.5 From Form to Function

Seminar No. **09431** Date **18.10.–23.10.2009**
Organizers: Darius Burschka, Heiner Deubel, Danica Kragic, Markus Vincze

At present we are on the verge of a new era when technical systems expand from typical industrial applications with pre-programmed, hard-wired behaviors into everyday life situations where they have to deal with complex and unpredictable events. The increasing demand for robotic applications in dynamic and unstructured environments is motivating the need for novel robot sensing and adaptable robot grasping abilities. The robot needs to cope with a wide variety of tasks and objects encountered in open environments. Since humans seem to have no difficulty to estimate a rough function of an object and to plan its grasping solely from the visual input, robot vision plays a key function in the perception of a manipulation system.

Our hypothesis is that the form and shape of objects is a key factor deciding upon actions that can be performed with an object. Psychophysical studies with humans confirm that affordance of grasping includes information about object orientation, size, shape/form, and specific grasping points. Affordances are discussed as one ingredient to close the loop from perception to potential actions.

The aim of this seminar is to bring together researchers from different fields related to the goal of advancing our understanding of human and machine perception of form and function. We set out to explore findings from different disciplines to build more comprehensive and complete models and methods. Neuroscientists and experimental psychologists will provide initial conceptual findings on the selective nature of sensor processing and on how action-relevant information is extracted. Cognitive scientists will tackle the modeling of knowledge of object function and task relations. Computer vision scientists are challenged to develop procedures to achieve context-driven attention and a targeted detection of relevant form features. All participants will profit from the ideas and findings in the related disciplines and contribute towards establishing a comprehensive understanding of brain and computing processes to extract object function from form features.

- Computer vision and perception needs to detect relevant features and structures to build up the shape/form of objects, to determine their orientation and size, and to define good grasping points. Currently, appearance has been successfully used for recognizing objects and codebooks of features assist in object categorization. Our goal is to move the data abstraction higher to define object function from the perception of edges, contours, surface properties, and other structural features, which still remains less explored. A main task of the workshop is to bring the key experts together to discuss how to advance the state of the art.

- Attention is the mechanism to enable fast and real-time responses in humans. Studies with humans show that grasping can be performed independent of object recognition. Hence it is timely to investigate how this direct link or affordances can be modeled and replicated for exploitation in robotic and cognitive systems.

- Prediction and the integration of bottom-up and top-down data flow is often discussed. Primate vision has been largely studied based on passively recording neuron functions when observing patterns (bottom-up stream). Only recently the importance of top-down triggers has been more closely shown. For example, from the retina but other brain areas including what is thought to be higher brain regions. Recent neuro-scientific findings state that predictions are a primary function of these connections. This indicates that the human brain uses predictions to focus attention, to exploit task and context knowledge, and hence scale an otherwise too wide space of inputs. For example, prediction indicates how a shape will be perceived when a certain action is executed on the target object. The task will be to identify what is the relevant information and how can it be computed in a machine vision system.

- Finally, humans seem to build up extensive knowledge about typical shapes and forms of whatever is seen in daily life. Seeing a partly occluded object often immediately triggers the respective model to complete the shape. Also in grasping it has been found that the grasping point on the backside of an object is typically invisible but it is inferred from a symmetry assumption. The search for objects (say cups) is focused on horizontal surfaces and exploits knowledge about object category to look in a kitchen rather than in the garage. Recent work created first databases and ontologies to describe such knowledge, yet it remains open to fuse these developments with the results listed above.

- In summary, the seminar brought together scientists from disciplines such as computer science, neuroscience, robotics, developmental psychology, and cognitive science to further the knowledge how the perception of form relates to object function and how intention and task knowledge (and hence function) aids in the recognition of relevant objects.

Chapter 6

Software Technology

6.1 Software Service Engineering

Seminar No. **09021** Date **04.01.–07.01.2009**
Organizers: Willem-Jan van den Heuvel, Olaf Zimmermann, Frank Leymann, Tony Shan

Seminar Topics and Objectives

Service-oriented architecture (SOA) as an architectural style based on common principles and patterns such as Business Process Choreography and Enterprise Service Bus (ESB) allows service engineers to effectively (re-)organize and (re-) deploy executable business processes, functional components, and information assets as business-aligned and loosely-coupled software services. SOA is unique in that it aims at unifying various related, yet up to now largely isolated domains such as business process management, distributed computing, enterprise application integration, software architecture, and systems management.

Given the loosely-coupled, heterogeneous, and dispersed nature of SOA, several of the key assumptions underlying traditional approaches to software engineering are being challenged; consequently, conventional software engineering methods and tools have to be carefully reevaluated and possibly extended to be applicable to analysis, design, construction, and delivery of software services. Due to the continuing evolution of SOA, SOA research so far has been mostly focused on certain parts of the service lifecycle, such as runtime service infrastructure and middleware. There is a lack of comprehensive methods and tools consistently supporting all phases of software service engineering ranging from analysis to implementation and evolution. Such methods and tools must be grounded both in scientific foundations and in industrial best practices. It was the overall goal of this seminar to assess existing methods, techniques, heuristics, and practices from related fields such as requirements engineering, software engineering, Object-Oriented Analysis and Design (OOAD), Component-Based Development (CBD), and method engineering to harness SOA methods and tools and to define a road map for the creation of a unified software service engineering method. More precisely, the first objective of the workshop was to understand assumptions and impact of emerging runtime platform models on the engineering process, e.g., SOA principles and patterns such as loose coupling and programming

without a call stack, ESB and service composition, Software as a Service (SaaS) and cloud computing, Web 2.0 and mashups, as well as mass programming. Based on the results of this analysis activity, the second objective of the workshop was to define distinguishing characteristics of Software Service Engineering (SSE) and to assess the state of the art in SOA and service design methods. The third and last goal was to identify focus areas for future work and a roadmap for SSE. In particular, the following three questions were addressed: Are new methods and tools required? How can the existing body of knowledge in software engineering and SOA design be extended? Is method unification a la Unified Modeling Language (UML) and Unified Process (UP) desirable and feasible?

Seminar Organization

Participating communities. With this seminar we brought together researchers and practitioners from various industrial domains and research areas that work in the emerging field of software service engineering. In particular, we established linkages and enduring collaborations between the following three communities that operated in isolation before:

1. Requirements and software engineering including patterns.
2. SOA middleware and platform standards.
3. Industrial adopters of SOA.

41 participants from 10 countries attended the seminar; industry participation was in the range of 40% to 60% (depending on how industrial research labs are counted). Areas of interest and expertise varied from business process modeling to SOA design principles, patterns, and platform, but also method engineering, software architecture, testing, legacy system analysis, semantic Web, and software product lines.

Conclusion and Outlook

SOA-enabled applications can be developed and evolved by applying aging software engineering paradigms, notably CBD and OO; however, the key advantages of SOA cannot be fully exploited when doing so. The main reason for this is that conventional software engineering paradigms typically adopt the closed world assumption, hypothesizing that applications have clear boundaries, and will be executed in fully controlled, relatively homogeneous, predictable and stable execution environments. This thesis is backed up by conclusions drawn from a decade-to-decade analysis of software engineering by Barry Boehm.

Instead, we claim that for SOA to be applied successfully, SSE has to embrace the open-world assumption, in which software services are composed in agile and highly fluid service networks that are in fact systems of software-intensive systems operating in highly complex, distributed, and heterogeneous execution environments. In addition, the service networks that are designed based on this assumption have to be continuously (re-)aligned

with business processes, and vice versa. Adoption of the open-world assumption is reflected in the three types of SSE tenets: architecture, process, and engineering.

Based on the research reported, we came up with an initial definition of SSE as: Software service engineering entails the science and application of concepts, models, methods, and tools to define, design, develop/source, integrate, test, deploy, provision, operate, and evolve business-aligned and SOA-enabled software systems in a disciplined and routinely manner. Clearly, SSE will benefit from timeless generic principles and lessons learned from its elderly parent software engineering; however, we herein argue that aging computing model specific principles and practices, e.g., distributed component technology, are in clear need for revision given the specific nature of SOA. In our view, SSE will be based on standards and will be frequently realized with Web services. Specifications such as SOAP, WSDL, BPEL, WS-Policy, and WS-Agreement already constitute the first step to realize the technical aspects in some of the SSE tenets, including engineering tenets 1, 2, 4, and 5. Other architectural styles and technology paradigms can also be used to realize software services. However, further research is required to more effectively satisfy the open-world assumption. This has also been reflected in the outcome of the brainstorm on the key open research challenges.

The results of this seminar are core results in nature. During the seminar it became clear that the discipline of software service engineering is still in its embryonic phase, and further work is required in several directions. Firstly, the list of tenets has to be further explored, validated, and potentially refined or expanded. The presented list is derived from a literature survey, as well as expertise and experience from real-world SOA projects and discussions with leading industry experts and renowned researchers in the field of software engineering, software patterns and SOA; however, more case studies have to be analyzed critically to further validate this initial list.

As a follow-up activity, we have published the results of this seminar in an ICSE workshop paper. The workshop paper extends the discussion in this executive summary and provides an example which illustrates the difference between SSE/SOA and traditional software engineering disciplines.

6.2 Self-Healing and Self-Adaptive Systems

Seminar No. **09201** Date **10.05.–15.05.2009**
Organizers: Artur Andrzejak, Kurt Geihs, Onn Shehory, John Wilkes

Summary

During the last few years, the functionality and complexity of software and systems in enterprise and non-commercial IT environments have increased a great deal. The result is soaring system management costs and increased likelihood of failures. There is a common understanding across researchers and engineers alike that enhancing systems with self-management capabilities is a promising way to tackle these challenges. These *self-managing capabilities* - frequently summarized under the term *autonomic computing* -

include self-configuration, self-healing, self-optimization and self-protection. Recent years have brought a notable increase in related research activities, the driving forces being major IT players including IBM, HP, SUN, and Microsoft.

The Dagstuhl seminar "Self-Healing and Self-Adaptive Systems" focused on self-healing IT systems in the broader context of self-adaptive systems. Self-healing refers to the automatic detection of failures and anomalies and their subsequent correction in a temporary or a permanent manner. Self-healing systems are of particular interest as they directly impact improvements in dependability. Self-adaptive systems are ones that monitor their execution environment and react to changes by modifying their behavior in order to maintain an appropriate quality of service. Obviously, there is a substantial intersection between self-healing and self-adaptiveness: self-healing systems may be viewed as a special kind of self-adaptive systems.

Goals and Content of the Seminar

The overall goals of the seminar were

- to bring together experts from various disciplines and organisations for exchanging different viewpoints on the state of the art of methods and technologies for designing, implementing and evaluating of self-healing and self-adaptive systems,
- to foster open discussions on selected topics of the design space of such systems, and
- to facilitate community building in this increasingly important subject area.

In the invitations to the seminar participants three research fields were suggested in order to provide some structure for the presentations and discussions: fault detection and diagnosis, recovery and repair techniques, and frameworks and architectures for self-adapting systems. In order to establish a link between industrial practice and academic research, two focused application-oriented topics were intended to complement the seminar.

Conclusion

The Self-Healing and Self-Adaptive Systems seminar was a fertile meeting in which a diverse population of researchers have met. It included industry and academia, senior and junior researchers, multinational representation, and people coming from several disciplines. This diversity resulted in interesting and useful discussions, new understandings of the fundamental concepts and problems in the field, and in new collaborations on an array of problems which were not well defined or identified prior to this seminar.

Several work groups during the seminar not only generated new insights into specific topics in the field of self-healing and self-adaptive systems, but also initiated ongoing joint work, with group members continuing the work they started at the seminar.

The seminar included multiple presentations and discussions. Technical issues included all elements of the self-healing cycle, including monitoring, detection and diagnosis; recovery

and repair techniques; testing, quality trust issue; and, architectures, infrastructure and use cases. The participants identified the need for better terminology and taxonomy for the field. They further indicated the need for case studies and benchmarks. Several participants stressed the need for trustworthy solutions. It was widely agreed that the potential of self-healing and self-adaptive systems is high, even though much of the existing work in this field is rather academic in nature, and industrial take-up has been relatively slow, with a few notable exceptions.

This seminar clearly illustrated the diversity, relevance, and fertility of the topics we presented and discussed. The intensity of the participants' involvement leads us to believe that the interactions fostered by the seminar will generate a lot of follow-up research, and eventually lead to practical use as well.

6.3 Design and Validation of Concurrent Systems

Seminar No. **09361** Date **30.08.–04.09.2009**
Organizers: Cormac Flanagan, Susanne Graf, Madhusudan Parthasarathy, Shaz Qadeer

While concurrency has always been central to embedded and distributed computing, it has recently received increasing interest from other fields related to software engineering such as programming languages, compilers, testing, and verification. This recent interest has been fuelled by a disruptive trend in the evolution of microprocessors — the number of independent computing cores will continue to increase with no significant increase in the speed of each individual core. This trend implies that software must become increasingly concurrent in order to exploit current and future hardware.

At the same time, we are seeing an unprecedented penetration of embedded and distributed systems into everyday human life. Embedded devices, such as cell phones and media players, are ubiquitously used for communication and entertainment, and distributed control systems in cars and airplanes are increasingly safety-critical. Today, systems and software engineers face the challenging task of developing efficient and reliable software for concurrent, embedded, distributed, and multi-core platforms.

The presence of concurrency in a system severely increases its complexity due to the possibility of unexpected interactions among concurrently-executing components. System designers are invariably forced to make trade-offs between productivity, correctness, and performance. Current practice includes "correct-by-construction" design methods that yield safe implementations; these implementations are unlikely to be the most efficient. Conversely, highly flexible design methods can yield efficient distributed or multithreaded implementations; these methods are labor intensive and may require expensive post-design validation. We believe these two approaches delimit a continuous spectrum of design and validation techniques. It is important to develop techniques that provide a principled but pragmatic tradeoff between the rigidity of "correctness-by-construction" and the difficulty of post-hoc verification of arbitrary systems.

This workshop brought together academic and industrial researchers who are interested in design and validation techniques for concurrent systems. We had a broad participation

reflecting the various approaches to the problem, including language design, compiler construction, program analysis, formal methods, and testing. We believe this mix of participants generated interesting and lively discussions. Concretely, we addressed the following set of inter-related questions during the seminar:

- Specification and programming languages: How can a programmer specify correctness properties of a concurrent system? What are the right idioms for reasoning about concurrent programs? What concurrency-control mechanisms should be provided by the programming language? How do we enable programmers to write well-reasoned code that can be compiled for efficient execution on a multi-core platform? What kind of abstractions from the hardware/OS/runtime are useful and efficient?
- Design methods: How should a programmer choose the right design approach given the constraints, such as quality-of-service and reliability, that may be imposed on a given application domain? What are common patterns of non-interference, e.g. race-freedom, atomicity, and determinism, that help programmers avoid common concurrency-related pitfalls?
- Validation: How do we verify applications built using a given set of concurrency primitives? How do we verify implementations of algorithms realizing these primitives? How do we design efficient algorithms for verifying various forms of non-interference, and for explaining existing interference in terms understandable to the programmer? How do we test concurrent applications that may exhibit a high degree of, possibly uncontrollable, nondeterminism?

6.4 Refinement Based Methods for the Construction of Dependable Systems

Seminar No. **09381** Date **13.09.–18.09.2009**
Organizers: Jean-Raymond Abrial, Michael Butler, Rajeev Joshi, Elena Troubitsyna, Jim C. P. Woodcock

With our growing reliance on computers, the total societal costs of their failures are hard to underestimate. Nowadays computers control critical systems from various domains such as aerospace, automotive, railway, business etc. Obviously, such systems must have a high degree of dependability – a degree of trust that can be justifiably placed on them. Although the currently operating systems do have an acceptable level of dependability, we believe that they development process is still rather immature and ad-hoc. The constantly growing system complexity poses an increasing challenge on the system developers and requires significant improvement on the existing developing practice. To address this problem, we investigated how to establish a set of refinement-based engineering methods that can provide the designers with a systematic methodology for development of complex systems.

The seminar brought together academicians that are experts in the area of dependability and formal methods and industry practitioners that are working on developing dependable systems. The industry practitioners have described their experience and challenges posed by formal modeling and verification. The academicians tried to address these challenges while describing their research work. The seminar proceeded in a highly interactive manner and provided us with an excellent opportunity to share experience.

One of the outcomes of that seminar was the identification of the following list of challenging issues faced by industrial users of formal methods:

- Team-based development
- Dealing with heavy model re-factoring
- Linking requirements engineering and FMs
- Abstraction is difficult
- Refinement strategies are difficult to develop
- Guidelines for method and tool selection
- Keeping models and code in sync
- Real-time modelling
- Supporting reuse and variants
- Proof automation
- Proof reuse
- Handling complex data structures
- Code generation
- Test case generation
- Handling assumptions about the environment

The seminar has encouraged knowledge transfer between several major initiatives in the area of formal engineering of computer-based systems. We have got a good understanding of the advances made within the EU-funded project Deploy "Industrial deployment of system engineering methods providing high dependability and productivity". The project aims at integration of formal engineering methods into the existing development practice in such areas as automotive industry, railways, space and business domains. The participants described advantages and problems of refinement-based development using Event-B and Rodin tool platform. The advances made within the Grand Challenge in Verified Software initiative have been described by the researchers working on the Mondex system and a verified file store. Several large-scale experiments on system development and software verification were presented by the various researchers working in the software industry.

Discussions of such topics as foundations of program refinement, verification, theorem proving, techniques for ensuring dependability, automatic tool support for system development and verification, modeling concurrency and many others resulted in several new joint research initiatives and collaborative works.

6.5 Quantitative Software Design

Seminar No. **09432** Date **20.10.–23.10.2009**
Organizers: Astrid Kreissig, Iman Poernomo, Ralf Reussner

Quantitative software design is a field of research that is not yet firmly established. A number of challenging open research issues are only recently being addressed by the academic research community (see below). The topic is also gaining increasing emphasis in industrial research, as any progress towards a more systematic and goal-driven software design promises the reduction of costs and risks of software projects, by avoiding current trial-and-error approaches to design. The whole field is therefore of high industrial relevance, though it is far from providing ready-to-use solutions.

The research area of quantitative software design is not yet firmly established. Its subject is the investigation of the relationship of the design of a software system on quantitatively measurable quality attributes. Such quality attributes include internal quality attributes (such as maintainability), but also externally measurable attributes (such as performance metrics, reliability or availability). This also includes quality attributes where quantitative metrics are under current investigation, such as security. While there is no debate on the fact that the software design (mainly its architecture) is the main influencing factor on the quality of the resulting software system, an understanding of how an architecture influences on the quality is currently primarily anecdotal. Much progress was made on recent years in the area of model-based and model-driven quality prediction where software architectures are used as an input for the prediction of the quality of the system, namely various performance metrics, such as throughput, response time or reaction time. However, several important scientific questions remain unanswered:
trade-off decisions between antagonistic quality attributes,
quantitative metrics for relevant quality attributes such as security,
software design as an optimisation problem,
lifting classical maintainability metrics to the architectural level.

The aim of the seminar was to bring together industrial and academic experts from relevant areas to establish the field of quantitative software design. We were fortunate enough to have a group whose expertise cut across the relevant domains:software architecture, component-based software engineering, modelbased software, quality of service and business informatics. The seminar was organized into smaller discussion groups who attempted to define and problematise the relevant sub areas of the field.

6.6 Evolving Critical Systems *(Dagstuhl PerspectivesWorkshop)*

Seminar No. **09493** Date **02.12.–04.12.2009**
Organizers: José Luiz Fiadeiro, Michael G. Hinchey, Bashar Nuseibeh

The need is becoming evident for a software engineering research community that focuses on the development and maintenance of Evolving Critical Systems (ECS). This community must concentrate its efforts on the techniques, methodologies and tools needed to design, implement, and maintain critical software systems that evolve successfully (without risk of failure or loss of quality).

In recent years a number of factors have changed in the software engineering landscape to highlight the importance of Evolving Critical Systems (ECS). There are new difficulties and new attitudes that may have been specific to particular industries and software engineering sub-fields but are now widespread across the discipline. We have identified the following five "game changers":

- Universality of Software: This means that software failures are more likely to affect ordinary people.
- Pervasiveness of Software: As software embeds itself into the fabric of society failures affect more of society. This increases the criticality of even very simple software.
- Increased Interactions with People: As software is deployed to control systems in which human actors participate, the issue of people in the loop becomes more important. As it is more common for software and (non-technical) humans to interact the implications for modelling the system and for criticality have become more common.
- Increasing Complexity: Software itself is more complex and much real-world software is becoming entangled and dependent on software developed by third-party operators.
- Increased Tempo of Evolution: The tempo of software evolution is increasing as users become accustomed to demanding more from software.

We believe that the software engineering community needs to concentrate efforts on the techniques, methodologies and tools needed to design, implement, and maintain critical software systems that evolve successfully (without risk of failure or loss of quality). The Perspectives Workshop on Evolving Critical Systems held in Schloss Dagstuhl in December 2009 brought key software engineering researchers and practitioners (19 participants from 8 countries) who are in positions to influence their organisation's research direction together to discuss ECS. Similar issues and questions must be addressed within ECS as in other (non-ECS) software engineering research, but with the added (and conflicting) requirements of predictability/quality and the ability to change.

The fundamental research question underlying ECS research is: How do we design, implement, and maintain critical software systems that are highly reliable, and that retain this reliability as they evolve without incurring prohibitive costs. We discussed an incomplete list of demands that must be met before the ideals of ECS can be fully realised, including:

- Architectural models: We must determine the characteristics that make a successful architectural model and/or technique for ECS.
- Changing Development Environment: We must be able to maintain the quality of critical software in spite of constant change in its teams, processes, methods and toolkits. We must improve our existing software design methodologies so that they facilitate the support and maintenance of ECS, e.g., how can we use agile development methodologies to evolve critical software?
- Capturing Requirements: We must be able to specify what we want to achieve during an evolution cycle and to be able to confirm that we achieved what we intended, and only what we intended; in other words, we must be able to capture and elucidate the requirements for change in such a manner that allows that change to take place correctly.
- Effort Estimation: We must develop techniques for better estimating specific evolution activities a priori and only attempt software change when we are certain that evolution will be successful and that the benefit outweighs the cost. Too many software change activities run over time and budget and ultimately many are abandoned.
- Model Based Evolution: We must develop strategies to make model-driven, automatic evolution a viable alternative to manual change. In cases where it is not appropriate to mechanise change we must develop heuristics for determining when such an approach is viable. Where it is necessary for humans to perform the change we must develop support tools that make this a less risky enterprise.
- Traceability: We must develop new tools for traceability that keep the various software artefacts (e.g., documentation and source code) in sync throughout the evolution cycle. Where regulatory compliance is required, these tools must ensure that evolution results in compliant software.
- Evolving in Runtime: During runtime evolution we must ensure that run time policies must be adhered to. We must develop techniques that can monitor and model changing requirements in dynamic environments (especially autonomic and adaptive software). We must develop strategies for evolution that are tolerant of uncertainty in the operational environment, where the environment changes in a deterministic, non-deterministic, or stochastic manner. We must ensure that software never evolves into a state where it exhibits unstable behaviour.

One of the outcomes of the workshop was a special issue of IEEE Computer Magazine on Evolving Critical Systems in May 2010.

Chapter 7

Distributed Computation, Networks, Architecture

7.1 Management of the Future Internet

Seminar No. **09052** Date **27.01.–30.01.2009**
Organizers: Olivier Festor, Aiko Pras, Burkhard Stiller

The goal of this Dagstuhl Seminar on "Management of the Future Internet" was to discuss the development of the provisioning of high quality Future Internet services to everybody by means of modern Network Management methods. This was achieved in an effective manner, since the discussion and presentation of adequate management aspects capable of configuring, monitoring, and controlling the Future Internet services delivered was performed. Such a management plane has been the focus of research and development in the context of traditional data and voice networks. And it was shown that the development of tomorrow's Future Internet — providing integrated voice and data services over multiple access networks — puts new major challenges to this area in terms of scalability, dynamicity, security, and automation.

Within the "Management of the Future Internet" Dagstuhl Seminar, the functionality of existing work on management of the Internet technology, traditional management approaches, and economic management approaches, especially with respect to its capabilities to allow for an integrated approach of design and deployment of future networks that incorporate new services, have been considered.

More specifically, the following areas of interest have been partially addressed:

- Management Mechanisms for the Future Internet
- Fault, Configuration, and Security Operation in the Future Internet
- Intra and Inter-Domain Autonomic Management in the Future Internet
- Economic Network and Service Management in the Future Internet
- Commercial operator-oriented mechanisms (Traffic Management)

7.2 Delay and Disruption-Tolerant Networking (DTN) II

Seminar No. **09071** Date **08.02.–11.02.2009**
Organizers: Kevin Fall, Cecilia Mascolo, Jörg Ott, Lars Wolf

Today's Internet architecture and protocols, while perfectly suitable for wellconnected users, may easily experience serious performance degradation and entirely stop working in more challenged networking environments. Such environments are manifold, ranging from mobile users experiencing occasional or frequent disconnections to communication services for remote areas, to vehicular network communication in large areas, sensor networks to habitat or wildlife monitoring, and to space and underwater communications. These scenarios all share two commonalities: that an end-to-end path between two communicating nodes may not exist at any single point in time and that communication delay may be significant. Luckily, in most cases, delay in the delivery of the data can be tolerated. However, with the continued expansion of the Internet into new areas and the increasing penetration of communication technologies into more areas of life and technology, these environments become commonplace and are no longer restricted to exotic sensing applications but are quickly becoming relevant to consumers in everyday life.

Many attempts over recent years of incrementally fixing the Internet protocols in a bottom up fashion have only achieved partial successes: while mobile IP, HIP, transport, session, and cross-layer approaches may support changes of network attachments and short-term disconnections, a more fundamental approach is needed to address networking environments in which delays and disconnections may last for significant periods of time, and are the rule rather than the exception.

Delay-tolerant Networking (DTN) has taken a more encompassing approach to dealing with virtually all types of connectivity challenges, from bit rate to errors to delays to disruptions. By providing a novel communication abstraction that relies exclusively on asynchronous hop-by-hop message passing with no need for instant end-to-end connectivity, DTN concepts enable communications even under adverse conditions. This comes, however, at the cost of interactivity of communications, rendering any kind state synchronization or validation more difficult and raising new challenges. These include routing protocols – that need to operate under often unknown future conditions, security mechanisms – that can no longer carry out instant key derivation or validation even if a security infrastructure was in place, and application protocols and paradigms – that can no longer rely on simple lower layer abstractions promising (mostly) instant and reliable interactions.

Overall, the Dagstuhl seminar DTN II has provided the participants with a forum for fruitful discussion of present and future work on emerging networking applications and paradigms. The seminar has contributed to furthering the understanding of the perspectives of future development and real-world deployments of delay-tolerant networking as well as helped identifying issues – as research and engineering directions – to be resolved on this way.

7.3 Bandwidth on Demand

Seminar No. **09072** Date **08.02.–11.02.2009**
Organizers: Panayotis Antoniadis, David K. Hausheer, Kohei Shiomoto, Burkhard Stiller, Jean Walrand

The rapid technological progress in the area of network virtualization, mainly driven by new optical fiber technology and virtual router infrastructures, is generating a new trend for "on demand" provisioning of bandwidth or even whole networks for applications that require short-term bandwidth assignments at large scale, such as large sporting events or cultural open air activities. Network virtualization, in addition to numerous benefits that it offers in terms of security, flexibility, and reliability, enables the transparent sharing of physical network equipment between different customers of the same network provider. The current trend is backed by new optical network management systems which enable the provisioning of end-to-end light-paths across multiple independent optical network domains.

At the same time, the proliferation of wireless technology has enabled users "to be connected" anytime and anywhere in the world. What's more, wireless devices allow users to offer network connectivity to each other, *e.g.* via mobile ad-hoc networks or neighborhood wireless mesh networks. The support of bandwidth allocation in a fully decentralized manner, such as based on emerging peer-to-peer (P2P) concepts, shows further advantages in terms of robustness and scalability for large-scale systems.

Despite (or due to) these recent technical advances, the provisioning of the right amount of bandwidth at the right location and at the right time remains a challenge. Suitable business models for "on demand" bandwidth services have not yet evolved. Moreover, the design of resource allocation policies and incentive mechanisms for cooperation in this context are very challenging and interesting research questions. Resource allocation mechanisms should aim, ideally, to maximize the overall social welfare of the system. However, participants may not have the incentive to disclose truthfully their private information. Auctions are a standard way to achieve such objectives, but the distributed environment and the different types of resources involved poses significant challenges on their design and implementation.

Wireless technology enables interesting new business models such as FON and Boingo. However, suitable incentive mechanisms are required for such systems to operate efficiently and avoid free riding and other types of undesirable behavior in terms of resource sharing and the overall distributed management of the system, which depends on the cooperation and resource contributions of all participants. These incentive issues reduce the overall value that could be generated thanks to the positive externalities that appear in P2P systems. It is of interest to study the potentials and limits of P2P bandwidth sharing systems and understand to what extent they could harm the ISP business. In addition, legislative and regulative issues related to these concepts have to be tackled.

Therefore, the purpose of this Dagstuhl Seminar on "Bandwidth on Demand" was to bring together researchers and practitioners from different disciplines to discuss and develop partially technical, economic, and regulatory mechanisms for the provisioning of bandwidth

on demand services. The key topics tackled by this seminar included but were not limited to:

- The technical design of scalable, robust, and cost-effective bandwidth allocation and provisioning schemes, including fully decentralized and market-based mechanisms such as auctions
- Economic studies and modeling of market and business models in carrier and service provider networks, including cost and revenue models as well as game theoretical bandwidth on demand models
- Resource allocation and provision in non-profit systems such as neighborhood wireless mesh networks and network testbed infrastructures like GENI or PlanetLab
- Industrial developments of new technologies that facilitate or create impediments to bandwidth on demand, including network virtualization technologies and wireless mesh networks
- Legislative and regulatory issues related to the Telecom Act and in comparison to other commodities markets such as the electric grid, as well as legal issues of P2P trading infrastructures

Discussion and Conclusions

A closing discussion was held at the end of the seminar to draw conclusions. The discussion was opened by Fernando Beltran. He thought that Bandwidth on Demand becomes more feasible in wireless than in wired areas, because of the uncertainties in mobility, since users are generators of unpredictable service demand. But he argued that an agreement on standards would be needed. In the future he envisions that agents may be working for us, *e.g.*, driving around looking for the best service.

Peter Reichl enjoyed that the topic brings together different aspects. But he asked what "on demand" would really mean. He argued that BoD will not differ so much between fixed and wireless networks, since the key difference is between user-to-ISP and ISP-to-ISP relationships. Later on, Martin Waldburger opposed that the most important question is whether we address consumers or businesses. The Kindle example shows that making it work is the way to go.

Aiko Pras found that BoD is similar to a water-pipe, but end customers are not interested in capacity, they are interested in data. With optical networks there is the possibility to provide bandwidth on demand, but we should look mainly at big customers. Torsten Braun agreed that users request for service on demand. Therefore, it is more important to focus on services than on bandwidth, and network providers must care how these services can be provided. However, the future importance of virtual networks is unclear.

Athanassios Androutsos believes that bandwidth is an important input to the QoS provisioning process. However, appropriate pricing models are needed to allocate bandwidth in an efficient manner. For example, long-term bandwidth provisioning could reach a

cheaper price. Architectural frameworks should be considered to enforce a certain pricing model, but enforcing mechanisms are needed in order to support Self-* and dynamic SLA provisioning. Shigeo Urushidani added that we need more user experience. Without it could be difficult to improve network provisioning. The critical mass of users that use BoD service has to be increased and they are only willing to pay, if they get addicted. However, Burkhard Stiller was concerned that a viable business model may not be achieved. It could be applicable to services, resources, and bandwidth, but protocols may be different. The key is to standardize the BoD interface.

Panayotis Antoniadis had the impression that the BoD discussion is about the future of the Internet, i.e. we try to predict the scenario of the future. But it is unclear to what extent there will be scarcity, therefore, we don't know the type of problem. Bruno Tuffin added that it is clear now that it is not clear what BoD is. But he observed that much more tools are available for the user-to-ISP than for ISP-to-ISP relationship. Isabelle Hamchaoui initially thought that academic people were happy about BoD. She emphasised that BoD is not a main issue for the customer, but it is important for the ISP-to-ISP relation and, thus, there is a need for new solutions in the inter-domain context. She was surprised to see discussion about signalling protocols like GMPLS and new things like energy networks, which shows that academic people are open to operator problems. George Huitema replied that energy on demand is different compared to bandwidth on demand for telcos. In the energy sector it is more interesting to look at the user to energy provider interaction.

Adrian Farrel thought initially that BoD is applied down in the network. But he is now convinced that we should have Qo*, and part of this is Quality of Business. The relation between Qo* and *oD is important. The users know best what they mean by Qo*, but they don't understand *oD. Hopefully, this can be parameterized. Giancarlo Ruffo believes that the problems can be solved with a layered architecture, rather than with peering. But contracts between ISPs and national governments are necessary. Consumers are at the end of the value chain and cannot be cut off, since they are part of the long-tail that contribute to the service. In that respect, he still sees many problems at different levels.

In summary, the seminar was very successful. With 25 attendees it did lead to many fruitful discussions and scientific exchange. A future BoD workshop is planned to be organized again in colocation with a large conference.

7.4 Naming and Addressing in a Future Internet *(Dagstuhl Perspectives Workshop)*

Seminar No. **09102** Date **01.03.–04.03.2009**
Organizers: Jari Arkko, Marcelo Bagnulo Braun, Scott Brim, Lars Eggert, Christian Vogt, Lixia Zhang

The purpose of the perspectives workshop on naming and addressing in a future Internet is to generate input to the research and engineering community on how to evolve the Internet architecture to satisfy the naming and addressing requirements of the existing and future

Internet. The workshop will bring together key researchers and engineers from the realm of Internet naming and addressing to find a common understanding on a preferred evolution of the Internet architecture. The outcome of the perspectives workshop will be presented to the Routing research group of the Internet Research Task Force, as an input into the group's current effort in developing a scalable routing architecture.

The IP addresses that are used to deliver data in today's Internet encompass three functions:

- Name. IP addresses are used by protocols above IP as node identifiers.
- Locator. IP addresses name the topological locations of nodes.
- Forwarding directive. IP addresses need to be aggregatable based on network-topological locations to make a routing system scale. Thus, IP addresses for network-topologically close locations can be aggregated into a common forwarding directive.

Overloading the functions of name, locator, and forwarding directive – the latter two of which are commonly subsumed as "addressing" – into IP addresses suits an Internet when it is small, and where neither network topology nor host attachments change often. It was hence a wise design choice at the time it was devised because it avoided the (back then unnecessary) complexity that a secure split between these functions would have entailed. In today's Internet, however, there is increasing pressure to decouple the three functions of IP addresses:

- Hosts are oftentimes present at multiple locations in the network, be it sequentially due to mobility, or simultaneously for better performance or fault tolerance. IP addresses that serve as locators then can no longer serve as a stable and unique name.
- Networks at the Internet edge are also increasingly present at multiple places in Internet topology, be it sequentially due to provider changes, or simultaneously because they access the Internet via multiple providers for better performance or fault tolerance. "Network-topological closeness" of two IP addresses, which used to be the basis for efficient IP address aggregation, is then no longer clearly defined. The consequence is an abandoning of the function of IP addresses as forwarding directives, hence less scalable data forwarding.

A future Internet architecture must hence decouple the functions of IP addresses as names, locators, and forwarding directives in order to facilitate the growth and new network-topological dynamisms of the Internet. Although there have been various research efforts in the past that addresses these issues, (see proposals such as FARA, DONA, Plutarch, Triad, I3, SNF, TurfNet, IPNL, or HIP), they have made little impact on practice, perhaps with the Host Identity Protocol the only exception.

The purpose of this workshop on naming and addressing in a future Internet is to bring together researchers and engineers to develop a crystal problem statement regarding naming

and addressing issues in the existing Internet architecture, to examine solution directions to meet the needs of the future Internet, and to identify immediate next steps towards an evolutionary path to address the architectural issues. The outcome of this workshop will be presented to the Routing research group of the Internet Research Task Force as one input to its effort in developing a scalable routing architecture, and, if appropriate, to the Internet Engineering Task Force as an informational document to be referenced in future protocol development.

7.5 Architecture and Design of the Future Internet *(Dagstuhl Perspectives Workshop)*

Seminar No. **09162** Date **14.04.–17.04.2009**
Organizers: Georg Carle, David Hutchison, Bernhard Plattner, James P. G. Sterbenz

This workshop brought together thirty seven experts from Europe, North America and Asia to discuss the way ahead for the Internet.

The aims of the workshop were as follows:

- to understand the 'state of the art' in Future Internet research by reviewing current programmes in the USA, Europe and Asia;
- to identify gaps and new opportunities in Future Internet research so that we can recommend new programme or project topics to appropriate bodies;
- to discuss potential collaborative activities amongst programmes or projects in order to make the most of current research: these activities could include testbeds and workshops.

During the workshop, the participants made contributions in:

- Defining and discussing the problem space. It was broadly agreed that three aspects are crucial: technological, economic, and societal/political;
- Describing relevant Future Internet activities in the USA, Europe and Asia;
- Remaining challenges in network essentials: naming/addressing, routing, mobility;
- How to provide for key network properties: security, resilience, performance;
- Management, policy, economic, green, and other Future Internet issues;
- Architecture questions: evolution vs revolution, virtualization, testbeds.

All of the above took place in plenary sessions, with a view to identifying the key issues that would be debated in smaller groups on the last day of the workshop. These key issues were as follows:

1. Sacred cows;
2. Management issues;
3. Social, economic, green etc.;
4. Programmability and virtualization;
5. Personalization and context for Future Internet.

The participants self-organised into groups, which produced a summary of their discussion. Each group reported back in the final plenary session and a closing discussion followed.

Group (1) covered the IP address architecture, routing structure, TCP and the end-to-end argument, dumb core and smart edges – as the 'sacred cows' of the current Internet, and debated three things: layering principles, and whether there's a need for management & control; re-routing as the primary approach to failure recovery, and whether overlays solve all problems; and virtual circuits (CO / CL).

Several of the comments indicated that we seem to be re-visiting these topics yet again, but perhaps in the light of new application or user needs (such as resilience) this is actually appropriate. Also, we don't yet know best how to make the right choices from the above sets.

Group (2) offered some basic observations and issues that still require study: nested control loops, and stability provision; humans in the (control) loop – or not; knowledge, and the amount of data required, to produce satisfactory management – how to do inferences (we still don't know how).

Group (3) covered a range of topics including the digital divide – which exists in all countries, the balance between security and privacy, network neutrality as a growing concern (or not), how to be 'greener' in networking, and the cultural and objective differences between academic and industry (for example ISPs) – where there will always be some tension.

Group (4) was concerned with whether programmability is now having its time, for example with the advent of multiple cores and the prospect of virtualization; associated research imperatives include router architecture, protocol architectures for massive parallelism, and the architecture of networks where the main routers have multiple cores.

Group (5) asked questions about personalization and context in the Future Internet: what is the typical usage; will virtual and physical worlds become more integrated, raising possible issues of privacy, social exchange etc.; which application areas will need further support, e.g. emergency, mobile video streaming, using social relationships to support communities, using public transport to support mobile users, etc.. Context was defined variously as what's 'around', situation awareness, and it was agreed that context-modelling is an upcoming issue.

Conclusion

Unsurprisingly, in such a short time, the workshop participants did not manage to point definitively towards a new architecture or design for the Internet of the future. Rather, what the workshop did was to identify the beginnings of a number of promising tracks of investigation, listed above as (1-5), which are distinct from the activities currently being undertaken in the USA, Europe and Asia. These issues could form the basis of specifically-targeted seminars at Dagstuhl, and elsewhere, that have a much more interdisciplinary flavour and participant balance than the one reported here. It is clear that many diverse influences are being brought to bear on the possible Future Internet 'shape', and we should put together people from these diverse backgrounds in a focussed programme of discussions designed to elicit some more concrete outcomes. This would inevitably mean that the proportion of participating computer scientists and electrical engineers would have to be reduced considerably, balanced by a larger, carefully selected set of participants from other disciplines.

7.6 Fault Tolerance in High-Performance Computing and Grids

Seminar No. **09191** Date **03.05.–08.05.2009**
Organizers: Franck Cappello, Laxmikant Kale, Frank Mueller, Keshav Pingali, Alexander Reinefeld

The objective of this seminar was to bring together researchers and practitioners from the HPC and Grid communities to discuss medium to long-term approaches to address fault tolerance (FT). The focus of solutions was on the practical, system side and with the intent to reach beyond established solutions.

Overall, the objective of the workshop is to spark research activities in a coordinated manner that can significantly enhance FT capabilities of today's and tomorrow's HPC systems and Grids. The benefits of this work extend to the community of scientific computing at large, well beyond computer science. Due to the wide range of participants (researchers and industry practitioners from the U.S., Europe, and Asia), forthcoming research work may significantly help enhance FT properties of large-scale systems, and technology transfer is likely to eventually reach general-purpose computing given the increasing trend to multi-core parallelism and server-style computing, such as Google. Specifically, the work should set the seeds for increased collaborations between institutes in Europe and the U.S./Asia. If successful, a follow-up seminar may be organized in the following year.

This meeting was the first of its kind at Dagstuhl and provided a foundation to create a community platform with a cohesive outlook on FT in HPC and Grids. The presentations of participants concentrated on fundamental issues related to FT in HPC applications, runtime systems, operating systems, networking, I/O and scheduler. The program consisted of an introductory session for all participants, 22 presentations well as four "open mic" sessions where time was set aside for spontaneous discussions, brain storming and

community-building plans. The seminar brought together a total of 31 researchers and developers working in the areas related to fault tolerance from universities, national research laboratories and computer vendors. The goals were to increase the exchange of ideas, knowledge transfer, foster a multidisciplinary approach to attacking this very important research problem with direct impact on the way in which we design and utilize parallel systems to make applications resilient to faults in hardware or software.

7.7 From Quality of Service to Quality of Experience

Seminar No. **09192** Date **05.05.–08.05.2009**
Organizers: Markus Fiedler, Kalevi Kilkki, Peter Reichl

For at least a decade, Quality of Service (QoS) has been one of the dominating research topics in the area of communication networks. Whereas the Internet originally has been conceived as a best-effort network, the introduction of QoS architectures like Integrated Services or Differentiated Services was supposed to pave the way for high-quality real-time services like Voice-over-IP or video streaming and thus to increase the competitiveness of packet-based TCP/IP networks.

Originally, the notion of end-to-end QoS was, according e.g. to ITU-T, aiming at the "degree of satisfaction of a user of the service". In the course of time, however, the dominating research perspective on QoS has become more and more a technical one, focussing on monitoring and improving network performance parameters like packet loss rate, delay or jitter. But end users usually are not bothered at all about technical performance; what they really care about is the experience they are able to obtain, and the Internet provided, even without any QoS mechanisms, a lot of new experiences, like web-browsing, e-mail and search engines.

Based on this insight, we have recently observed an important paradigm shift as far as service quality is concerned. While the prior "grand challenges" of QoS research have begun to disappear from the research agenda, e.g. due to large-scale overprovisioning in today's core networks, a counter movement has started to become visible, with the aim of interpreting "end-to-end quality" in the proper sense of regarding the human being as the end of the communication chain. As a result, the notion of Quality of Experience (QoE, abbreviated also as QoX) has appeared, describing quality as perceived by the human user instead of as captured by (purely technical) network parameters.

Currently, there are several attempts to define QoE, but the ultimate definition is still lacking. According to [2], Quality of Experience may be defined as "overall acceptability of an application or service as perceived subjectively by the end-use". Hence, Quality of Experience is a subjective measure from the user's perspective of the overall value of the service provided, and thus does not replace, but augment end-to-end QoS by providing the quantitative link to user perception. As such, it extends the current QoS perspective described above towards the actual end user, including technical QoS as well as the expectations of the end users, the content of the service, the importance of service for the end user, the characteristics of the device, the usability of the human-computer interfaces,

the joyfulness of interaction, the perception of security, and maybe even the price of the service, to name but a few new ingredients.

Today, research on Quality of Experience faces the challenge of creating a unifying interdisciplinary framework that is able to combine these diverse aspects under a common umbrella in a way that we are able to predict the behaviour of end users when new services are offered to them and to ensure service provisioning and management that actually meets user expectations. Therefore, understanding the transition from Quality of Service to Quality of Experience will become an indispensable prerequisite for taking the subjective user experience into proper account while designing and providing successful future communication services.

The Dagstuhl Seminar 09192 was an important "kick-off" to reconsider the concept of QoE, leaving more questions open than there were before the seminar, and for the formation of a community which already has taken first steps to drive the questions further. As Dagstuhl offers perfect surroundings for creative and open discussions, both community and organisers would be very much interested in a follow-up Dagstuhl Seminar.

7.8 Visualization and Monitoring of Network Traffic

Seminar No. **09211** Date **17.05.–20.05.2009**
Organizers: Daniel A. Keim, Aiko Pras, Jürgen Schönwälder, Pak Chung Wong

The seamless operation of the Internet requires being able to monitor and visualize the actual behaviour of the network. Today, IP network operators usually collect network flow statistics from critical points of their network infrastructure. Flows aggregate packets that share common properties. Flow records are stored and analyzed to extract accounting information and increasingly to identify and isolate network problems or security incidents. While network problems or attacks significantly changing traffic patterns are relatively easy to identify, it tends to be much more challenging to identify creeping changes or attacks and faults that manifest themselves only by very careful analysis of initially seemingly unrelated traffic pattern and their changes. There are currently no deployable good solutions and research in this area is just starting. In addition, the large volume of flow data on high capacity networks and exchange points requires to move to probabilistic sampling techniques, which require new analysis techniques to calculate and also visualize the uncertainty attached to data sets.

Goals

The aim of the seminar is to bring together for the first time people from the networking community and the visualization community in order to explore common grounds in capturing and visualizing network behaviour and to exchange upcoming requirements and novel techniques. The seminar also targets network operators running large IP networks as well as companies building software products for network monitoring and visualization.

We believe that bringing experts from two usually separate fields together makes this seminar unique and we expect that the intensive exchange in a Dagstuhl seminar setting has high potential to lead to joint follow-up research.

Research Questions

The following research questions were suggested for discussion:

- What are suitable data analysis and visualization techniques that can operate in real-time and support interactive online operation?
- How can monitoring and visualization techniques be made scalable?
- How can distributed monitoring systems be self-organizing and adapt dynamically to changes in network and service usage?
- How can algorithms aggregate data within the network and trade accuracy of the measurement results against data collection overhead?
- What are suitable sampling techniques and how does sampled data impact data analysis techniques and data visualization?
- Which filtering, zooming, and correlation techniques can be applied in real-time?
- What are good techniques for visualizing unusual traffic patterns or very rare patterns?
- What are effective methods to detect and visualize intrusions, like (distributed) scan attempts and denial of service attacks.

While this item list was helpful as an orientation, not all of the items were actually covered during the seminar. Moreover, other concerns, such as NetFlow storage and retrieval, were emphasized in the presentations and discussions.

Conclusions

The Visualization and Monitoring of Network Traffic seminar was a fertile meeting in which researchers from diverse background met. It included industry and academia, senior and junior researchers, multinational representation, and people coming from several disciplines. This diversity resulted in interesting and useful discussions, new understandings of the fundamental concepts and problems in the field, and in new collaborations on an array of problems which were not well defined or identified prior to this seminar.

Several work groups during the seminar not only generated new insights into specific topics in the field of visual network monitoring, but also initiated ongoing joint work, with group members continuing the work they started at the seminar. The seminar included multiple presentations and discussions. In particular, the largely disjoint research communities

Networking and Visualization exchanged their methods and unsolved problems resulting in fruitful discussions and awareness of the respectively other field.

This seminar clearly illustrated the diversity, relevance, and fertility of the topics we presented and discussed. The intensity of the participants' involvement leads us to believe that the interactions fostered by the seminar will generate a lot of follow-up research, and eventually lead to practical use as well.

7.9 Algorithmic Methods for Distributed Cooperative Systems

Seminar No. **09371** Date **06.09.–11.09.2009**
Organizers: Sándor Fekete, Stefan Fischer, Martin Riedmiller, Suri Subhash

A standard scientific method for *understanding* complicated situations is to analyze them in a top-down, hierarchical manner. This approach also works well for *organizing* a large variety of structures; that is why a similar hierarchical, centralized approach has worked extremely well for employing computers in so many aspects of our life: Data is gathered, processed, and the result is administered by one central authority.

On the other hand, the structures in our modern world are getting increasingly complex. The resulting challenges have become so demanding that it is impossible to ignore that a large variety of systems are based on a very different principle: The stability and effectiveness of our modern political, social and economic structures relies on the fact that they are based on decentralized, distributed and self-organizing mechanisms. This paradigm shift is also reflected in a variety of modern computing systems, which work in a distributed manner, based on local (and thus: incomplete) information and interaction, and implement the results in a localized fashion; as opposed to a variety of social or economic situations, we may assume that the individual components of such a system are not primarily selfish, but pursue a joint goal that is to be reached in collaboration.

The purpose of this workshop was to bring together researchers from different disciplines whose central interest is in both algorithmic foundations and application scenarios of distributed cooperative systems. In particular, participants from the following communities were present:

AF **Algorithmic Foundations.** When developing a systematic method for solving an algorithmic problem by a cooperating set of loosely coupled processors, the result is a distributed algorithm. One of the resulting consequences is incomplete information, for which an increasing number of algorithmic aspects have been studied. Moreover, a number of additional issues are considered, such as communication complexity, timing issues, and the amount and type of information that is available to individual processors.

SN **Sensor networks.** In recent time, the study of wireless sensor networks (WSN) has become a rapidly developing research area, both from the theoretical and the

practical side. Typical scenarios involve a large swarm of small and inexpensive processor nodes, each with limited computing and communication resources, that are distributed in some geometric region; communication is performed by wireless radio with limited range. From an algorithmic point of view, these characteristics imply absence of a central control unit, limited capabilities of nodes, and limited communication between nodes. This requires developing new algorithmic ideas that combine methods of distributed computing and network protocols with traditional centralized network algorithms. In other words: How can we use a limited amount of strictly local information in order to achieve distributed knowledge of global network properties? Just now, an important set of additional challenges for sensor networks is starting to emerge from mobile nodes, making it necessary to deal with additional problems arising from network dynamics.

RT **Multi-robot systems.** Multi-robot systems consist of several individual robots, either identical or heterogeneous. Among the scenarios for teams or swarms of autonomous robots are robot soccer (RoboCup), rescue missions, exploration and other complex tasks that can be carried out in a distributed fashion. Beyond the technical aspects of perception, behaviors, learning, and action, the most interesting issue in the context of this interdisciplinary seminar are various modeling aspects that are getting quite close to those faced in areas such as SN: after all, a sensor-equipped robot becomes quite similar to a mobile sensor node, and both face similar difficulties, but also possibilities.

AP **Application scenarios.** In order to provide further scenarios for challenges and discussions, we included a selection of researchers from other application areas; among these was

- Traffic: making use of car-to-car communication, it has become possible to provide online, up-to-date local information and coordination. What challenges can be tackled by making use of these possibilities?
- Biology: swarm behavior of animals has developed over millions of years. What lessons can be learned from such behavior?

These four aspects were subdivided into algorithmic foundations (provided by AF), two specific areas (SN and RT) that form the link between pure theory and real-life applications, and a variety of real-life challenges (provided by AP) that can serve as goals and benchmarks for the other scientific work.

Quite naturally, there was some amount of overlap between these four areas in terms of individual researchers, as various scientists combine theory and practice, to a varying degree. Despite of the previous distinction between the different fields, a variety of aspects implied that similar problems were faced, so that a dialogue between the researchers turned out to be quite fruitful.

Each of the fields both benefitted and contributed:

AF Theoretical methods of distributed algorithms and algorithms with incomplete information form the basis for the algorithmic work on the application scenarios (AP), as

well as the problems arising in both SN and RT. Fundamental insights and results turn out to be useful for the development of practical methods, but also show basic obstacles for obtainable results. Conversely, application scenarios help to focus the theoretical algorithmic work, and lead to the identification of new kinds of problems.

SN The practical side of sensor networks gives rise to a number of quite specific scenarios for distributed algorithms. Many of these problems consider stationary nodes, for which a variety of aspects enjoy an increasing amount of understanding; however, there is growing demand for dealing with mobile sensor nodes (in particular when dealing with scenarios from the application areas AP), requiring extensions of theoretical work (AF), but also leading to a growing similarity with scenarios faced by RT.

RT The ever-improving technology and control for autonomous robot systems has become more and more sophisticated, and advanced to the point where there is an increasing demand for higher-level, algorithmic methods (AF). The aspects of dynamics, communication and outside information give rise to a number of quite challenging scenarios (AP); moreover, problems like the exploration of unknown territory by a swarm of robots still require a lot of algorithmic work (AF). This is where the similarity to systems of mobile nodes in a sensor networks (SN) are striking; it is obvious that the exchange between both communities was quite beneficial, as some of the fundamental challenges are surprisingly similar.

AP The application areas described above provide both a collection of grand challenges and a reality check for the theoretical methods by AF and the specific methods developed by SN and RT; on the other hand, solutions and insights by AF, SN, and RT gave rise to completely new possibilities for mastering those challenges.

The workshop brought together 35 researchers from nine countries. The 20 presentations, varying in length, covered a large variety of topics. Owing to the combination of different research areas, there were a number of survey talks and discussion session, but also a variety of individual research presentations. In addition, there was sufficient time for informal discussion and small-scale interaction.

Overall, participants agreed that the interdisciplinary nature of the workshop was quite fruitful. There was strong interest in repeating this event with a similar combination of fields and researchers in the not-too-distant future.

Chapter 8

Scientific Computing

8.1 Combinatorial Scientific Computing

Seminar No. **09061** Date **01.02.–06.02.2009**
Organizers: Uwe Naumann, Olaf Schenk, Horst D. Simon, Sivan Toledo

The activities of the seminar focused on combinatorial issues in high-performance scientific computing. The activities included:

- eight one-hour invited talks
  - Bruce Hendrickson,[1] Sandia National Laboratory: Combinatorial Scientific Computing: A View to the Future
  - Alex Pothen,[2] Purdue University: Graph Matchings in CSC
  - Rob Bisseling,[3] Utrecht University: Combinatorial Problems in HPC
  - Paul Hovland,[4] Argonne National Laboratory: Combinatorial Problems in Automatic Differentiation
  - Iain Duff,[5] Rutherford Appleton Laboratory and CERFACS: Combinatorial Problems in Numerical Linear Algebra
  - Ruud van der Pas,[6] SUN Microsystems: Present and Future of High-Performance Scientific Computing
  - David Bader,[7] Georgia Institute of Technology: Emerging Applications in Combinatorial Scientific Computing

[1] URL: `http://www.sandia.gov/~bahendr/`
[2] URL: `http://www.cs.purdue.edu/people/faculty/apothen/`
[3] `http://www.math.uu.nl/people/bisselin/`
[4] `http://www.mcs.anl.gov/~hovland`
[5] `http://www.numerical.rl.ac.uk/people/isd/isd.html`
[6] `http://blogs.sun.com/ruud/`
[7] `http://www.cc.gatech.edu/~bader/`

  - Michael Mahoney,[8] Stanford University: Combinatorial and scientific computing approaches to modern large-scale data analysis

- thirteen 20-minute contributed talks by participants from seven countries
- six software tutorials
  - Eric Boman, Sandia National Laboratory: Zoltan[9]
  - Francois Pellegrini, LABRI: Scotch and PT-Scotch[10]
  - Jean Utke, Argonne National Laboratory: OpenAD[11]
  - Andreas Wächter, IBM Research: Ipopt[12]
  - Andrea Walther, Technical University Dresden: ADOL-C[13]
  - John Gilbert, University of California Santa Barbara: Star-P[14]
- three round table discussions
  - Graph coloring for parallel computation (organizer: Assefaw Gebremedhin, Purdue University)
  - Multilevel algorithms for discrete problems (organizer: Eric Boman, Sandia National Laboratory)
  - Data-Flow Reversal in Adjoint Codes (organizer: Uwe Naumann, RWTH Aachen University)

Participants also enjoyed one of two recreational activities in a free afternoon (a long hike or a trip to Trier). As usual in Dagstuhl seminars, participants also engaged in lively informal professional (and personal) discussions during breaks, at meal times, and late into the night in the cafeteria / the wine cellar.

Invited Talks

Some of the invited talks surveyed different problem areas within combinatorial computing (all of these also described very recent research). These talks included Henrickson's talk, which surveyed the entire field and attempted to define it in both old and new ways. The talks by Pothen, Bisseling, Hovland, and Duff each focused on one problem area within the field. Pothen's talk focused mostly on linear-time, highly parallel approximation algorithms for weighted graph matching and for graph coloring. Bisseling's talk focused mostly on parallel graph partitioning. Hovland surveyed the area of automatic differentiation, with a focus, of course, on the combinatorial problems that arise in such

[8] http://cs.stanford.edu/people/mmahoney/
[9] http://www.cs.sandia.gov/Zoltan/
[10] http://www.labri.fr/perso/pelegrin/scotch/
[11] http://www-unix.mcs.anl.gov/OpenAD/
[12] https://projects.coin-or.org/Ipopt
[13] http://www.math.tu-dresden.de/~adol-c/
[14] http://www.interactivesupercomputing.com/

computations. Duff talk focused on algorithms for sparse-matrix factorizations; it covered both traditional topics like the multifrontal method and elimination data structures, as well as on recent development, like the use of graph matching to enhance efficiency and parallelism in sparse factorizations. Bader focused on the use of high-performance accelerator processors (the Cell BE and GPUs) to speed up the solution of large-scale combinatorial problems in science.

The relatively large number of invited survey talks served as a community-building tool. The combinatorial scientific computing has been holding specialized meetings for less than 5 years, and it is still important for us to teach each other about the specific problems that each one of us address and about and techniques that we use to solve them. We have recognized that we share significant commonality, but we still make an effort to more precisely define the community and to enhance the scientific connections between its members. Our situation is quite different from that of computation geometry, say, a community that has been conducting Dagstuhl seminars for 18 years (and other meetings even earlier).

In fact, the common threads of combinatorial scientific computing revealed themselves in several ways during the seminar. Hendrickson, in the first talk of the seminar, suggested that the field is defined not only by the focus on combinatorial problems and algorithms in scientific computations, but also by a shared aesthetic; a common sense of what makes a problem important and beautiful, and what makes a proposed solution as a success. For example, researchers in combinatorial scientific computing favor problems that have a tangible impact on society (better science, better medicine, not just cleverer math); they therefore tend to assess results in a relatively practical way; they tend to assess algorithms and implementations as a whole, not to focus on one or the other; and so on. Another common thread was the focus on high-performance computing. It was widely recognized by participants that one cannot usually make a significant progress in scientific computing without paying attention to parallelism, because the only computers that can solve large-scale problems are highly parallel. This thread quickly led to another common concern, regarding the ability to implement our algorithms in a way that achieves high performance but without forcing us to spend significant amount of time tuning the implementation to each parallel machine (a fairly common behavior in high-performance computing).

Two long talks were given by people from outside the community. The talk by van der Pas focused on programming tools for high-performance computing. As mentioned above, this is an area that many members of the community feel passionately about, because they program high-performance machines and they struggle with the tools (compilers, profilers, programming languages). The talk was, therefore, very well received.

The other talk from outside the community was by Mike Mahoney, a theoretical computer scientist who talked about new techniques in large-scale data analysis. The examples in the talk drew both from non-science applications (discovering online communities from records of personal interactions on the internet) and from scientific computations (discovering significant SNPs in gene databases). The talk was exciting in that some of the problems that Mahoney talked about were clearly related to scientific computations, they were clearly combinatorial, yet the techniques that he used were very different from those used by our community. This is to a large extent due to the fact that in biological data analysis

the problems are not always well defined mathematically, as opposed to the physical sciences where problems are usually well defined. But the talk inspired some participants to search for applications of the techniques that Mahoney described.

Both Mahoney and van der Pas attended the whole seminar; they participated in many technical discussions with other participants and enriched our community and the seminar.

Contributed Talks

Virtually all the contributed talks described very recent reseach. Three were given by PhD students (Donfack, Daitch, and Langguth). They covered results in many areas of combinatorial scientific computing: automatic differentiation (Gebremedhin, Lyons, Steihaug), sparse factorizations (Davis and Li), ordering for sparse factorizations (Reid, Donfack, Scott), combinatorial preconditioning (Daitch), and huge-scale parallel PDE solvers (Arbenz, Ruede). Two interesting contributions came from long-time members of the community who are now working in new problem areas: genome sequencing (Catalyurek) and graph visualization (Hu).

Tutorials

The seminar included several hands-on tutorials on software packages that researchers in the community have been developing. These tutorials were intended to give other researchers first-hand experience in using the software, and in allowing them to continue using them on their own more easily. In other words, they were part demos to show the tools and part tutorials to make it easy to learn the tools.

The tutorials were conducted using a large cluster of laptops that were brought for this purpose from Aachen University by organizer Uwe Naumann and his colleagues from Aachen. All the relevant software packages have been installed on the laptops prior to the tutorials.

The demo part of the tutorials went fantastically well. It was a joy to see the developers of the software demonstrate it using simple examples in an interactive way. In many cases, the audience asked the tutorial presenter to try other things than he or she prepared, and the interactive nature of the demonstration was very lively. The fact that each participant had access to a laptop running the same software also helped, as participants were able to run simple examples and to examine the structure of the software and to look at a few source files.

The goal of actually teaching participants to use the software was perhaps a bit ambitious for 2-hour tutorials, but they certainly gave participants an opportunity to get started more easily than in their own offices back home, without the benefit of having the lead developer right there to answer questions.

Feedback from participants has been very positive both in terms of the anonymous survey and in terms of what participants told the organizers. The results of the anonymous survey are largely consistent with those of other Dagstuhl seminars (as summarized in

the last-60-days statistics), with many questions on which this seminar scored higher than average and a few on which it scored lower than average.

Many of the suggestions and constructive comments that participants filled in the survey reflect the wide range of expectations and wishes of participants. For example, some wished for fewer talks and more unstructured time for discussions; but many participants expressed a wish to give a talk. The organizing committee tried to balance these expectations. Similarly with respect to the length of talks: short talks are harder to follow, but leave more time for other activities.

Post-Seminar Book

Two of the organizers, Uwe Naumann and Olaf Schenk, are working on a proposal to CRC Press for publication of a special collection of articles on Combinatorial Scientific Computing in their Chapman & Hall/CRC Computational Science series. The purpose of the book is to provide the first collection of references for a diverse community of researchers working in different aspects in the exiting field of CSC. The content is strongly motivated by this seminar including survey articles as well as tutorial-style software guides. Potential readers include graduate students, young researchers, scientists in mid-career, and senior investigators from both academia and industry. Some are experts on graph combinatorial aspects, some are focusing on theoretical analysis, and some are more directed towards software development and concrete applications. Outreach into areas of science and engineering that face similar combinatorial problems as the CSC community is a major objective.

8.2 The Future of Grid Computing *(Dagstuhl Perspectives Workshop)*

Seminar No. **09082** Date **15.02.–20.02.2009**
Organizers: Dieter Kranzlmüller, Andreas Reuter, Uwe Schwiegelshohn

In February 2009, the participants of a Dagstuhl Perspectives Workshop addressed the future of grid computing. The detailed results are published in the journal Future Generation Computing Systems. In general, it can be observed that grid computing has been promoted for more than 10 years as the global computing infrastructure of the future. Some scientists like Jeremy Rifkin considered it as one of sources of the impact of scientific and technological changes on the economy and society. This claim is based on the observation that the usage of large data volumes becomes important to many disciplines, from natural sciences via engineering even to the humanities. To amortize the substantial costs of generating and maintaining these data volumes, they are typically shared by many scientists of different institutions leading to so called virtual research environments (VRE). We consider the support of these VREs to be the key property of computing grids. Further, the exploitation of these large data volumes usually leads to large simulation tasks

that require large IT systems. However, this affects also some disciplines whose members have traditionally little experience in administrating and managing these systems. These users can neither afford to manage suitable IT systems by their own nor to establish a sufficiently large local compute. Therefore, the concept of a computing infrastructure similar to the electrical power infrastructure is particularly appealing to them. However, despite significant investments in the grid concept, the number of users is not increasing. Possibly also for this reason, grid computing recently receives less attention although the basic observations still hold. Instead, new concepts, like cloud computing, seem to replace the grid computing approach.

Unfortunately, the simple electrical power grid analogy does not only provide a simple motivation for efforts to install computational grids but it also has raised hopes for a fast realization. But here this analogy really falls short as it ignores significant differences between electrical and computation power, like, for instance, the heterogeneity of resources and complexity of services.

Other experts claim grid computing to be the next evolutionary step in internet development thereby implicitly suggesting an analogy between internet/web development and grid evolution. Certainly, IT networks are a precondition for any remote computing paradigm including grid computing. Moreover, some web services already show traits of grid computing. The relationship was discussed in detail during the workshop. We believe that the internet significantly benefits from its huge user communities while there is no indication of a rapid adoption of grids to the mass market at the moment. Instead, grid users have more complex requirements, for instance, in the areas of security and reliability, leading to a slow and more evolutionary proliferation.

Further, the original reasons for grid computing still hold or even gain more importance: the number of applications exploiting large scale data resources will continue to increase as, for instance, the trend towards a virtual representation of the real world is still unbroken. Further, the smart combination of online data from sensor networks and arbitrary archives on the one hand and computing facilities on the other hand will provide novel services that do not only benefit scientific fields, like particle physics or climate research, but also reach into industrial and societal domains.

We also realized that the success of the internet is based to a large extend on the definition of a simple common protocol that allows seamless interoperation between the various networks. But so far, a general need for interoperable grids has not been demonstrated. Nevertheless, at least due to the high dynamicity in IT infrastructure, more specific forms of interoperability are of interests and can be achieved on application and middleware levels. The realization of these forms of interoperability requires a mature and reliable middleware. Unfortunately, current grid middleware implementations do not only fail to interoperate seamlessly but they are also too complex to allow quick appropriate modifications. For the sake of reducing this complexity, it can be expected that some of the necessary grid functionality can be moved to different layers. Security and data integration are key requirements which could be moved from the middleware to the operating system. Other functions, like meta-scheduling and brokering, can be moved up to the community or even to the application levels. In our view, it is necessary that grid researchers and software engineers to establish large scale grid production systems.

In the past, improvements in efficiency, simplicity of use and reduction of cost have always been published reasons for grid computing. In the meantime, commercial players have removed the VRE paradigm of grid computing and provided a new distributed computing on demand concept termed Cloud computing. Due to its simple business model and its less complex technical requirements, Cloud computing has become commercially successful and partially replaced grid computing in public attention. We believe that future grid systems may incorporate Cloud computing on the resource level. But even if a suitable technology is available, still many legal and administrative hurdles must be overcome to achieve these goals.

8.3 Service Level Agreements in Grids

Seminar No. **09131** Date **22.03.–27.03.2009**
Organizers: Hans Michael Gerndt, Omer F. Rana, Wolfgang Ziegler, Gregor von Laszewski

Grid computing allows virtual organizations to share resources across administrative domains. In its early days, Grid computing was inspired by the need for transparent access to supercomputing resources and by the idea to even couple the resources in a meta-computing environment to create even more powerful computational resources. Currently the focus is on service-oriented architectures (SOA) where a wide variety of services from multiple administrative domains can be accessed by service clients.

One of the most important tasks of current Grid middleware centers on efficient resource management. Resource providers offer their resource to virtual organizations and publish detailed information about the resources. Recent efforts have also focused on exposing computational and data resources as "services" – thereby providing a single abstraction that could be applied at different levels of software deployment. Based on this information appropriate resources for Grid applications are selected, and jobs are finally submitted to these resources.

Service Level Agreements (SLA) are attracting more and more attention in Grids as a means to guarantee quality of service terms for grid applications and to enable the establishment of novel business models. A wide range of research and development questions have to be addressed in this context. This covers the creation of languages for formulating SLAs that are powerful enough to express the relevant QoS terms, but can also be used to automatically manage the negotiation, execution, and monitoring of SLAs. Brokering systems are required that can select resources for job execution based on the SLA templates offered by the resource owners. Scheduling algorithms that can optimize for different goals in the context of multi-item, multi-attribute, and multi-unit optimization problems are also necessary. Flexible local resource management algorithms are required for provisioning the resources at the provider's side to meet signed SLAs.

The seminar brought together people working on SLAs in the context of grid computing mainly from computer science, but also from information systems and application areas. These researchers come from different areas and brought in a wide range of research work. The topics covered by the seminar are:

- Languages and protocols for creation of SLA
- Business models
- Grid economy
- SLA management
- Resource management
- Job scheduling
- Application deployment mechanisms
- Negotiation strategies

Agenda

The seminar included the following sessions:

1. SLA application
2. SLA implementations, technologies and approaches
3. SLA negotiation approaches
4. SLA policies and legal issues
5. Interoperability: Standards for describing and creating SLAs
6. SLA applications, status, monitoring and billing

In addition to the above session, working group meetings on WS-Agreement Profiles and on WS-Agreement-Negotiation of the GRAAP working group of the Open Grid Forum were held in conjunction with this seminar.

Chapter 9

Modelling, Simulation, Scheduling

9.1 Sampling-Based Optimization in the Presence of Uncertainty

Seminar No. **09181** Date **26.04.–30.04.2009**
Organizers: Jürgen Branke, Barry Nelson, Warren Powell, Thomas J. Santner

Motivation

There are numerous industrial optimization problems in manufacturing, transportation and logistics, security, energy modeling, finance and insurance, and the sciences where decisions have to be evaluated by a process that generates a noisy result. The process might be a discrete-event simulation, a Monte Carlo evaluation of a complex function, or a physical experiment (e.g., how many cancer cells were killed by a particular compound?). There might be a small number of discrete decisions (the location of an emergency response facility, the design of a compound, or a set of labor work rules), or a large vector of decision variables (the allocation of a fleet of vehicles, choosing a set of research projects or allocating assets among investments). There are applications in virtually any area of business, government, science and engineering. Algorithms to support decisions in these diverse environments are urgently needed. This Dagstuhl seminar focused primarily on problems where this measurement is expensive (for example, some computer models can take a day or more for a single data point), in which case the number of samples that could possibly be generated is rather limited. When the goal is to efficiently identify an optimal (or at least a very good) solution, the search for good solutions, and the collection of information to guide the search, are tightly coupled. It is necessary to strike a balance between collecting information (exploration or global search) and making decisions that appear to be the best given what we know (exploitation or local search). This is particularly true when measurements are expensive (long simulations, field experiments). Because of its wide-ranging applications, sampling-based optimization has been addressed by different communities with different methods, and from slightly different perspectives. Currently, communities are largely tied to problem categories (e.g., finite vs. infinite number of alternatives; discrete vs. continuous decision variables; desired statement at termination).

This Dagstuhl seminar brought together researchers from statistical ranking and selection; experimental design and response-surface modeling; stochastic programming; approximate dynamic programming; optimal learning; and the design and analysis of computer experiments with the goal of attaining a much better mutual understanding of the commonalities and differences of the various approaches to sampling-based optimization, and to take first steps toward an overarching theory, encompassing many of the topics above.

Seminar week

The seminar brought together 31 internationally renowned researchers from 11 countries. After an introductory session the seminar started with four **tutorials** on the various involved communities:

- Jack Kleijnen: Design and Analysis of Experiments: An Overview
- Steven Chick: Ranking and Selection Tutorial
- Barry L. Nelson: A Brief Introduction to Optimization via Simulation
- Warren Powell: Tutorial on optimal learning

Other planned events included two **feature talks**, 15 **regular talks**, and a **panel discussion** on 'Barriers to Application' (panelists Steve Chick, Genetha Gray, Tom Santner and Warren Powell).

A significant amount of time of the seminar was spent in **working groups**.Based on suggestions made by the participants, four working groups were formed to discuss some important and cutting-edge research questions in more detail:

1. Multiobjective optimization under uncertainty
2. Optimization with expensive function evaluations
3. Approximate dynamic programming/optimal learning
4. Cross-fertilization of experimental design, ranking & selection and optimization.

These working groups met each day for 1-2 hours, and presented their results to the general audience on the last day.

Besides the official programme, there were plenty of opportunities for informal discussions, e.g., during lunches, a short hike on Wednesday afternoon and a wine& cheese party on Wednesday evening. Overall, the seminar was a great success and offered many possibilities for cooperation. It was generally agreed that such a workshop should be repeated in two years time.

9.2 Models and Algorithms for Optimization in Logistics

Seminar No. **09261** Date **21.06.–26.06.2009**
Organizers: Cynthia Barnhart, Uwe Clausen, Ulrich Lauther, Rolf H. Möhring

Logistics is the cost aware planning, design, and control of material flow and related information flow (persons, energy, money, information, ...) in production processes. The notion is often used as a synonym for transportation, distribution, or warehouse management. The topic is of a rich variety, has great practical importance, and attracts researchers from the computer science (CS), mathematical programming (MP), and the operations research (OR) communities alike.

Today, problems from logistics are widely studied as parts of the disciplines of mathematical programming and operations research; algorithmics and theoretical computer science; and computer systems. The specific models and methods, as well as the objectives to be optimized, differ in the various disciplines; nevertheless, there are remarkable similarities (as well as significant differences) in the general framework adopted by researchers in logistics in these disparate disciplines.

The primary objectives of the seminar were to bring together leading and promising young researchers in the different communities and practitioners to discuss problems that arise in current and future technology, to expose each community to the important problems addressed by practice and the different communities, and to facilitate a transfer of solution techniques from each community to the others.

There were approximately fifty participants at the seminar, nearly evenly split between computer science, mathematical programming, and engineering and industry.

Six special invited presentations served as introductory lectures on important research areas and application domains and created a common understanding. They were given by George Nemhauser on maritime inventory routing, Jens Baudach and Ronny Hansmann on waste disposal logistics, Ozlem Ergun on humanitarian logistics, Alexander Martin on the power of discrete optimization in logistics, Cynthia Barnhart on trends in airline optimization, and by Patrick Jaillet on probabilistic analysis of routing problems.

This was complemented by an industry day on Tuesday, on which participants from industry and industry-near research institutes presented their research, problems and viewpoints for future research in logistics.

In discussion with the different communities, we organized 27 medium length talks on various recent research results. There was a plenary session on Friday morning to discuss interesting directions for future research and future collaborations. The discussion identified and collected specific needs for future topics such as enabling real time decisions in optimization, a better integration of heuristics and integer programming, dealing with non-observable information through better use of statistic methods, and to exploit game-theoretic aspects in logistics networks.

This seminar was essentially a first meeting of practitioners with the mathematical programming and theoretical computer science community. The general consensus was that both communities learned a lot about the other communities and that it is worthwhile and challenging to continue this form of workshop.

Chapter 10

Cryptography, Security

10.1 Symmetric Cryptography

Seminar No. **09031** Date **11.01.–16.01.2009**
Organizers: Helena Handschuh, Stefan Lucks, Bart Preneel, Phil Rogaway

Topics

Cryptography is the science that studies secure communication in adversarial environments. Symmetric Cryptography deals with two cases:

- either sender and receiver share the same secret key, as for encryption and message authentication;
- or neither sender nor receiver use any key at all, as, e.g., in the case of cryptographic hash functions.

Specifically, Symmetric Cryptography deals with symmetric primitives (block and stream ciphers, message authentication codes and hash functions), and complex cryptosystems and cryptographic protocols employing these primitives. Since symmetric cryptosystems are one to two orders of magnitude more efficient than asymmetric systems, most security applications use symmetric cryptography to ensure the privacy, the authenticity and the integrity of sensitive data. Even most applications of public-key cryptography are actually working in a *hybrid* way, separating an asymmetric protocol layer for key transmission or key agreement from secure payload transmission by symmetric techniques.

Presentations

The seminar brought together about 40 researchers from industry and academia, leading experts as well as exceptionally talented junior researchers. Most of the presentations did concentrate on one of the following three research directions:

1. studying the design and analysis of *stream ciphers*;
2. presenting and attacking recent proposals for *cryptographic hash functions*; and
3. advancing the field of complex symmetric cryptosystems and protocols and their *provable security*.

The great interest in stream ciphers relates to the recently terminated eSTREAM project, under the umbrella of the European Network of Excellence ECRYPT. This initiative has brought remarkable advances in stream cipher design, a.o. by recommending a portfolio of 8 stream ciphers which are believed to be promising for further study. The cryptanalysis of hash functions has made a quantum leap in recent years. As a result, the National Institute of Standards and Technologies (NIST, USA) initiated a competition for a new hash function standard "SHA-3". The list of SHA-3 first-round candidates and their submission documents have been published about one month ahead of the Dagstuhl seminar. That was just enough time for the seminar participants to gain some first insights into strengths and weaknesses of some of the candidates. This constellation was ideal for the Dagstuhl seminar, as it lead to a fruitful exchange of ideas for cryptanalysing SHA-3 candidates, and to intense discussions about the relevance of several weaknesses. Provable security is based on the idea of formally specifying the security requirements a cryptosystem should satisfy, and formally proving that these security requirements are met if certain assumptions hold. In recent years, the research community in Symmetric Cryptography had shown a growing interest in provable security; in the SHA-3 competition, provable security plays an essential role to study the relation between the security of the building blocks and the hash function itself.

Discussion

In an *open discussion session*, many questions where raised, regarding the state of the art in Symmetric Cryptography in general, how the field has evolved in the past and how it will likely evolve in the future, how the community would like it to evolve, whether the research community actually concentrates on the right questions, and so on. One major issue, which raised substantial interest among the participants was the following:

> There is a broad range of abstract techniques to study the security of symmetric primitives, such as Differential Cryptanalysis, Linear Cryptanalysis, Algebraic Attacks and so on. But in many cases, a researcher who is trying to apply these techniques needs tools (typically software), e.g., to compute the difference distribution table of a cipher, a round function or an S-box or to find the best linear or differential characteristic of an iterated cipher. It turns out however that each researcher or each group of researchers develops such tools on their own from scratch.

The general agreement was that the research community would benefit from establishing a culture of tool reuse, by encouraging researchers to share not only their ideas, but also the software they developed for the purpose of analyzing cryptosystems.

Summary

Research in Symmetric Cryptography is quickly evolving. The seminar was the second of its kind, the first one took place in 2007. We observe a steadily increasing interest in Symmetric Cryptography, as well as a growing practical demand for symmetric algorithms and protocols.

The seminar was very successful in discussing recent results and sharing new ideas. Furthermore, it inspired the participants to consider how Symmetric Cryptography has evolved in the past, and how they would like it to evolve in the future. The hospitality and support of the Dagstuhl team did contribute significantly to the success of the seminar.

10.2 Web Application Security

Seminar No. **09141** Date **29.03.–03.04.2009**
Organizers: Dan Boneh, Ulfar Erlingsson, Martin Johns, Benjamin Livshits

Security of Web applications has become increasingly important over the last decade. This is not at all surprising: Web applications are now ubiquitous, spanning the spheres of e-commerce, healthcare, finance, and numerous other areas. More and more Web-based enterprise applications deal with sensitive financial and medical data, which, if compromised, in addition to downtime can mean millions of dollars in damages. It is crucial to protect these applications from malicious attacks. Yet, to date, a great deal of attention has been given to network-level attacks such as port scanning, even though, about 75% of all attacks against Web servers target Web-based applications, according to recent surveys. Traditional defense strategies such as firewalls do not protect against Web application attacks, as these attacks rely solely on HTTP traffic, which is usually allowed to pass through firewalls unhindered. Thus, attackers typically have a direct line to Web applications. Furthermore, traditional vulnerabilities such as buffer overruns, pervasive in applications written in C and C++, that have been the subject of intense for over a decade are now largely superseded by Web applications vulnerabilities such as cross-site scripting, SQL injection, and session riding attacks.

Web applications have progressed a great deal in the last decade since their humble beginnings as CGI scripts. Todays Web applications are sophisticated multi-tier systems that are built on top of complex software stacks. Web applications are also distributed: a Web application typically includes both a server-side component running on top of an application server such as JBoss, as well as a client-side component that usually consists of HTML and JavaScript. Consequently, Web application security touches upon many aspects of systems research. The topic of Web application security has attracted researchers from diverse backgrounds in recent years. In addition to core security experts, this includes specialists in programming languages, operating systems, and hardware. Similarly, the research directions proposed so far range from improving security through Web browser changes to low-level hardware-level support and in-depth analysis of server code. Last but not least, much work remains to be done in social engineering for security as applied to Web applications.

The last several years have seen dramatic changes in Web application development. We are now in the middle of the Web 2.0 revolution, triggered by demand for better, more interactive user experience and enabled by Ajax (asynchronous JavaScript and XML). However, extra functionality of rich-client applications is generating new security concerns. A good example of that is JavaScript worms, which first emerged in 2005 and have grown increasingly popular in the last year or so. JavaScript worms take advantage of the ability of the Web client to programmatically issue server requests through Ajax to propagate malicious payload.

The seminar was well attended with 38 participants. A good balance of European and American researchers was present. Furthermore, the group represented a nice mix of participants of academia and industry (including members of companies such as Mozilla, Microsoft, SAP, and Google).

This was the first Dagstuhl seminar on Web application security. In addition, academic research on this topic is a rather young discipline. For this reason, the seminar's organisation favored presentations over open workgroups or plenum style discussions. This way, a good, comprehensive view on current activities and open problems in the realm of Web application security could be achieved.

Since the seminar took place, the underlying research of most talks has been presented at conferences and the corresponding papers have been published in the associated proceedings. Hence, we list a comprehensive list of publications that are directly associated with the seminar's content in the bibliography of this document.

The seminar was perceived as highly inspiring by the participants. In consequence, it had a fertilizing effect on follow-up activities: Besides various informal collaborations that resulted from discussions in Dagstuhl, we would like to single out two results which directly can be attributed to the seminar: For one, during the seminar the observation was made, that Europe at that point in time did not offer a compelling venue for academic Web application research. For this reason, a set of present participants decided to pursue this issue. The result of this effort was the OWASP AppSec Research conference, which had its first iteration in June 2010 in Stockholm. Furthermore, based on initial discussions during the seminar, a consortium formed for further collaboration in a larger research project. This resulted in a successful proposal for a EU FP7 project. Out of the five primary drivers of the proposal, four (in the form of the seminar participants from SAP, Chalmers, KU Leuven, and Uni Passau) had met at the seminar. The project is called WebSand and will start in October 2010 its three year run. It will target research questions in the field of Web application security in multi-party scenarios.

The dominant result of the seminar was that the field of Web application security research simply does not exist. Instead, the topic is approached from a highly heterogeneous set of directions, ranging from low-level vulnerability countermeasures, through ad-hoc run-time enforcement mechanisms, over security protocol analysis, to fully formalized typing approaches. Research in this field has to be agile and versatile as even the most fundamental building blocks of the young application paradigm are still evolving and constantly changing – sometimes for the better, sometimes for the worse from a security point of view. The fight for secure Web applications is still an uphill battle. We live in interesting times.

10.3 Foundations for Forgery-Resilient Cryptographic Hardware

Seminar No. **09282** Date **05.07.–08.07.2009**
Organizers: Jorge Guajardo Merchan, Bart Preneel, Ahmad-Reza Sadeghi, Pim Tuyls

Motivation

The rapid expansion of global connectivity, distributed applications and digital services over open networks and across organizational domains requires secure IT systems that adhere to well-defined policies. Cryptography and technical IT security mechanisms support the establishment of secure channels and authorized access. However, many of today's IT applications demand sophisticated security and privacy mechanisms in both software and hardware that go beyond secure channels and authorization and include truly secure liaisons: Enterprises or manufacturers outsource their computations, data storage, and production to potentially untrusted parties over which they have limited control. Medical records are transmitted through and processed by various IT systems such as Handhelds, PCs or hospital servers. Biometric data are carried by individuals on their ID card or electronic passport. Fake and counterfeited pharmaceuticals or automotive and avionic spare parts are packaged in some countries and distributed illegally to worldwide destinations.

IT system security is, however, not only based on strong cryptographic primitives and protocols but also on technological support for secure implementation of the corresponding algorithms. In particular, this concerns security functionality provided by the underlying hardware, which is commonly deployed in the form of cryptographic hardware. The study of how to model, design, evaluate and deploy such cryptographic hardware was the focus of our seminar.

The recent trend of deploying security functionality in hardware typically assumes trust in the various parties involved in the design and manufacturing of the hardware. The life-cycle of cryptographic hardware begins with the IC design step, which results in IC blue-prints being shipped for production to (typically overseas) low-cost manufacturer's facilities. This trend is driven by economic and strategic reasons as well as by globalization. Although this model has many advantages, it also has the disadvantage that it becomes much easier for attackers to compromise hardware devices commonly used in critical infrastructure, which includes commercial, health and defense applications.

As a result, today many ICs and components are overbuilt (over-produced in an unauthorized manner). This, in turn, allows such devices and components to enter the market through gray channels and erode the revenues of legitimate Intellectual Property (IP) owners. In addition, there is a high risk that the functionality on the chip is (deliberately) modified or supplemented with a hidden trapdoor circuit, e.g., a hardware Trojan. For instance, keys which were never supposed to leave a security chip might be leaked (e.g., via padding), the tamper or leakage protection circuits of a chip may be disabled or weakened, a True Random Number Generator may be biased or the IC might have a kill switch that makes it stop functioning under certain conditions. Even in the non-malicious case, overseas manufacturers may try to cut costs by omitting or reducing security measures from

the original design. Any single one of these manufacturing attacks or malpractices will have serious consequences for any security application, allow industrial espionage, privacy violations, and finally even threaten national security.

Current methods for assuring the trustworthiness of cryptographic hardware rely heavily on the skills of an evaluator. The lack of standardized methodologies and tools requires that the evaluator correctly identifies and manually evaluates each risk area. The evaluator must be aware and execute all known attacks while also formulating and exercising new forms of attack. The evaluator knows only what was found, not what's left to be found. More resources are used to obtain higher levels of assurance with the ultimate measure of assurance being what happens once the product is in production or it has been deployed. Advances in commercially viable approaches to assure the security of hardware is critical. From defining systematic approaches for assurance to identifying tools to automate and continuously improve assurance levels, significant new research is required. Moreover, commercial hardware engineering practices are well behind software engineering when it comes to establishing a set of best practices that will yield high-quality security products. Existing methods developed for high assurance hardware, typically for use by governments, either break down when considering the size of designs (e.g., microprocessors) or are unacceptable from an economic perspective. Thus, a systematic approach with a solid scientific basis is required to ensure that hardware as the security anchor (or root of trust) for computing will deliver the necessary security guarantees.

Objectives and Goals of the Seminar

Based on the previous discussion, it is clear that there is an urgent need to design and develop methods that increase the security and trust in current hardware solutions. The purpose of this seminar was to bring together researchers from academia and industry and from different disciplines (cryptography, information theory, theoretical and experimental physics, hardware architectures and processor design), and allow them to investigate a whole new set of security and cryptographic methodologies which will allow for the development of reliable and trustworthy hardware components. Such trustworthy components will constitute the "root of trust" for future generation security devices and applications.

We have identified as the main challenges to provide strong, cost-effective and easily deployable methodologies and technological means to solve the following issues:

Exploiting inherent nano-scale physical properties (randomness) in hardware as a new key feature for a new level of security:

- The randomness caused by inherent variations in the hardware manufacturing process can be exploited to uniquely identify devices. In this context the most promising and interesting recent development based on primitives called Physically Unclonable Functions (PUFs), which are functions embodied in a physical structure. Due to their random structure a physical stimulus/challenge generates an unpredictable response which can be used for the purpose of device authentication. Regardless of their particular instantiation, the unclonability, tamper-evidence and tamper-resistance properties of PUFs are very useful tools

in anti-counterfeiting, secure secret key storage or binding software components to the underlying hardware.

- Investigating what sources of randomness we can exploit for this purpose, and how to use them efficiently.
- Integrating components based on unique physical properties into cryptographic primitives and security protocols, and investigating the security properties achieved by such systems.
- Investigating the construction of cost-effective and easy to use "Reconfigurable Physically Unclonable Functions (rPUF)" that can be physically reconfigured.

A framework offering provable security which is based on physical properties: We aim to discuss appropriate models and methodologies to realize and to analyze the security of resulting cryptographic primitives and security protocols that concern the following aspects:

- Manufacturing security: Preventing/detecting overproduction and ensuring security in the commercial manufacturing environment also under insider threats.
- Identification and evaluation of malicious (Trojan) and unspecified functionality in hardware: Ensuring the trustworthiness and full functionality of security sensitive ICs. Recent research results indicate that new hardware components are required to achieve this goal.
- Anti-Counterfeiting, verifiability and auditability of security critical devices: Investigating hardware and system components that are needed and economically implementable to prevent or detect counterfeited devices.
- Trade-off unique device identification versus privacy: Unique identification of objects stands clearly in contrast with privacy. In particular, in the medical device setting, it is, on the one hand, important to uniquely identify devices for reasons of security and safety, and on the other it is important to provide mechanisms enabling access control to this unique identifying information. This can include merely protecting the existence of the device, device type, or its ID, or the confidential information stored on it or broadcasted by it.
- Dynamic and distributed Trusted Computing: Designing security modules with dynamic trusted computing functionality, i.e., a minimum root of trust both for PC and mobile scenario where various cryptographic functionalities can be securely generated and loaded when needed. In particular we aim at investigating the questions such as what functionality does it really need to be included inside the trust boundary, how can we verify the trusted functionality in a meaningful way and how can we distribute trusted functionality over several ICs on the platform?

The relevance of the previously mentioned problems is only made clearer by looking at recent developments and trends in the commercial deployment of cryptographic hardware. Prominent examples include Intel's Trusted Execution Technology and next generation CPUs, AMD's Presidio, and the TPM (Trusted Platform Module) proposed by the Trusted

Computing Group (TCG). Moreover, future generations of CPUs are expected to provide a variety of cryptographic functions, all embedded into a single chip set. Their deployment is also the subject of large European projects such as OpenTC or TECOM.

The goals and challenges mentioned above comply with the objectives and challenges of secure, dependable and trusted infrastructures and bridge the gap between the current black-box security models and the real world we live in. Given recent important advancements and developments in the area of cryptographic hardware that concern many various disciplines, we expected this Dagstuhl seminar to be an appropriate platform for experts from various disciplines to benefit from the mutual exchange of ideas across these research communities. In addition, we hoped that the results of the discussions and interactions during the seminar would become the corner stone in theoretical and practical foundations for forgery-resilient cryptographic hardware.

The participants

The seminar counted with the participation of 30 researchers, who are currently working in the following countries:

> Belgium(8), Canada (1), Germany (10), Great Britain (1), Israel (1), The Netherlands (2), Poland (1), Switzerland (1) , United States (5)

These researchers brought to the seminar a rich variety of backgrounds in computer science and engineering. These included theoretical and practical cryptography, algorithms design, chip design, VLSI, low power design, system security, security evaluation, side-channel countermeasures and attacks, design of cryptographic primitives for constrained environments, and standardization. The diverse backgrounds created a stimulating atmosphere and allowed for interesting discussions.

Concluding Remarks

We found the seminar to be fruitful in the sense that several modeling issues were raised, which we expect will lead the community to understand better the security issues and requirements of forgery resilient hardware. In addition, the participation of both, theoretical computer scientists and more implementation oriented scientists, allowed for a better understanding from both sides: what models are realistic, what needs to be formalized to be able to prove security of an implementation, and what emerging applications of security hardware exist.

Moreover, it appears that the formal modelings of hardware primitives and the subsequent deployment of such hardware will remain hot topics for the next few years. In the future, we plan further workshops to encourage continued interdisciplinary interactions.

10.4 Classical and Quantum Information Assurance: Foundations and Practice

Seminar No. **09311** Date **26.07.–31.07.2009**
Organizers: Samuel L. Braunstein, Hoi-Kwong Lo, Kenny Paterson, Peter Y. A. Ryan

From 26 July 2009 to 31 July 2009, the Dagstuhl Seminar 09311 "Classical and Quantum Information Assurance Foundations and Practice" was held in Schloss Dagstuhl – Leibniz Center for Informatics. The workshop was intended to explore the latest developments and discuss the open issues in the theory and practice of classical and quantum information assurance. A further goal of the workshop was to bring together practitioners from both the classical and the quantum information assurance communities. To date, with a few exceptions, these two communities seem to have existed largely separately and in a state of mutual ignorance. It is clear however that there is great potential for synergy and cross-fertilization between and this we sought to stimulate and facilitate.

The program included tutorials from both communities aimed at bringing members of the the other camp up to speed:

- Intro to modern cryptography (Bart Preneel)
- Intro to provable security (Kenny Paterson)
- Intro to the modelling and formal analysis of cryptographic protocols (Peter Ryan)
- Intro to the theory of quantum cryptography (Charles Bennett)
- Towards quantum key distribution with testable assumptions: a tutorial (Hoi-Kwong Lo)
- Introduction to Universal Composability (Dominique Unruh)
- Practical aspects of QKD (Gregoire Ribordy)

The workshop generated stimulating and at times heated debates on the merits and demerits of quantum cryptography. A participant from the conventional cryptography community claimed that quantum cryptography is essentially useless in practice because of its high cost, low key rate, short distance, limited applications and the need to distribute the initial authentication key material. Moreover, his view was that quantum cryptography is not an effective counter-measure against the threat of quantum computing. He believed that public key cryptographic systems such as NTRU and McEliece could be used, if a quantum computer were ever built in future.

The quantum community countered as follows. First, there is a need for top secret long-term security and quantum cryptography can never reduce security. Second, to break a quantum cryptographic system, one needs to eavesdrop today because there is no classical transcript for a quantum communication. This means an eavesdropper has to invest in

quantum technologies in order to eavesdrop. Third, current technological limitations of quantum cryptography such as key rate and distance may be overcome in future. For instance, quantum repeaters could, in principle, extend the distance of quantum cryptography arbitrarily. Fourth, the cost of the quantum cryptographic systems may be absorbed through savings in multiplexing of optical channel in telecom fibers. Fifth, since few quantum people are working on breaking NTRU or McEliece crypto-systems these days, the security of those systems against quantum attacks is largely unknown.

Perhaps, a more balanced view to take is that it is important to explore future cryptographic infra-structure. Quantum cryptography, while probably not the only solution, may well play a part in such a future infra-structure.

Chapter 11

Data Bases, Information Retrieval

11.1 Interactive Information Retrieval

Seminar No. **09101** Date **01.03.–06.03.2009**
Organizers: Norbert Fuhr, Nicholas Belkin, Joemon M Jose, Cornelis J. van Rijsbergen

Introduction

Interactive information retrieval (IIR) systems are a commodity nowadays; however, the scientific foundation of this type of system is rather limited. Information retrieval (IR) theory has widely ignored this area, and cognitive IR approaches have not yet led to detailed specifications for IIR systems. Within this context, we organized a week long seminar at Dagstuhl during March 2009 and the activities and recommendations are described in this report.

The general idea was to collect the state of the art in IIR research, and to define a research agenda for further work in this area. For this purpose, we brought together experts from the related areas such as information science, cognitive science, interactive IR, theoretical IR and humancomputer interaction (HCI). We took a broad approach to the problem of IIR by highlighting latest results and naming crucial research issues. Based on these contributions, we identified open research problems and then point out steps towards resolving these problems.

Organization of the seminar

We had 32 participants from across the globe working on issues related to interactive information retrieval. The workshop started on a Monday and finished on Friday. The organization of the workshop included keynote talks (3), short talks from participants, demonstration sessions, special topic sessions, and breakout sessions. In addition, we had an afternoon visit to a nearby ancient city.

The technical activities focused around: evaluation methodology in information retrieval; adaptive and personalized retrieval; context and interfaces; and, semantic search. The three keynote talks were on: Cognitive & Context Modeling for Interactive IR; Evaluation of Interactive Retrieval systems; User Interfaces for Interactive Information Retrieval

Nick Belkin gave the first keynote talk with a historical overview of how cognitive (and other) models of users of interactive IR (IIR) systems have been elicited, constructed and used. Cognitive models in the domain of interactive information retrieval (IIR) are understood as models that a system (or a person) constructs of a(nother) person's "information need"; these are called "user models". Context models are models that a system (or a person) constructs of the conditions that led a(nother) person to engage in information seeking behavior, various characteristics of that person, and various aspects of that person's environment, broadly construed. Both types of models are used in personalization of IIR.

Belkin discussed the topic by reviewing Robert Taylor's 1968 article [Taylor 1968] in which he proposed four levels of "information need" or "query", and five "filters" according to which the librarian and the information seeker identify and clarify (i.e. model) various aspects of the user, the user's goals, the topic of interest and so on. He then presented various approaches to understanding of why people engage in information seeking behaviour, and of automatically constructing cognitive and contextual models. With respect to interactive developments, significant change points include the "cognitive turn" in the early to mid 1970s, modeling the human intermediary in the 1980s, and the "interactive turn" in the 1990s. He then outlined current state of cognitive and contextual modeling in IIR, which include cognitive modeling of need, intention and recently incorporating inclusion of individual characteristics becoming more significant. In addition, contextual modeling of environmental factors being used and also contextual modeling of social factors becoming recognized as significant. There are also attempts to model longterm needs.

Maristella Agosti presented the second keynote and highlighted the need for modeling, organizing and managing scientific data produced in an evaluation campaign. In general, user studies and logs are used in a separate way, since they are adopted with different aims in mind. It seems more scientifically informative to combine logs together with observation in naturalistic settings. A systematic use of triangulation of different data collection techniques is needed as a general approach in order to get better knowledge of the Web information search process [Pharo & Järvelin 2004]. Taking inspiration from this general approach, a method of combining implicit and explicit user interaction data to gain information to be used for personalization purposes is outlined. The argument is that data log analysis can be combined with the results of data derived from user studies to evaluate information access services. Further, Agosti argued for using digital library systems as a tool to do this and presented a case study demonstrating this idea.

Harald Reiterer presented the third keynote talk giving an indepth survey of interactive interfaces used for information access highlighting the lessons learned from these activities. Recent developments in interface technology are surveyed. Subsequently, the group discussed the role of interfaces in information seeking. Information is only useful when people interpret it in the context of their goals and activities. In order to design technologies that better support information work, it is necessary to better understand the details of user activity. In this context, the need for further studies on user information activity is needed.

Each of these keynotes was followed by short presentations from participants. Subsequently we formed shorter discussion groups, which are described briefly in the following.

Evaluation methodology

Starting from the Cranfield paradigm of evaluation methodology, we critically looked into the effect of searcher behaviour and the searcher's goals. The ultimate goal of information retrieval (IR) is to support humans to better access information in order to better carry out their tasks. Therefore IR research should provide methods and techniques to improve the retrieval/access process. In this regard Kalervo Järvelin led the debate and argued that IR evaluation methodology, in particular that based on Cranfield methodology, is not focusing its efforts properly to serve the usercentred goals. He argued further that the Cranfield paradigm of evaluation tends to lose its power as soon as one looks at human performance instead of system performance. Also, searchers using IR systems make use of rather unorthodox queries (from the Cranfield pointofview) and sessions. Their search strategies have not been sufficiently described and cannot therefore be properly understood, supported or evaluated. Moreover, searchers engage in an information seeking process, which they have found effective enough based on their previous searching experiences. They try not to optimize the search result alone but the entire process (and effort) and its expected contribution to their primary task. Järvelin argued that this can be better understood in terms of the management science theory "incrementalism" than in terms of rationalism.

At the same time, the merits of current evaluation methodology in benchmarking various systems have been highlighted. We looked into the role of test collections in IR and emphasized their role as corner stones of IR evaluation. Sanderson led the debate and highlighted the papers that support test collection based evaluation [Sanderson 2009]. It is also argued that user experiments are slow to set up and expensive. Often they are not large enough to support any conclusive evidence in support or against the hypotheses. However, there are unsolved issues with respect to test collections. On the one side, there is strong evidence for relevance feedback and pseudo relevance feedback from test collectionbased evaluations [Mitra et al. 1998]. Unfortunately such techniques have not been taken up for public utilization and we need to study why they haven't.

We also discussed simulated evaluation methodologies. One of the difficulties in interactive evaluation is the time needed to setup experiments and the costs involved in terms of experimentation. These get exacerbated if we need to test multiple retrieval models. On the other hand classical IR evaluation methodology fails to consider interactive elements. An alternative is simulated evaluation in which the idea is to simulate all possible user interactions that might have happened in an actual usage of these systems. Using the ground truth given in test collections these strategies can be run and measurements can be taken. This allows one to benchmark various interactive retrieval models. However, this process will not consider the cognitive issues involved in user interaction. Hence it is important to conduct follow on user testing possibly with reduced number of interactive models. It is very clear that this methodology needs further consideration in terms of simulation methods, measures etc. A possible alternative to this type of evaluation is to identify one or more specific types of users, limiting the user models to lead to prototypical interactive behaviors. The current difficult with this approach is in developing credible and valid user models and associated behaviors.

Personalization, Adaptation and Context

Ann Blandford [Blandford 2009] pointed out that many of the current IIR systems are based on wrong assumptions about users and their behaviour. Thus, a better understanding of the human activity is needed when building new systems, aiming at covering the whole process of interacting with information: starting with the information need, followed by the information acquisition stage, then the found information is interpreted and finally used. Effective support of this process is only possible if the system takes the usage context into account and also allows for personalization and adaptation

We spent serious effort on discussing personalization research in the context of information retrieval. Adaptive IR may include adaptation of system features based on nonuser factors, on the other hand personalization of IR is a subset of adaptive IR and is explicitly concerned with userbased factors. Personalization may be the more interesting, more difficult, and more fruitful approach.

Belkin highlighted the facets of personalization: Relevance/usefulness/interest; Task; Problem state; Personal characteristics; Personal preferences; and Context/situation. There are lots of studies on investigating single facets in personalization however, not much study on integrating multiple facets and recommended further investigation is needed along these lines.

Context is an important factor in the information seeking process. There are many definitions of context and it is important to define this concept and highlight its role in information retrieval process. Context models are models that a system constructs of the conditions that led a user person to engage in information seeking behaviour. There are many facets of context and it can support understanding as well as retrieval. However, there are many technical challenges that need to be addressed. These include: what features of context can be used? How to elicit and represent those features? How to combine these features into a retrieval process? How to evaluate such a system?

Theoretical modeling of Interactive Information Retrieval Systems

We also discussed the issues in modeling IIR systems. The classical Probability Ranking Principle (PRP) forms the theoretical basis for probabilistic Information Retrieval (IR) models, which have dominated IR theory for about 20 years. However, the assumptions underlying the PRP often do not hold, and its view is too narrow for interactive information retrieval (IIR).

Norbert Fuhr presented a new theoretical framework for interactive retrieval [Fuhr 2009]. The basic idea is that during IIR, a user moves between situations. In each situation, the system presents to the user a list of choices, about which s/he has to decide, and the first positive decision moves the user to a new situation. Each choice is associated with a number of cost and probability parameters. Based on these parameters, an optimum ordering of the choices can the derived the PRP for IIR. Fuhr highlighted the relationship of this rule to the classical PRP and pointed out issues for further research. Massimo Melucci introduced a geometric model and its investigation for contextual information

retrieval [Melucci 2009]. The geometric model leverages recent advances of vector space-based information retrieval.

The group observed that there is a lack of research activities in modeling of interactive IR systems and recommended this as one of the necessary action points.

Recommendations

In the closing session of the workshop, the group identified many research areas and highlighted the following recommendations:

1. Further effort is needed to define an evaluation methodology that can effectively evaluate context sensitive information retrieval systems. In this regard, the role of interactive test collections needs to be explored. In addition, the simulated evaluation methodology needs to be studied further.

2. Related to this the development of new evaluation measures, which evaluate system performance in terms of the entire information seeking interaction, rather than only in terms of the response to a single query.

3. There is an urgent need to define the concept of context and to study its exploitation in interactive information retrieval systems.

4. Theoretic models of interactive retrieval systems are very important. Serious efforts are needed to develop models that fits various interactive search scenarios

5. In order to reduce the effort for performing useroriented evaluations, cooperation between research groups should be enforced and appropriate evaluation initiatives be launched.

References

Blandford, A. (2009). Interacting with information. In this volume.

Fuhr, N. (2009). A Probability Ranking Principle for Interactive IR. In this volume.

Melucci, M. (2009). Contextual Information Retrieval Using Vector Spaces. In this volume.

Mark Sanderson (2009). It's nice and warm in the cave. In this volume.

Mitra M, Singhal A, Buckley C. (1998). Improving Automatic Query Expansion. In: ACM SIGIR '98 Proceedings, pp 206-214

Taylor, R.S. (1968). Question-negotiation and information seeking in libraries. College and Research Libraries, 28:178-194.

Pharo, N. & Järvelin, K. (2004). The SST method: a tool for analyzing web information search process,. Information Processing & Management, 40: 633-654

Chapter 12

Machine Learning

12.1 Similarity-based Learning on Structures

Seminar No. **09081** Date **15.02.–20.02.2009**
Organizers: Michael Biehl, Barbara Hammer, Sepp Hochreiter, Stefan C. Kremer, Thomas Villmann

The seminar centered around different aspects of similarity-based clustering with the special focus on structures. This included theoretical foundations, new algorithms, innovative applications, and future challenges for the field.

For finding the structure in the data set's smothers
many tools are related like sisters and brothers.
We conclude in the sequel:
All methods are equal!
(But some are more equal than others.)

Goals of the Seminar

Similarity-based learning methods have a great potential as an intuitive and flexible toolbox for mining, visualization, and inspection of large data sets across several disciplines. While state-of-the-art methods offer efficient solutions for a variety of problems such as the inspection of huge data sets occuring in genomic profiling, satellite remote sensing, medical image processing, etc. a number of important questions requires further research.

The detection, adequate representation, and comparison of structures turned out to be one key issue in virtually all applications. Frequently, learning data contain structural information such as spatial or temporal dependencies, higher order correlations, relational dependencies, or complex causalities. Thus, learning algorithms have to cope with these data structures. In this context, various qualitatively different aspects can be identified: often, data are represented in a specific structured format such as relational databases, XML documents, symbolic sequences, and the like. Similarity based learning has to identify appropriate preprocessing or similarity measures which facilitate further processing.

Several problem formulations are ill-posed in the absence of additional structural information e.g. due to a limited availability of labeled examples for high dimensional data. The dimensionality of microarray data, mass spectra, or hyperspectral images, for example, usually exceeds the number of labelled examples by magnitudes. Structural information can offer effective means for regularization and complexity reduction. More and more learning tasks require additional structural information instead of simple vectorial outputs such as multiple output values, hierarchies, dependencies, or relational information, as required for the inference of biological networks or the analysis of social graphs, for example.

The aim of the seminar was to bring together researchers who develop, investigate, or apply machine learning methods for structure processing to further advance this important field. The focus has been on advanced methods which have a solid theoretical background and display robust and efficient performance in large-scale interdisciplinary applications.

Structure

32 experts from 12 countries joined the seminar representing a good mixture of established scientists and young researchers. According to the interdisciplinary topic researchers from computer science, mathematics, physics, and related subjects as well as people working in industry came together to discuss and develop new paradigms in the area of structural data processing and learning on structures. During the week 29 talks and short presentations were given which adress different aspects of similarity based learning on structures which could be grouped into clusters on the following topics:

- Structural data processing for biology and medicine
- theoretical aspects of learning for high-dimensional and structered data
- discrete methods for structured data
- stability and quality assessment of data processing in the context of structures
- mathematical aspects of uncertain decisions
- structure-adapted non-standard metrics
- prototype based classification and learning algorithms for stuctures and structured data

The talks were supplemented by vivid discussions based on the presented topics and beyond. Additionally, the talks were complemented by expert software demonstrations which immediately gave an impressive view onto the ability of the presented methodologies. The evening wine and cheese sessions as well as the Wednesday excursion to a local brewery and the manufactory Villeroy&Boch gave ample opportunity to deepen scientific discussions in a relaxed and stimulating atmosphere.

Results

A variety of open problems and challenges came up during the workshop week. In particular, the following topics and their interplay were in the focus of several discussions:

- **feature extraction:** Feature selection is one of the main recent topics in classification. Thereby, the task dependent adaptation of predefined data structure models was in the foci of several talks. The methods range from metric adaptation to information theory based selection schemes. The latter follow the naturally inspired paradigm of sparseness while information flow is maximized. The former metric adaptation based approaches optimize the feature set by minimization of classification accuracy. Thereby, classification accuracy has to be defined carefully to cope with the discontinuity of the usual classification error.

- **cluster generation/evaluation:** Cluster generation and evaluation strongly depend on the underlying predetermined similarity/dissimilarity measure applied to the data. The data may be similar according to one measure but may differ heavily with respect to another one. Hence, the choice is crucial for adequate detection of relevant structures and has to be in agreement with the task at hand. The contributions during the seminar presented various approaches taylored according to the needs of different application related problems. Examples are the metric adaptation in quadratic forms for discriminative low-dimensional class representation, development of adaptive kernels or metric adapted multi-dimensional scaling under the specific aspect of high throughput.

- **graph methods for discrete data:** Clustering and classification on graph structures typically require a huge amount of computational costs. Therefore, adaptive methods for approximative solutions are highly desirable. Here the contributions in the seminar were mainly dedicated to the problem of clustering of graphs under the specific restrictions of optimized granularity (in terms of modularity) and the additional requirement of minimization of crossing edges after projection into the plane. The application areas of such problems range from visualization of social networks to dynamics of epidemies.

- **complexity reduction by utilization of structure:** The complexity of data processing of structured data frequently could be reduced if the specific structural information is taken into account. For example, vectorial representation of functions differs from usual data vectors by the spatial dependencies within the vectors. Yet, the utilization of the Euclidean metric disregards this information. During the workshop several approaches for functional metrics were discussed and how they can be incorporated into adaptive methods for data processing.

All in all, the presentations and discussion (often until late at night) revealed that similarity based learning on structures constitutes a highly evolving field. Significant progress has been achieved in recent years and was highlighted during the seminar. Although promising results and approaches were developed, many important problems still await satisfactory

practical solutions. For example, the functional aspect of data is not sufficiently exploited in many data processing methods. Another challenge is the sparseness of data in high-dimensional data analysis and adequate processing tools.

12.2 Machine Learning Approaches to Statistical Dependences and Causality

Seminar No. **09401** Date **27.09.–02.10.2009**
Organizers: Dominik Janzing, Steffen Lauritzen, Bernhard Schölkopf

The 2009 Dagstuhl Seminar "Machine Learning approaches to Statistical Dependences and Causality", brought together 27 researchers from machine learning, statistics, and medicine.

Machine learning has traditionally been focused on prediction. Given observations that have been generated by an unknown stochastic dependency, the goal is to infer a law that will be able to correctly predict future observations generated by the same dependency. Statistics, in contrast, has traditionally focused on data modeling, i.e., on the estimation of a probability law that has generated the data.

During recent years, the boundaries between the two disciplines have become blurred and both communities have adopted methods from the other. However, it is probably fair to say that neither of them has yet fully embraced the field of causal modeling, i.e. the detection of causal structure underlying the data. This has different reasons.

Many statisticians would still shun away from developing and discussing formal methods for inferring causal structure, other than through experimentation, as they would traditionally think of such questions as being outside statistical science and internal to any science where statistics is applied. Researchers in machine learning, on the other hand, have too long focused on a limited set of problems neglecting the mechanisms underlying the generation of the data, including issues like stochastic dependence and hypothesis testing – tools that are crucial to current methods for causal discovery.

Since the Eighties there has been a community of researchers from statistics, computer science, and philosophy, who in spite of the pertaining views described above have developed methods aiming at inferring causal relationships from observational data, building on the pioneering work of Glymour, Scheines, Spirtes, and Pearl. While this community has remained relatively small, it has recently been complemented by a number of researchers from machine learning. This introduces a new viewpoint on the issues at hand, as well as a new set of tools, such as nonlinear methods for testing statistical dependencies using reproducing kernel Hilbert spaces, and modern methods for independent component analysis.

The goal of the seminar was to discuss future strategies of causal learning, as well as the development of methods supporting existing causal inference algorithms, including recent developments lying on the border between machine learning and statistics such as novel tests for conditional statistical dependences.

The Seminar was divided into two blocks, where the main block was devoted to discussing state of the art and recent results in the field. The second block consisted of several parallel brainstorming sessions exploring potential future directions in the field. The main block contained 23 talks whose lengths varied between 1.5 hours and 10 minutes (depending on whether they were meant to be tutorials or more specific contributions).

Several groups presented recent approaches to causal discovery from non-interventional statistical data that significantly improve on state of the art methods. Some of them allow for better analysis of hidden common causes, others benefit from using methods from other branches of machine learning such as regression techniques, new independence tests, and independent component analysis. Scientists from medicine and brain research reported successful applications of causal inference methods in their fields as well as challenges for the future.

In the brainstorming sessions, the main questions were, among others, (1) formalizing causality, (2) justifying concepts of simplicity in novel causal inference methods, (3) conditional independence testing for continuous domains.

Regarding (1), the question of an appropriate language for causality was crucial and involved generalizations of the standard DAG-based concept to chain-graphs, for instance. The session on item (2) addressed an important difference between causal learning to most of the other machine learning problems: Occam's Razor type arguments usually rely on the fact that simple hypotheses may perform better than complex ones even if the "real world" is complex because it prevents overfitting when only limited amount of data is present. The problem of causal learning, however, even remains in the infinite sample limit. The discussion on conditional independence testing (3) focused on improving recent kernel-based methods.

Chapter 13

Bioinformatics

13.1 Formal Methods in Molecular Biology

Seminar No. **09091** Date **22.02.–27.02.2009**
Organizers: Rainer Breitling, David Roger Gilbert, Monika Heiner, Corrado Priami

The Life Sciences, and in particular Molecular Biology, are a rather new application area for advanced computational concepts. Living systems, from cells to entire organisms, function by the complex, dynamic interaction of a large number of components (proteins, nucleic acids, metabolites). The set of "molecular players" continues to be explored in genome sequencing projects and related experiments. Their physical and regulatory relationships are determined in detailed molecular studies and represented in cellular "wiring diagrams" and "flow charts". Such schematic pictures are used by biologists to reason about the expected behavior of biological systems, e.g. in response to disease processes or drug treatment. They can also be translated into quantitative mathematical descriptions of the system. With the recent explosion of biological knowledge, such approaches need to become more common and more formalized.

Formal logical models play an increasing role in the newly emerging field of Systems Biology. Compared to the classical, well-established approach of modeling biological processes using continuous and stochastic differential equations, formal logical models offer a number of important advantages: **Easy compositionality**, which allows the generation and management of large cellular models from a number of pre-defined and reliably manipulated building blocks; **model checking** for the rigorous exploration of model consistency, including the comprehensive exploration of state-space and the identification of necessary additions to an existing system description; **unambiguous visualization** based on the strictly enforced syntax of the modeling language. In addition, a number of recent studies have explored the combination of formal logical models with continuous and stochastic differential equation models, showing important relationships between the two approaches and further expanding the expressivity of the resulting models.

Many different formal modeling paradigms have been applied to molecular biology, each with its own community, formalisms and tools. The present seminar is intended to stimulate closer interaction within the field and to create a common platform for discussion.

The program covered a large fraction of the diversity of formal modeling in molecular biology, including sessions on

- ordinary differential equation models,
- process calculi,
- state machines
- process algebras,
- logics,
- constraints-based modeling.

A major area of interest was the debate over the relative merits of the different approaches to modeling that were presented in the meeting, and the emerging interest in directly executable specifications in terms of the analytical techniques that can be used. In addition to computational modelers, the participants also included a number of high-profile systems biologists who presented important new developments at the experimental side of the life sciences in keynote speeches and provided crucial critical feedback on the validity of the formal modeling concepts. The meeting was particularly friendly and productive, and had a good mix of young and established researchers. Numerous new collaborations were established across the fields and are now followed up in longer-term research projects.

Modeling Competition

A central feature of the seminar was a modeling competition (with a highly collaborative flavor) of various modeling paradigms. This provided a unique opportunity for participants to directly compare their approaches and find common ground. It turned out that Dagstuhl is an ideal place to encourage this kind of productive and challenging interaction: new teams started to form already on the first day and many new analyses or collaborations took place during intense personal interactions and in small groups in front of the computer.

All contributions to the competition were evaluated by a committee of judges, supplemented by a public vote, based on informal presentations during the conference. This turned out to be a challenging task, as many contributions were of excellent quality, including some by teams that had just met for the first time at the seminar.

All votes were statistically evaluated with software based on the algorithm presented in [BAAH04], revealing an extremly good correlation between the total assessment by the committee of judges and the total assessment by the public vote.

Chapter 14

Applications, Multi-Domain Work

14.1 Knowledge Representation for Intelligent Music Processing

Seminar No. **09051** Date **25.01.–30.01.2009**
Organizers: Eleanor Selfridge-Field, Frans Wiering and Geraint A. Wiggins

The ubiquity and importance of music have made it an obvious candidate for applications of new technology throughout history, but most notably since the late 19th Century, when analogue electronics and then digital computers were brought to bear. There was initially an emphasis on the production of audible sound, but as computers became powerful, they were used in the generation of scores, and in recent years digital technology has approached the difficult problem of the understanding of music, both as what is heard and what is imagined.

This seminar aims to promote the computational study of music at levels of abstraction higher than the audio waveform. Doing so will enable automation of the kind of reasoning applied explicitly by music composers, analysts, researchers and performers as consciously-developed skills, and implicitly by informed listeners as high-level cognitive processes.

Many music encoding systems have been created since the 1960s, and large quantities symbolic musical data have been produced across the world, as the output of disparate projects, and represented for storage in ways which are not interoperable. Music knowledge representation research, as opposed to musical data encoding, emerged in the 1970s. Only after several decades of research, consensus on generally appropriate features for music representation was reached, and approaches—for example MEI, MusicXML, and MPEG7 Notation—have been developed which do model music more fully. Only recently, attempts have been made to represent music in ways which conform to the principles of Knowledge Representation, in that their specifications explicitly include inference systems. The inference aspect is fundamentally important: a computer encoding of data is meaningless without a method for interpreting it.

An important area of application is in digital critical editions of music. Whereas paper editions have the drawback of presenting a selective and static image of a composition,

digital editions potentially provide a more complete representation of the source materials and allow different 'views' of these to be generated automatically. Suitable knowledge representations for these sources would allow inference of missing information that is considered essential for modern study and performance, such as accidental pitch changes in Renaissance music, voice leading in lute tablatures, realisation of implied chords in basso continuo accompaniment, and also suggest solutions for unclear, illegible, corrupted and lost passages. Finally they would allow the compositions to be processed by means of a wide range of music-analytical or music retrieval methods.

14.2 Model-Based Design of Trustworthy Health Information Systems

Seminar No. **09073** Date **11.02.–14.02.2009**
Organizers: Ruth Breu, John C. Mitchell, Janos Sztipanovits Alfred Winter

The Dagstuhl Seminar "Model-Based Design of Trustworthy Information Systems" took place from February 11th to February 14th, 2009, at the International Conference and Research Center Schloss Dagstuhl. The goal of the seminar was to bring together experts from the domains of health care, software engineering and security in order to discuss the challenges of emerging health care scenarios. The seminar combined presentations with discussions in groups.

New technologies for Health Information Systems (HIS) offer a revolutionary new way for the interaction between medical patients and Healthcare providers. Although healthcare like other information-intensive industries has developed and deployed standards-based, secure information infrastructures it is still dependent upon paper records and fragmented, error-prone approaches to service delivery. Thus healthcare has been characterized as a "trillion dollar cottage industry". One of the main concerns is security and privacy that needs to be organically integrated into HIS architectures. Widely cited reports of the U.S. Institute of Medicine and National Research Council have documented weaknesses in information security related to healthcare, the costs and impact of medical errors (a substantial proportion of which involve a component of information mismanagement), lack of a systems approach to complex, team-oriented interdisciplinary care, and the unrealized potential of using the Internet to improve the quality and availability of healthcare services.

How can Health Information Systems help?

Complementing the recognition of the weaknesses are three major drivers that push the healthcare industry towards radical change: (1) the dramatic increase of genetic information and the opening opportunity to provide personalized healthcare, (2) the economic pressures to move healthcare from institutions toward homes, and (3) the rapidly increasing use of Internet and information appliances in society. This fundamental change will be enabled by advanced information technology, including ubiquitous communication and sensing, extensive use of web portals as a central point of access for communication and documentation of health care efficiency. Quality of patient specificity will be achieved

via extensive use of clinical decision support systems combined with automated event monitors.

What are the key challenges?

HIS shall support patients and also doctors, nurses, paramedicals and other health care providers in diagnosing, treating and supporting patients. Health care is not only a health but also a life and death issue. In this existential situation patients have to trust on caregivers and both patients and caregivers depend on the trustworthiness of the information systems used. Not only the highly delicate relation between caregivers and patients but also the data related to this situation need particular protection from misuse. But unfortunately privacy and security requirements are frequently expressed in vague, contradictory and complex laws and regulations; it is a major concern that requires new approaches in systems design. Trustworthy HIS need to provide effective, high quality support for providing the best care for patients but without compromising their privacy, security and safety.

How to solve these challenges?

End-to-end architecture modeling integrated with privacy and security models offer new opportunities for system designers and end users. Model-based approaches to HIS are investigated extensively in Europe and in the US. While initial results show promise, many fundamental problems remained unsolved, such as modeling of privacy and security policies, and verification of their consistency, and compliance to requirements. HIS requires new architectures that are sufficiently flexible to support personalized health care without causing harm and can be adapted to changing policies.

Goals and Expected Results

The goal of this seminar was to help the computer science community understanding the unique challenges of this field and offer insight for HIS developers in the state of the art in model-based design technologies. The objective was to understand the challenges and promising approaches in HIS design as the intersection of five major areas: health information systems, model-based software and systems design, reliability, security and privacy science, enterprise information systems and legal policy. The seminar combined presentations with discussions in groups and in the plenary.

14.3 Algorithms and Number Theory

Seminar No. **09221** Date **24.05.–29.05.2009**
Organizers: Johannes A. Buchmann, John Cremona, Michael E. Pohst

This seminar on number-theoretical algorithms and their applications was the sixth on this topic at Dagstuhl over a period of seventeen years. This time 39 people from 10 countries participated.

One of the major goals of these seminars has been to broaden interactions between number theory and other areas. For instance, there has been an effort to bring together people

developing the theory of efficient algorithms with people actually writing software. There has also been continuing interest in cryptography. These aspects were both emphasized by the topics of special interest in this year: Number Theoretical Software and Algorithms for the Post Quantum Era.

About half of the 24 talks given were in these areas showing rapidly growing interest. One fourth of the talks were on curves, most with an eye to applications in cryptography.

The other talks focused on more classical topics of algorithmic algebraic number theory. We just mention the calculation of global fields and of class groups.

Even though we had less participants than at the last meeting the group seemed to be more homogenous. The variety of topics of the talks was stimulating to the audience. Their smaller number gave more room for discussions. It did not come as a surprise that these were most intensive in our emphasized topics.

For example, number theoretical software was not only discussed but also developed during the meeting. The participants did indeed a lot of coding. We would like to mention hat M. Stoll has a C program called ratpoints which is very fast at finding rational solutions to $y^2 = f(x)$. During the conference this program was incorporated into Sage, a process which included finding several bugs (memory leaks) in ratpoints so that M. Stoll could fix them right there.

The reaction of the participants was very positive and we believe that we succeeded in having an effective meeting that was able to appeal to a broad audience. We made sure to allow for adequate breaks between sessions, and - as already metioned - there were many opportunities for discussions that the participants took advantage of. The pleasant atmosphere of Schloss Dagstuhl once again contributed to a very productive meeting. Even more positively, several younger people who were there (for the first time) told us that they not only found it a very good meeting indeed, but the best venue they had been to for a conference.

14.4 Computational Creativity: An Interdisciplinary Approach

Seminar No. **09291** Date **12.07.–17.07.2009**
Organizers: Margaret Boden, Mark d'Inverno, John McCormack

Artistic creativity remains a mysterious, enigmatic subject — a "grand challenge" for Computer Science. While computers have exceeded the capabilities of humans in a number of limited domains (e.g. chess playing, music classification, theorem proofs, induction), human creativity generally remains unchallenged by machines and is considered a fundamental factor in our intellectual success. There is a sense that artistic creativity is somehow "special" in a way that could not be captured in an algorithm, hence implemented on a machine. This seminar aims to show that creativity is indeed special, but that it can be an emergent property of mechanical processes.

The seminar will address problems in computational creative discovery where computer processes assist in enhancing human creativity or may autonomously exhibit creative behavior independently. The intention is to develop ways of working with computation that achieve creative possibilities unattainable from any existing software systems. These goals will be developed in the context of artistic creation (visual art and music composition), however the results may be applicable to many forms of creative discovery.

The specific seminar aims are to:

- Contribute to fundamental research on our understanding of artistic creativity in humans and machines;
- Develop new methodologies for creative design in digital media, with particular emphasis on evolutionary ecosystem dynamics, where new algorithms for creative discovery are inspired by biological processes;
- Bring together researchers from a variety of disciplines and backgrounds, with coverage across the arts and sciences, but with a common goal of furthering our understanding of how computers may generate creative behaviour, using an interdisciplinary approach.

Creativity is a vast and complex topic, investigated by many disciplines. In broad terms it involves the generation of something novel and appropriate (i.e. unexpected, valuable). In this seminar we focus on artificially creative systems, either simulated in software or software process working in synergetic tandem with a human artist. The necessary conditions for any artificial creative system must be the ability to interact with its environment, learn, and self-organise, and this is the basis of the seminar's approach. Darwinian evolution has been described as the only theory with the explanatory power for the design and function of living systems, accounting for the amazing diversity and astonishing complexity of life. Evolutionary synthesis is a process capable of generating unprecedented novelty, i.e. it is creative. It has been able to create things like prokaryotes, eukaryotes, higher multicellularity and language through a non-teleological process of replication and selection. We would like to investigate, on a metaphoric level, the mechanisms of biological evolution in order to develop new approaches to computational creativity.

Questions Addressed by the Seminar:

- How can we further our formal understanding of the artistic creative process in a variety of disciplines, including visual art and music? By inviting a group of leading creative practitioners in visual art and music, we hope to gain insight on how computational systems can be creative from the perspective of creative artists.
- How can evolutionary algorithms be extended to encompass heterogeneous environments and more complex process cycles? If evolution is a process for creative discovery, how can this process be adapted to human-creative domains (as opposed to biological ones)?

- What are the appropriate mappings and metaphors if we are to use biological process and systems as a basis for developing creative systems with computers?
- What is the best approach in developing autonomous creative systems (i.e. machines that exhibit independently creative behaviour).
- What are the appropriate methods and measures to objectively verify and validate creative behavior in artificial systems.
- Which is the better approach to understanding creativity in discrete devices: combinatorial emergence (the understanding that creativity is the creative combination or recombination of previously existing elements) or creative emergence (creativity begins with knowledge, skill and abilities, and emerges from these faculties through interaction with the environment; new primitives emerge in the underlying system, leading to a transformation of the conceptual space).

Objectives and results expected to be produced by the seminar:

We intend to invite a gathering of many of the world's leading researchers in this area with a view to making a significant contribution to state-of-the-art knowledge in this area. We have selected researchers with expertise in artificial intelligence, agent-based modelling, evolutionary algorithms, fine art, music composition, cognitive science and philosophy. We will request each participant to develop a position paper that addresses one or more of the research questions stipulated above. Following on from the seminar we would expect to edit a new scholarly book in the area of computational creativity that explores the topics discussed at the Dagstuhl seminar in greater depth.

14.5 Democracy in a Network Society *(Dagstuhl Perspectives Workshop)*

Seminar No. **09402** Date **27.09.–02.10.2009**
Organizers: David Chaum, William H. Dutton,, Miroslaw Kutylowski, Tracy Westen

The workshop was a meeting forum for experts in the area of computer security and social sciences. The main idea of the seminar was to discuss new challenges for democracy during the transition from traditional society into a society where network communication influences so much social and political life.

The workshop participants discussed the key issues behind success or failure of electronic systems in e-democracy. While advances of technology play a central role in evolution of e-democracy, the main threats and failures are due to insufficient collaboration and lack of understanding among technologists, social scientists, public officials and other stakeholders. In the past, major failures can be attributed to a narrow view of the systems supporting e-democracy. For this reason many fundamental mistakes have been made.

Some major problems arise when technical sciences and social sciences meet. On the one hand, computer specialists are often unaware of real requirements for the emerging systems, on the other hand the specialists from social sciences might be unaware of technical limitations due to hermetic language of computer security professionals. Nevertheless, the workshop participants succeeded immediately in building up a working group focused on identifying the most crucial issues for development of future e-democracy systems.

The result of the workshop is a set of recommendations for decision-makers regarding e-democracy systems. The list does not consider all problems that may arise, but brings focus to those that in our opinion have the biggest impact.

Recommendations

1. *Encourage Interdisciplinary Collaboration.* Severe design errors may result from making decisions based on partial expertise, or from separate groups working in isolation. As design processes for technologies used in democratic systems should include a wide range of competencies, it is vital that lawyers, public officials and social scientists are engaged as well as computer scientists and engineers. Unfortunately, the workshop participants observe that this is not a common practice today and many fundamental errors in the past resulted from partial expertise.

2. *Ensure Effective Take-up of E-democracy Solutions.* At present, government-driven processes (like elections, disclosure of information) are often so conservative that they fail to take full advantage of new technologies and approaches, despite that they have proved effective elsewhere. The reason for this phenomenon is a discrepancy between available solutions that are ready to use and specific requirements of e-democracy. Substantial amount of research is necessary to adapt emerging technologies to meet the diverse requirements of e-democracy.

3. *Deploy Appropriate Design Models.* The lesson we have learnt during the last decades is that the really successful systems are in practice the flexible ones that were not designed by a single organization but have instead developed through collaborative efforts of many participants driven by their interests and needs. Therefore we feel that new technical systems supporting e-democracy should be small, flexible, modular and based on proven off-the-shelf technical components, rather than be large, centralized special-purpose systems.

4. *Promote Best Practice.* There are examples of excellent solutions which are implemented and used in practice. However, dissemination of such best practices is limited. A survey should be conducted of best practices. This is particularly important for making government information accessible online inexpensively, efficiently and in forms that are easy to use by the public. Today, inefficient access to information is one of the major weaknesses of democracies, despite many efforts. Pilot projects should then be funded to implement these best practices in a number of different jurisdictions. Information on best practices and pilot projects should be made available to the public in easily accessible formats.

5. *Support Open-Audit Systems.* Research on electronic voting systems has shown that our approach to security assurance should be redefined. Traditional certification by trusted bodies should be continued, however in order to provide undeniable evidence open-audit concepts should be developed. In particular, current field trials of open-audit voting systems should be carefully assessed and documented. When they are successful, larger-scale trials should be encouraged.

6. *Learn from Web 2.0 Innovations.* Public officials and system designers should draw on the experience of Web-based social networks. There are substantial technical and social challenges related to Web 2.0, but there are opportunities as well. This should be taken into account when planning online systems for democratic decision making.

7. *Address Conflicting Requirements.* Quite often, requirements for e-democracy systems are in conflict. A prominent example are e-voting systems, which have to provide strong privacy of vote casting and voters' identification at the same time. Since according to the present state-of-the-art the answers for many fundamental questions are still missing, more research should be directed towards new technologies that have the potential to reconcile between such conflicting requirements. This concerns in particular privacy enhancing technologies, identity management and cryptographic protocols.

8. *Gain Public Acceptance.* One of preconditions for introducing technical systems supporting democratic processes is gaining understanding, acceptance and confidence by the lay, non-scientific public. A failure to do so would immediately undermine the citizens' will to engage in the process. Therefore technical solutions for e-democracy that support democratic processes should be made simple enough, or must be so widely endorsed by the scientific community and other trusted societal leaders. Democratic technologies should be designed with widespread public acceptance as a key design parameter.

9. *Fund Civic Engagement Experiments.* Since in the field of e-democracy we are entering unknown grounds, a lot can be learned from examples. For this reason, governments should be encouraged to fund experimentations with technologies that support greater online civic engagement in democratic processes (voting, information acquisition, collaborative participation in government decisions). On the one hand, such government funding will encourage technological research as well as provide computer scientists with the priorities they require. On the other hand, these experimentations will allow the citizens to influence design evolution so that it goes in the right direction.

10. *Share Knowledge Between Disciplines.* Lack of interaction and sometimes even barriers for interdisciplinary work is one of the main risk factors for development of e-systems supporting democracy. Therefore, various contributions made by different disciplines to e-democracy development can be strengthened through forums that encourage (not only verbally) dialogue between multidisciplinary groups of computer and social scientists, legal scholars, practitioners and policy experts.

For more about results of the seminar see the article in Social Sciences Research Network Machiavelli Confronts 21st Century Digital Technology: "Democracy in a Network Society" (`http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1521222`) published by the workshop participants.

Chapter 15

Other Work

15.1 Preventing the Brainware Crisis *(Dagstuhl Perspectives Workshop)*

Seminar No. **09142** Date **31.03.–03.04.2009**
Organizers: Stephan Diehl, Mike Fellows, Werner Hartmann, Ulrike Stege

Computer science (CS) and engineering research develops rapidly, and their successes influence economy and our daily life tremendously. CS has become an important part of other sciences, social sciences, arts, and engineering; a big part of research activities and jobs in CS are highly interdisciplinary.

From the 1960s until 2000, engineering and CS have been popular. No extra mile had to be taken to attract students to universities. While nowadays our youth is using new technologies fluently, the number of CS students in first-year university has declined alarmingly in North-America and Europe. Although according to some German statistics the number of graduating students increased from 2005 to 2006, enrollments have dropped 49% from their height in 2001/02. The number of female undergraduate students in CS is low. The alarming trend of declining enrollment exists despite a desperate need for computer scientists in industry and a popular debate on the topic in the media. Although the IT industry is booming and the job opportunities for its graduates are excellent, the public perception seems to reflect a contrary attitude.

The participants of the perspectives workshop included researchers from academia and industry, teachers, science journalists as well as employment officers. This allowed them to discuss from many different aspects the problem of how to prevent the brainware crisis, i.e. to stop the increasing lack of computer science students which will result in a massive shortage of computer science researchers and IT professionals in the long run.

In their manifesto the participants elaborate on their three main recommendations which can be briefly summarized as: make computer science programs more attractive to women, make curricula more engaging and interdisciplinary, and make the public more aware of the results and impact of computer science research.