



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report

2013



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2013

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Deutschland
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. Dr. h.c. Stefan Jähnichen
Geschäftsführung	Prof. Dr. Dr. h.c. Dr. h.c. Reinhard Wilhelm (Wissenschaftlicher Direktor) Dr. Christian Lindig (Technisch-administrativer Geschäftsführer)
Gesellschafter	Gesellschaft für Informatik e.V., Deutschland Universität des Saarlandes, Deutschland Technische Universität Kaiserslautern, Deutschland Karlsruher Institut für Technologie (KIT), Deutschland Technische Universität Darmstadt, Deutschland Universität Stuttgart, Deutschland Universität Trier, Deutschland Johann Wolfgang Goethe-Universität Frankfurt am Main, Deutschland Institut National de Recherche en Informatique et en Automatique (INRIA), Frankreich Centrum voor Wiskunde en Informatica (CWI), Niederlande Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V., Deutschland
Einbindung	Mitglied der Leibniz-Gemeinschaft Mitglied im Leibniz-Forschungsverbund Science 2.0
Kontakt	Geschäftsstelle Schloss Dagstuhl Universität des Saarlandes Campus E1 1 66123 Saarbrücken Germany Telefon: +49 681 302–4396 Fax: +49 681 302–4397 E-Mail: service@dagstuhl.de http://www.dagstuhl.de
Redaktion	Dr. Marcel R. Ackermann, Laura Cunniff, Dr. Andreas Dolzmann
Mitarbeit	Dr. Marcel R. Ackermann, Dr. Roswitha Bardohl, Heike Clemens, Laura Cunniff, Dr. Andreas Dolzmann, Jutka Gasiorowski, Dr. Marc Herbstritt, Oliver Hoffmann, Antónia Košťová, Dr. Christian Lindig, Petra Meyer, Thomas Schillo, Christina Schwarz, Dr. Michael Wagner, Prof. Reinhard Wilhelm und die Autoren der Zusammenfassungen in Kapitel 4
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
Druck	Digitaldruck Pirrot GmbH, Saarbrücken-Dudweiler, Deutschland
Druckauflage	180
ISSN	2199-1995
Online-Version	http://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	April 2013
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	© Creative Commons BY-ND 3.0 DE

Vorwort

Foreword

Dies wird mein letztes Vorwort zu einem Dagstuhl-Tätigkeitsbericht sein, weil ich Ende April 2014 als wissenschaftlicher Direktor des Leibniz-Zentrums für Informatik in Schloss Dagstuhl ausscheide. Das soll mir erlauben, ein wenig über die 24 Jahre seit der Gründung zu philosophieren.

Als das „Internationale Begegnungs- und Forschungszentrum für Informatik“, so wurde es damals genannt, 1990 gegründet wurde, schien es schon aus der Zeit gefallen zu sein. Wer braucht ein Zentrum, irgendwo weit ab vom Schuss und schwer erreichbar, um dort Wissenschaftler für eine Woche einzusperrern, damit sie miteinander arbeiten, wenn zur gleichen Zeit die Informations- und Kommunikationstechnologien die direkte, unvermittelte Kommunikation überflüssig machen? Weshalb müssen Wissenschaftler noch die Kosten und Strapazen von Reisen auf sich nehmen, um sich zu treffen, wenn es eine Videokonferenz auch tut? Es war die feste Überzeugung, dass technisch vermittelte Kommunikation die direkte Kommunikation nicht überflüssig macht und dass ein klösterliches Beisammensein der gemeinsamen Forschungsarbeit förderlich ist, welche die Gesellschaft für Informatik damals dazu gebracht hat, das vom Mathematischen Forschungsinstitut in Oberwolfach vorgelebte Konzept zu übernehmen.

40 Wissenschaftler, alt und jung, gestanden und frisch, aus aller Herren Länder, sollten für eine Woche zusammen leben und an einem gemeinsamen Thema arbeiten. Es wurden einige Mechanismen eingeführt, welche für die optimale Kommunikation sorgen sollten: Die Gäste wurden bei den Mahlzeiten (mehr oder weniger) zufällig an die Tische gruppiert, um Cliques aufzulösen. Die Wanderung am Mittwochnachmittag ermöglichte die zwanglose Kommunikation. Es gab keinen Fernseher auf dem Zimmer, um die Gäste auch nach Ende der Sitzungen in die Gemeinschaftsräume zu locken. Es gab Workstationräume, in die man sich begeben musste, wenn man an das Internet wollte – damals brachte noch niemand einen Laptop mit und Handys wogen 4 kg und waren für Wissenschaftler unerschwinglich. Es gab den Weinkeller, in dem der Geist, und nicht nur der Weingeist, noch freier schweifen konnte

This will be my last foreword for a Schloss Dagstuhl Annual Report, since my tenure as Scientific Director of the Leibniz-Center for Informatics at Schloss Dagstuhl will come to an end on April 31, 2014. A little philosophizing about the 24 years since the center's founding is therefore perhaps in order.

When the “Internationales Begegnungs- und Forschungszentrum für Informatik,” (“International Meeting and Conference Center for Computer Sciences”) as it was then called, was created in 1990, it already seemed to have fallen out of time. Who needed a center in some remote corner far off the beaten track, in which to imprison scientists for a week so that they could work together, when information and communication technologies had succeeded in making such direct, unmediated communication superfluous? Why should scientists still bear the costs and hardships of traveling in order to meet each other, when a video conference could do the same? It was the firm conviction that technically mediated communication does not make direct communication superfluous, and that a monastic gathering of joint research promotes collective research, which motivated the Gesellschaft für Informatik to launch at Schloss Dagstuhl a new center for the informatics community, modeled on the concept of the the Mathematical Research Institute in Oberwolfach.

Forty scientists, young and old, seasoned and fresh, from all over the world, should live together for a week and work together on a common theme. Some mechanisms were put into place to ensure optimal communication: guests were (more or less) randomly seated at meals in order to break up cliques; the traditional Wednesday afternoon hike allowed for informal communication; sleeping quarters without television sets enticed people to socialize with each other between sessions in the common areas; work stations were set up for those who wanted to access Internet – at that time, no one brought a laptop and mobile phones weighed 4 kilos and anyway were unaffordable for scientists. And there was the wine cellar, where the spirit – not just the alcoholic kind – could roam more freely than in the lecture halls. And look! The principle worked! With

als in den Vortragssälen. Und sieh an! Das Prinzip funktionierte! Es wurde ziemlich schnell klar, dass das Konzept Wissenschaftskloster sich nicht überlebt hatte.

Danach wurde das Konzept harten Herausforderungen unterworfen: Handy wurden populär, und die anfänglich schlechte Empfangsqualität in der Umgebung des Schlosses wurde besser; WLAN musste eingerichtet werden und brachte die Gäste in schwere Versuchung, auch wenn meine ironische Frage auf dem Begrüßungsposter, „Why are you here? To read your email or to communicate with your peers?“ meist verstanden wurde. Es gilt nach wie vor die Beobachtung, dass die Qualität der gemeinsamen Arbeit mit der Anzahl der offenen Laptops negativ korreliert.

In diesem Kernbereich unserer Tätigkeit, der Organisation von Dagstuhl-Seminaren und Perspektiven-Workshops, ist das Leibniz-Zentrum für Informatik immer erfolgreicher und populärer geworden. Man sieht das an der zunehmenden Auslastung des Zentrums und an der großen Zahl an Anträgen auf solche Veranstaltungen.

Neben diesem Kerngeschäft beteiligt sich das Leibniz-Zentrum für Informatik am Trend zum open-access-Publishing. Die Öffentlichkeit als Geldgeber fordert den kostenfreien Zugang zu Forschungsergebnissen, die traditionell in Zeitschriften veröffentlicht werden. So entsteht eine neue Dynamik in der Beziehung zwischen Wissenschaftlern, kommerziellen und nichtkommerziellen Verlagen und Bibliotheken.

Der Konsument, der Wissenschaftler, braucht eine Unterstützung zum effektiven Zugriff auf das riesige Volumen an wissenschaftlichen Publikationen. Dies leistet unter anderem die bibliographische Datenbank dblp, die über 20 Jahre von Michael Ley an der Universität Trier entwickelt wurde. Seit mehreren Jahren kooperiert das Leibniz-Zentrum für Informatik mit der Universität Trier in der Weiterentwicklung dieser weltweit gerühmten Datenbank. Diese Kooperation hat sich in einer rasant gewachsenen Abdeckung der Informatikliteratur, aber auch in einer solideren politischen und strategischen Fundierung der Arbeit bemerkbar gemacht. Diese Arbeit wird weiter geführt, derzeit mit dem Ziel die Autorenidentifikation zu verbessern.

so much communication, it quickly became clear that the concept of a science monastery was not outdated.

After that the Dagstuhl concept was subjected to some serious challenges: mobile phones became popular and the initially poor reception quality in the vicinity of the castle improved; WLAN had to be set up and became a severe temptation to visitors, even if my ironic question on the welcome poster, “Why are you here? To read your email or to communicate with your peers?” was usually understood. It was generally observed that the quality of the collective work negatively correlated to the number of open laptops.

Despite such setbacks, in this key area of our activity – organizing Dagstuhl Perspectives Workshops and Dagstuhl Seminars – the Leibniz-Zentrum für Informatik has become increasingly successful and popular. You can see this in the increasing utilization of the center and in the growing number of seminar and workshop proposals.

In addition to this core business, the Leibniz-Zentrum für Informatik has gotten involved in a trend towards open access publishing. The public as a donor demands free access to research results that are traditionally published in journals. This creates a new dynamic in the relationship between scientists, commercial and non-commercial publishers and libraries. The consumer, the scientist, needs support in order to gain effective access to the huge volume of existing scientific publications and achieves that, inter alia, via the bibliographic database dblp, which was developed over 20 years by Michael Ley at the University of Trier. In recent years, the Leibniz-Zentrum für Informatik has cooperated with the University of Trier to further the development of this world-acclaimed database. This collaborative initiative has manifested itself in the rapidly growing coverage of computer science literature, and in a more solid political and strategic foundation of the dblp project. The work will be continued in the future with the aim of improving author identification.

Reinhard Wilhelm

Im Namen der Geschäftsführung

On Behalf of the Managing Directors

Prof. Dr. Dr. h.c. Dr. h.c. Reinhard Wilhelm
Wissenschaftlicher Direktor

Dr. Christian Lindig
Technisch-administrativer Geschäftsführer

Inhaltsverzeichnis

Contents

Vorwort		
<i>Foreword</i>		iii
1	Das Zentrum Schloss Dagstuhl	
	<i>Schloss Dagstuhl Center</i>	1
1.1	Wissenschaftlicher Austausch in anregender Umgebung	
	<i>Stimulating Exchanges in Relaxed Surroundings</i>	2
1.2	Neuigkeiten in 2013	
	<i>News from 2013</i>	4
2	Das wissenschaftliche Programm 2013	
	<i>Scientific Program 2013</i>	13
2.1	Dagstuhl-Seminare	
	<i>Dagstuhl Seminars</i>	14
2.2	Dagstuhl-Perspektiven-Workshops	
	<i>Dagstuhl Perspectives Workshops</i>	15
2.3	Einreichung der Anträge und Begutachtungsverfahren	
	<i>Proposal Submission and Review Process</i>	16
2.4	Nähere Betrachtung des Dagstuhl-Seminarprogramms 2013	
	<i>A Closer Look at the Dagstuhl Seminar Program in 2013</i>	16
2.5	Angaben zu Teilnehmern und Organisatoren	
	<i>Participant and Organizer Data</i>	17
2.6	Themen und Forschungsgebiete	
	<i>Topics and Research Areas</i>	19
2.7	Weitere Veranstaltungstypen	
	<i>Further Event Types</i>	21
2.8	Qualitätssicherung	
	<i>Quality Assurance</i>	22
2.9	Auslastung des Zentrums	
	<i>Utilization of the Center</i>	22
3	Resonanz	
	<i>Feedback</i>	25
3.1	Resonanz von Seminarteilnehmern	
	<i>Feedback from Seminar Participants</i>	26
3.2	Resonanz unserer Seminarorganisatoren	
	<i>Feedback from Seminar Organizers</i>	27
3.3	Öffentliche Resonanz im Web	
	<i>Feedback in Social Media</i>	29
3.4	Resonanz im Fragebogen	
	<i>Seminar Survey Feedback</i>	29
3.5	Resonanz zu Dagstuhl Publishing	
	<i>Feedback on Dagstuhl Publishing</i>	33
3.6	Resonanz zur Bibliographiedatenbank dblp	
	<i>Feedback on the dblp Computer Science Bibliography</i>	34

4	Die Seminare in 2013	
	<i>The 2013 Seminars</i>	35
4.1	Symbolic Methods in Testing	37
4.2	Engineering Resilient Systems: Models, Methods and Tools	38
4.3	Computational Counting	40
4.4	Civilian Crisis Response Models	42
4.5	Epidemic Algorithms and Processes: From Theory to Applications	44
4.6	Software Certification: Methods and Tools	46
4.7	Multicore Enablement for Embedded and Cyber Physical Systems	49
4.8	Fault Prediction, Localization, and Repair	50
4.9	Decentralized Systems for Privacy Preservation	51
4.10	Dependence Logic: Theory and Applications	52
4.11	Mechanisms of Ongoing Development in Cognitive Robotics	54
4.12	Consistency in Distributed Systems	56
4.13	Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices	58
4.14	Analysis, Test and Verification in The Presence of Variability	60
4.15	Computational Geometry	62
4.16	Scheduling	64
4.17	Bidimensional Structures: Algorithms, Combinatorics and Logic	65
4.18	Future Internet	66
4.19	Formal Verification of Distributed Algorithms	68
4.20	Correct and Efficient Accelerator Programming	69
4.21	Drawing Graphs and Maps with Curves	70
4.22	Interface of Computation, Game Theory, and Economics	73
4.23	Pointer Analysis	74
4.24	Customizing Service Platforms	76
4.25	VaToMAS – Verification and Testing of Multi-Agent Systems	78
4.26	Meta-Modeling Model-Based Engineering Tools	80
4.27	Tree Transducers and Formal Methods	82
4.28	Information Visualization – Towards Multivariate Network Visualization	84
4.29	Automated Reasoning on Conceptual Schemas	86
4.30	Computational Methods Aiding Early-Stage Drug Design	87
4.31	Belief Change and Argumentation in Multi-Agent Scenarios	88
4.32	Indexes and Computation over Compressed Structured Data	89
4.33	Virtual Realities	90
4.34	Parallel Data Analysis	91
4.35	Interoperation in Complex Information Ecosystems	92
4.36	Theory of Evolutionary Algorithms	94
4.37	Computer Science in High Performance Sport – Applications and Implications for Professional Coaching	96
4.38	Duality in Computer Science	97
4.39	“My Life, Shared” – Trust and Privacy in the Age of Ubiquitous Experience Sharing	98
4.40	Reinforcement Learning	100
4.41	The Critical Internet Infrastructure	102
4.42	Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time	104
4.43	Verifiably Secure Process-Aware Information Systems	106
4.44	ICT Strategies for Bridging Biology and Precision Medicine	108
4.45	Coding Theory	110
4.46	Interaction with Information for Visual Reasoning	111
4.47	Crowdsourcing: From Theory to Practice and Long-Term Perspectives	112
4.48	Cloud-based Software Crowdsourcing	114
4.49	Quantum Cryptanalysis	116
4.50	Integration of Tools for Rigorous Software Construction and Analysis	118
4.51	Algorithms and Scheduling Techniques for Exascale Systems	120
4.52	Collaboration and learning through live coding	122
4.53	Algorithm Engineering	124
4.54	Inter-Vehicular Communication – Quo Vadis	126
4.55	Automatic Application Tuning for HPC Architectures	128
4.56	Physical-Cyber-Social Computing	131

4.57	Deduction and Arithmetic	132
4.58	Genomic Privacy	134
4.59	Algorithms for Optimization Problems in Planar Graphs	136
4.60	Nominal Computation Theory	138
4.61	Real-World Visual Computing	140
4.62	Evaluation Methodologies in Information Retrieval	142
4.63	Computational Audio Analysis	144
4.64	Proxemics in Human-Computer Interaction	148
4.65	Electronic Markets and Auctions	151
4.66	Computational Models of Language Meaning in Context	152
4.67	Synchronous Programming	154
4.68	Global Measurement Framework	156
4.69	Unleashing Operational Process Mining	158
4.70	Forensic Computing	160
4.71	Computational Mass Spectrometry	162
4.72	Geosensor Networks: Bridging Algorithms and Applications	166
4.73	Approaches and Applications of Inductive Programming	168
4.74	Software Engineering for Self-Adaptive Systems: Assurances	170
4.75	Social Issues in Computational Transportation Science	171
5	Öffentlichkeitsarbeit	
	<i>Public Relations and Outreach</i>	173
5.1	Pressemitteilungen und Medienarbeit <i>Press Releases and Media Work</i>	174
5.2	Fortbildung <i>Educational Training</i>	175
6	Dagstuhl Publishing	
	<i>Dagstuhl Publishing</i>	177
6.1	Portfolio <i>Portfolio</i>	178
6.2	Infrastruktur <i>Infrastructure</i>	180
7	Bibliographiedatenbank dblp	
	<i>dblp computer science bibliography</i>	185
7.1	Schloss Dagstuhl und dblp <i>Schloss Dagstuhl and dblp</i>	186
7.2	Fortentwicklung der Datenbasis <i>Development of the Data Stock</i>	187
7.3	Technische Weiterentwicklung <i>Technical Improvements</i>	187
7.4	Erhebung zu wissenschaftlichen Publikationsorganen in der deutschen Informatik <i>Survey on Scientific Publication Venues in German Computer Science</i>	189
7.5	Ausblick <i>Outlook</i>	191
8	Einrichtung und Service	
	<i>Facilities and Services</i>	193
8.1	Tagungsräume <i>Conference Facilities</i>	194
8.2	Freizeit und Ambiente <i>Leisure Facilities</i>	194
8.3	Dagstuhls Küche <i>Dagstuhl's Kitchen</i>	194
8.4	Kinderbetreuung <i>Childcare</i>	198
8.5	Computer und Vernetzung <i>Computers and Networks</i>	198

8.6	Dagstuhl's Web-basierte Dienste <i>Dagstuhl's Web-based Services</i>	198
9	Bibliothek	
	Research Library	201
9.1	Bestand und Angebot <i>Inventory and Offering</i>	202
9.2	Spenden an die Bibliothek <i>Library Donations</i>	202
10	Kunst	
	Art	205
10.1	Dagstuhl als Galerie <i>Dagstuhl as Art Gallery</i>	206
10.2	Kunstankauf durch Spenden <i>Art Sponsorship and Donations</i>	208
11	Stiftung „Informatikzentrum Schloss Dagstuhl“	
	The Dagstuhl Foundation	211
11.1	Zielsetzung <i>Aims</i>	212
12	Organe und Gremien	
	Organs and Bodies	213
12.1	Struktur der Gesellschaft <i>Structure of the Center</i>	214
12.2	Organe und Gremien der Gesellschaft <i>Organs and Bodies of the Organization</i>	214
13	Statistik 2013	
	Statistics 2013	221
14	Veranstaltungen 2013	
	Schedule of Events 2013	227
14.1	Dagstuhl-Seminare <i>Dagstuhl Seminars</i>	228
14.2	Dagstuhl-Perspektiven-Workshops <i>Dagstuhl Perspectives Workshops</i>	233
14.3	GI-Dagstuhl-Seminare <i>GI-Dagstuhl Seminars</i>	233
14.4	Lehrveranstaltungen <i>Educational Events</i>	234
14.5	Sonstige Veranstaltungen <i>Other Events</i>	234

1 **Das Zentrum Schloss Dagstuhl** *Schloss Dagstuhl Center*

Wissenschaftlicher Austausch in anregender Umgebung

1.1

Stimulating Exchanges in Relaxed Surroundings

Schloss Dagstuhl, Leibniz-Zentrum für Informatik, hat zum Ziel, Informatikforschung von international anerkanntem Rang zu fördern, Weiterbildung auf hohem fachlichen Niveau durchzuführen und den Wissensaustausch zwischen Forschung und Praxis anzuregen. Das Zentrum veranstaltet hierzu Forschungsseminare, bei denen führende Wissenschaftler aus der ganzen Welt für eine Woche zu einem intensiven Wissensaustausch zusammengeführt werden. Die Seminare ermöglichen die Vorstellung neuer Ideen, die Diskussion aktueller Probleme sowie die Weichenstellung für zukünftige Entwicklungen.

Die Idee zur Gründung von Schloss Dagstuhl wurde Ende der 1980er Jahre geboren, zu einem Zeitpunkt, an dem die Informatikforschung – ursprünglich der Mathematik und den Ingenieurwissenschaften entsprungen – enormen Aufwind erfuhr. Die Gesellschaft für Informatik beobachtete damals die zunehmende Nachfrage von Informatikwissenschaftlern am weltbekannten Mathematischen Forschungsinstitut Oberwolfach und sah die Notwendigkeit, ein eigens auf die Informatik ausgerichtetes Zentrum einzurichten. Schloss Dagstuhl wurde schließlich 1990 gegründet und entwickelte sich rasch zu einem weltweit renommierten Treffpunkt in der Informatikforschung.

Das wissenschaftliche Programm von Schloss Dagstuhl umfasst die sogenannten Dagstuhl-Seminare sowie Dagstuhl-Perspektiven-Workshops. Dem verheißungsvollen Forschungsnachwuchs wird dabei die Möglichkeit gegeben, in seinem speziellen Fachgebiet mit exzellenten Experten zusammenzuarbeiten und neue Sichtweisen zu diskutieren. Das Programm eines Dagstuhl-Seminars wird absichtlich flexibel gestaltet, um eine gemeinschaftliche Atmosphäre zu schaffen, die in dynamischer Weise offene und kreative Diskussionen zulässt. Bei einem Dagstuhl-Perspektiven-Workshop hingegen diskutiert eine oftmals kleinere Gruppe von ausgewiesenen Experten ein Themengebiet und seine perspektivische Ausrichtung. Hierzu wird die aktuelle Situation eines Forschungsgebietes analysiert, um darauf aufbauend strategische Empfehlungen und richtungweisende Perspektiven für die weitere Zukunft zu entwickeln. Die Erkenntnisse werden in einem Manifest zusammengefasst, das auch an (politische) Entscheidungsträger weitergegeben wird.

Die Seminare und Perspektiven-Workshops werden jeweils von einer kleinen Gruppe ausgewiesener Wissenschaftler im entsprechenden Gebiet beantragt. Für die Begutachtung der Vorschläge und der Teilnehmerlisten ist das Wissenschaftliche Direktorium verantwortlich, bevor Anträge akzeptiert und in Dagstuhls wissenschaftliches Programm aufgenommen werden. Die Teilnahme an diesen Veranstaltungen ist nur mit einer persönlichen Einladung durch das Zentrum möglich. Um den besten internationalen Wissenschaftlern eine Teilnahme zu ermöglichen, wird ein Teil der Aufenthaltskosten von Dagstuhl übernommen.

Schloss Dagstuhl befindet sich in einer ländlichen Gegend im nördlichen Saarland, im Herzen des Dreiländerecks Deutschland, Frankreich und Luxemburg. Es

The mission of the Leibniz-Zentrum für Informatik at Schloss Dagstuhl is to promote world-class research in informatics, support cutting-edge continuing education and professional development, and encourage the exchange of knowledge and findings between the academic community and industry. The center hosts research seminars in which leading researchers from all over the world live together at Schloss Dagstuhl for several days in an intensive research climate. New ideas are showcased, topical problems are discussed, and the course is set for future development in the field.

The idea behind Schloss Dagstuhl came about during the late 1980s, when research in computer science grew rapidly worldwide as an offshoot of mathematics and engineering. At that time the German *Gesellschaft für Informatik* (German Informatics Society) became aware of the growing number of computer scientists at the world-famous Mathematics Research Institute in Oberwolfach, Germany, and recognized the need for a meeting venue specific to the informatics community. Dagstuhl was founded in 1990 and quickly became established as one of the world's premier centers for informatics research.

The center's scientific program includes the well-known Dagstuhl Seminars series and its complement, the Dagstuhl Perspectives Workshops series. Dagstuhl Seminars offer promising young researchers in a specific cutting-edge field of informatics the opportunity to work closely together and share their views and findings with the international elite of their field. The seminars thrive on an open-ended program that allows participants to take advantage of synergies as they come up over several days, creating a dynamic space for discussion and debate that often leads in unexpected directions. By contrast, Dagstuhl Perspectives Workshops bring together a group of well-established senior researchers to discuss a topic area and its perspectives. The goal is to analyze the overall state of the field in order to detect strategic trends and develop new perspectives on its continued evolution. The results are collected and published in a Dagstuhl Manifesto, which is made available to policymakers.

Each Dagstuhl Seminar and Dagstuhl Perspectives Workshop is headed by a small group of scientists of international standing in their respective fields. Proposals are reviewed by the Dagstuhl Scientific Directorate before their acceptance into the center's scientific program. Participation in these events is possible by way of personal invitation only by the center, which assumes part of the associated costs in order to enable the world's most qualified scientists to participate.

Located in the idyllic countryside of northern Saarland at the heart of the tri-country region formed by Germany, France and Luxembourg, Schloss Dagstuhl offers visitors a unique working environment that encourages guests to interact with each other in tandem with daily life. Lounges, formal and informal dining areas, a world-class research library, and an impressive range of work and leisure rooms

bietet den Gästen eine einzigartige Arbeitsumgebung, die den Austausch mit anderen Gästen in einer wohnlichen Atmosphäre fördert. Gemütliche Sitzecken, ansprechende Essräume, eine der besten Informatik-Fachbibliotheken weltweit, sowie eine Vielzahl von zusätzlichen Arbeits- und Freizeiträumen bieten vielfältige Möglichkeiten, damit sich die Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

Schloss Dagstuhl wird durch eine Bund-Länder-Förderung finanziert und beherbergt jedes Jahr mehr als 3 500 internationale Gäste. Seit 2006 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 89 Forschungsinstituten, Bibliotheken und Museen.¹

offer multiple possibilities for connecting one-on-one outside of the official conference rooms and meeting times.

Schloss Dagstuhl is jointly funded by the German federal and state governments and hosts over 3,500 research guests each year from countries across the globe. Since 2006, it has been a member of the Leibniz Association, a non-profit research consortium composed of 89 research institutes, libraries and museums throughout Germany.¹

¹ Stand: November, 2013
As of November, 2013



Fig. 1.1

The Schloss part of Schloss Dagstuhl. Photo reprinted with permission from 2013 seminar participant Robert Kosara.

Neuigkeiten in 2013

1.2

News from 2013

■ Das Team

Unter der Leitung des Wissenschaftlichen Direktors Professor Dr. Reinhard Wilhelm und des Technisch-administrativen Geschäftsführers Dr. Christian Lindig ist Schloss Dagstuhl – Leibniz-Zentrum für Informatik an drei Standorten tätig: die Geschäftsstelle auf dem Universitäts-campus in Saarbrücken, die Büros des dblp-Teams auf dem Campus 2 der Universität Trier sowie das eigentliche Konferenzzentrum im Schloss Dagstuhl in Wadern. Der wissenschaftliche Stab von Schloss Dagstuhl arbeitet mit jeweils drei Mitarbeitern an den Standorten Saarbrücken und Trier.

Schloss Dagstuhl hat 2013 Tina Schneider als Mitarbeiterin für die Küche und die Studentin Antónia Košťová zur zeitweisen Verstärkung der Geschäftsstelle eingestellt, um bei stetig wachsender Anzahl von Seminaren und Gästen einen gleichbleibenden Service zu gewährleisten. Im Berichtsjahr konnte Schloss Dagstuhl mit dieser vergleichsweise kleinen personalen Verstärkung und einem verbesserten Arbeitsablauf den gestiegenen Arbeitsaufwand bei der Seminarorganisation und in der Gästebetreuung auffangen. Das Zentrum unterstützt auch die Ausbildung junger Leute. In Küche und Hausdienst hat Lisa Pütz im Juli erfolgreich ihre Ausbildung beendet. Derzeit ist Sara Flis im selben Bereich in der Ausbildung. In der IT-Abteilung haben die vier Schüler Daniel Berwanger, Philipp Gehlen, Yannick Schillo und Felix Schonarth sowie der Student Dominik Michels ein Praktikum abgeleistet.

Das Leibniz-geförderte SAW-Projekt „LZI+DBLP“ endete im Juni 2013. Mit dem Projektende liefen auch die projektbezogene Stellen aus und Stefanie von Keutz, Katharina Hostert und Christopher Perrin beendeten ihre Tätigkeit. Schloss Dagstuhl wird aber weiterhin in enger Zusammenarbeit mit der Universität Trier die Literaturdatenbank dblp betreiben. Dazu wurden drei neue Stellen im wissenschaftlichen Stab von Schloss Dagstuhl geschaffen, so dass Dr. Marcel R. Ackermann, Oliver Hoffmann und Dr. Michael Wagner weiterhin mit dem Betrieb von dblp beauftragt sind. Dabei wird Dr. Wagner zu gleichen Teilen auch Aufgaben für das Dagstuhl Publishing übernehmen, während Dr. Ackermann zudem die Koordination der Zusammenarbeit im Leibniz-Forschungsverbund Science 2.0 übernimmt. Schloss Dagstuhl verfügt damit nun über einen wissenschaftlichen Stab bestehend aus sechs Vollzeit-Mitarbeitern.

Alle Mitarbeiterstellen von Schloss Dagstuhl wurden aus dem Kernhaushalt des Zentrums finanziert. Ausnahme ist das dblp-Team, das in Teilen durch eine großzügigen Spende der Klaus Tschira Stiftung in Höhe von 60 000 € unterstützt wurde. Ende 2013 beschäftigte Schloss Dagstuhl insgesamt 33,3 Vollzeitäquivalente.

■ The Team

Headed by Schloss Dagstuhl's Scientific Director, Professor Dr. Reinhard Wilhelm, and Technical Administrative Director Dr. Christian Lindig, the center now extends over three locations: the Dagstuhl Office on the University of Saarland campus in Saarbrücken; the dblp headquarters on Campus 2 at the University of Trier; and the Schloss Dagstuhl conference center itself in Wadern. The Schloss Dagstuhl scientific staff works out of the Saarbrücken and Trier sites, with three staff members based in each location.

To keep up with its growing number of seminars and guests, Schloss Dagstuhl reinforced its kitchen and household staff by one position, filled by Tina Schneider, and hired student Antónia Košťová to help the administrative team in the Dagstuhl Office. With these small personnel changes and an optimized workflow design, the center was able to meet the challenge of its newly expanded workload with respect to seminar organization and guest services during the year under review. Dagstuhl also supported the professional training of several young people, including two successive apprentices in the Kitchen and Household department and five student summer interns in the IT department. Sara Flis began training in the kitchen as a replacement for Lisa Pütz, who successfully completed her apprenticeship in July. The IT interns included Daniel Berwanger, Philipp Gehlen, Yannick Schillo, Felix Schonarth, and Dominik Michels.

Three new positions on the scientific staff of Schloss Dagstuhl were also created in 2013 so that Dr. Marcel R. Ackermann, Oliver Hoffmann and Dr. Michael Wagner could continue to carry out the work of dblp – operated by Schloss Dagstuhl in close cooperation with the University of Trier – following the end of the Leibniz-funded SAW project “LZI+DBLP” in June 2013. Dr. Wagner shares this work in equal measure with Dagstuhl Publishing activities, while Dr. Ackermann is also responsible for coordinating the center's cooperative work with the Leibniz Research Network Science 2.0. The finalization of “LZI+DBLP” marked the completion of the project work carried out by Stefanie von Keutz, Catherine Hostert and Christopher Perrin, who left our staff in 2013. The center is now supported by a total of six full-time scientific staff members.

All staff at Schloss Dagstuhl were funded from the center's core budget in 2013, with the exception of the dblp team, which was partially supported in 2013 by a generous donation of 60,000 € from the Klaus Tschira Foundation. A total of 33.3 full-time equivalent staff members worked at Schloss Dagstuhl in 2013.

■ Dagstuhl Publishing

Die Open-Access-Verlagsangebote von Schloss Dagstuhl wurden auch in 2013 vielfach genutzt. Die Anzahl an Veröffentlichungen in den Konferenzband-Serien LIPIcs und OASICS wurde im Vergleich zum Vorjahr moderat gesteigert. Erfreulich war die Veröffentlichung von drei Bänden in der Buchreihe *Dagstuhl Follow-Ups*, welche als umfassende Ergebnissammlung der jeweils zugrundeliegenden Seminare verstanden werden können.

Mit der Gründung der wissenschaftlichen Open-Access-Zeitschrift *Leibniz Transactions on Embedded Systems* (LITES) wurde in 2012 Neuland betreten. Im Februar 2013 wurde die Zeitschrift offiziell in Gang gesetzt und es dauerte nicht lange, bis die ersten Artikel zur Begutachtung eingereicht wurden. Die Veröffentlichung der ersten akzeptierten Artikel ist nun für das Frühjahr 2014 geplant.

Generell ist festzuhalten, dass Forschungsförderer weltweit verstärkt in geförderten Projekten Open-Access-Mandate erteilen. Dies wird z. B. bei den formalen Kriterien für das EU-Förderprogramm *Horizon 2020* deutlich, bei dem eine Open-Access-Veröffentlichung der Projektergebnisse entweder auf dem goldenen oder dem grünen Weg als Standard vorausgesetzt wird, und zudem als Pilot-Initiative zur Open-Data-Veröffentlichung von Forschungsdaten angehalten wird.

Mehr Informationen zu den Open-Access-Aktivitäten von Schloss Dagstuhl finden sich in Kapitel 6.

■ Bibliophiedatenbank dblp

Die Literatur-Datenbank dblp hat sich seit ihrer Gründung im Jahre 1993 zu einem mächtigen Werkzeug entwickelt, das von Informatikforschern weltweit bei ihrer täglichen Arbeit zur Suche nach Grundlagen, Ideen und Experten genutzt wird. Dr. Michael Ley (Universität Trier) reagierte mit der Gründung von dblp auf die spezielle Publikationskultur in der Informatik, in der die oft schwer zu recherchierenden Konferenzbeiträge ein höheres Gewicht als Veröffentlichungen in Fachjournalen haben. Bereits seit der Gründung sah dblp sich dem Ideal der „offenen Daten“ (Open-Data) verpflichtet. Alle Daten stehen der internationalen Forschung frei zur allgemeinen Nutzen zur Verfügung und erlauben einen einzigartigen Einblick in die komplexen Zusammenhänge der internationalen Informatikforschung. Die Datenbank genießt international ein hohes Ansehen und wurde in der Vergangenheit bereits mit verschiedenen Preisen wie dem „ACM SIGMOD Contributions Award“ ausgezeichnet.

Bereits seit 2011 besteht eine Kooperation zwischen Schloss Dagstuhl und der Universität Trier für den gemeinsamen Betrieb von dblp. Ziel war und ist es, die thematische Breite der Datenbank zu vergrößern und das Angebot langfristig zu verstetigen. Dabei legt das dblp-Team großen Wert auf Verlässlichkeit und Qualität der Einträge.

Der Erfolg der Zusammenarbeit kann sich auch in Zahlen sehen lassen: Jeden Monat erfolgen mehr als eine Million Zugriffe aus aller Welt auf die Sammlung; jedes

■ Dagstuhl Publishing

The Open Access publishing services offered by Schloss Dagstuhl were frequently used throughout 2013. The conference proceedings series LIPIcs and OASICS saw a moderately increased number of publications and articles and the center was also pleased to publish three volumes in the monographs series *Dagstuhl Follow-Ups*, which can be seen as comprehensive results from the underlying seminars.

With the creation of the scholarly Open Access journal *Leibniz Transactions on Embedded Systems* (LITES), Schloss Dagstuhl entered new territory in 2012. The journal was officially launched in February 2013 and it didn't take long until the very first articles had been submitted for review. The first accepted articles will be published in spring, 2014.

In general, it became clear in 2013 that funding agencies worldwide had strengthened the requirements for making results from funded projects freely available by formally stating Open Access mandates. This becomes apparent, for example, when looking at the formal criteria of the *Horizon 2020* European funding program, which explicitly commits grantees to publish project results either via the golden or the green Open Access road and, additionally, provides incentives for making research data freely available as part of a pilot project.

More information about the Open Access activities of Schloss Dagstuhl can be found in Chapter 6.

■ dblp computer science bibliography

Since its inception in 1993, the dblp computer science bibliography has evolved into a powerful tool that is used worldwide by computer science researchers in their daily work, e.g., when searching for research papers, emerging trends, or experts. Dr. Michael Ley (University of Trier) reacted with the creation of dblp to the specific publication culture in computer science, in which conference articles often have a higher impact than publications in journals. Even since the early days of dblp, the database was committed to the ideal of open data and has released its data to be freely available for the general benefit of the international research community. dblp enjoys a high reputation and has won several awards such as the “ACM SIGMOD Contributions Award.”

In 2011, Schloss Dagstuhl and the University of Trier agreed to operate and maintain dblp jointly. The goal of the collaboration is to improve and extend dblp, and to guarantee the dblp service to the international computer science community long-term. One major focus of the dblp team is the reliability and quality of the provided data.

The success of the collaboration is already quite impressive: More than a million web pages are visited each month from researchers all over the world. Each year, the database grows by more than 300,000 new entries. By the end of 2013, dblp had already indexed about 2.5 million scholarly publications from all disciplines of computer science.

Jahr wächst die Datenbank um mehr als 300 000 Einträge. Ende 2013 indexierte dblp bereits knapp 2,5 Millionen Fachartikel aus den verschiedenen Teilgebieten der Informatik.

Das Jahr 2013 markierte zudem den zwanzigsten Geburtstag des dblp-Dienstes. Dieses Jubiläum wurde am 4. Juli 2013 mit einem Festkolloquium an der Universität Trier gewürdigt.

Mehr Informationen zu dblp finden sich in Kapitel 7.

■ Leibniz-Forschungsverbund „Science 2.0“

Im November 2013 trat Schloss Dagstuhl dem Leibniz-Forschungsverbund „Science 2.0“ bei. Forschungsverbünde sind eine Initiative der Leibniz-Gemeinschaft zur Förderung von inter- und transdisziplinären Forschungsvorhaben, die einen wesentlichen Beitrag zur Erforschung von aktuellen sowie wissenschaftlich und gesellschaftlich relevanten Themen leisten. Forschungsverbünde sind mit einer Perspektive von fünf bis fünfzehn Jahren angelegt und offen für die Zusammenarbeit mit Universitäten, anderen außeruniversitären Forschungs- und Infrastruktureinrichtungen sowie ausländischen Forschungsgruppen.

Der Leibniz-Forschungsverbund „Science 2.0“ beschäftigt sich mit dem Studium neuer Gewohnheiten und Trends in Forschung und Entwicklung, welche durch die Nutzung moderner, partizipativer und kollaborativer Internet-Technologien in allen Bereichen der Forschung möglich geworden sind. Insbesondere konzentriert sich der Forschungsverbund dabei auf die Themenfelder

- *Neue Arbeitsgewohnheiten:* Wie verändert das Social Web Arbeitsgewohnheiten von Forschenden? Wie verändert das Internet die heutigen Forschungs- und Publikationsprozesse in den unterschiedlichen Wissenschaftsdisziplinen?
- *Technologieentwicklung:* Wie können die Forschungsprozesse durch Science 2.0 Unterstützung finden? Wie können Werkzeuge des Science 2.0 heutige Forschungsprozesse innovieren und beschleunigen?
- *Nutzungsforschung:* Welche neuen Formen der Wissenschaftskommunikation innerhalb der Forschungsgemeinschaft, aber auch zwischen Öffentlichkeit und Forschungsgemeinschaft werden durch Science 2.0 ermöglicht? Wie werden die dafür nötigen Werkzeuge genutzt?

Als Infrastruktureinrichtung kann Schloss Dagstuhl für viele der innerhalb des Forschungsverbundes initiierten Projekte eine ausgezeichnete Rolle als Use-Case-Partner spielen. Insbesondere Dagstuhls Erfahrung in den Bereichen Publikationswesen und Indexierung, sowie das Engagement für Open-Access und Open-Data, erlauben es Schloss Dagstuhl, diese Projekte zu unterstützen und sich dadurch als wertvolles Mitglied der Leibniz-Gemeinschaft stärker zu integrieren.

The year 2013 also marked the 20th anniversary of the dblp service. The anniversary of dblp was celebrated on July 4, 2013, with a honorary colloquium at the University of Trier.

More information about dblp can be found in Chapter 7.

■ Leibniz Research Alliance “Science 2.0”

In November 2013, Schloss Dagstuhl joined the Leibniz Research Alliance “Science 2.0.” Leibniz Research Alliances are a new initiative of the Leibniz Association to support inter- and transdisciplinary approaches to address current scientific and socially-relevant issues. The alliances are designed for a period of five to fifteen years and are open for collaboration with universities, other non-university research institutions and infrastructure facilities as well as research groups abroad.

The Leibniz Research Alliance “Science 2.0” deals with the study of emerging trends for research and development, originating from the application of new participative and collaborative Internet technologies in all phases of research. In particular, the research alliance focuses on the study of

- *new working habits:* How does the Internet with its new possibilities, particularly the social web, change working habits of researchers? How does it impact on today’s research and publication processes in different research disciplines?
- *technology development:* How can Science 2.0 support existing research processes? How can today’s research processes be innovated by Science 2.0 tools?
- *user behavior:* Which new forms of scientific communication – within the research community and between the public and the research community – does Science 2.0 facilitate? How do researchers use new Science 2.0 tools?

As a infrastructure facility, Schloss Dagstuhl is a perfect fit as a use-case partner for many projects within the research alliance. Particularly the experiences from the publishing and indexing departments and the commitment to the Open Access and Open Data movement enable Schloss Dagstuhl to contribute to these projects, thereby improving Dagstuhl’s integration within the Leibniz Association.

■ Öffentlichkeitsarbeit und Weiterbildung

Um junge Journalisten und Volontäre zu ermutigen, über anspruchsvolle Informatikthemen zu berichten, bietet Schloss Dagstuhl jährlich einen Workshop zum Thema Wissenschaftsjournalismus an. 2013 fand dieser parallel zu dem Dagstuhl-Seminar 13241 „Virtual Realities“ vom 9. bis 12. Juni statt. Als Dozenten für den Workshop konnten Tim Schröder (Wissenschaftsjournalist und Medientrainer, Oldenburg) und Gordon Bolduan (Pressesprecher des Exzellenz-Cluster „Multimodal Computing and Interaction“ an der Universität des Saarlandes) gewonnen werden. Alle Teilnehmer als auch die Dozenten waren höchst zufrieden mit den Inhalten und Ergebnissen des Workshops.

Schloss Dagstuhl engagiert sich im schulischen Bereich durch Organisation einer jährlichen Lehrerfortbildung, die sich an Informatiklehrer im Saarland und in Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Das Interesse an dieser Fortbildung stieg seit dem Beginn in 1991 stetig an und die 23. „Lehrerfortbildung in Informatik“, die vom 11. bis 13. Dezember 2013 stattfand, führte mehr Teilnehmer zusammen als jemals zuvor. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch häufen sich Anfragen zur Teilnahme von Lehrern aus anderen Bundesländern.

Schloss Dagstuhl war auch bei der Betreuung eines studentischen Projekts im Rahmen des Software Engineering Kurses im Wintersemester 2013/14 an der Fakultät für Mathematik und Informatik der Universität des Saarlandes unter der Leitung von Prof. Andreas Zeller beteiligt. Ziel des Projektes war, die Literaturdatenbank dblp in Publikationswerkzeuge einzubinden, die im täglichen Gebrauch der Publikationsabteilung von Dagstuhl sind.

Mehr Informationen zur Öffentlichkeitsarbeit und zu den Weiterbildungsaktivitäten finden sich in Kapitel 5.

■ Besuchergruppen

In Ergänzung zu dem wissenschaftlichen Programm, empfängt Schloss Dagstuhl auch interessierte Gruppen aus dem öffentlichen Leben. Ein besonderes Highlight in diesem Bereich war ein dreitägiges studentisches Kunstprojekt, Teil einer interdisziplinären Projektreihe von der Hochschule für Bildende Künste Saarbrücken (HBKsaar). Vom 10. bis 14. Februar 2013 beherbergte das Zentrum eine Gruppe von zehn Studenten für Fotografie und Bildende Kunst sowie ihre Professorin Gabriele Langendorf von der HBKsaar. Die Studenten, die bereits an ähnlichen Projekten mit der Deutschen Radio Philharmonie Saarbrücken Kaiserslautern und dem Landtag teilgenommen hatten, hatten den Auftrag, ihre Eindrücke der Dagstuhl-Seminare 13071 und 13072 darzustellen, ohne die Seminarteilnehmer zu stören. Sowohl Künstler als auch Seminargäste erzählten später, es sei für alle ein sehr positives Erlebnis gewesen und die Mischung habe die kreative Synergie beider Gruppen noch verstärkt.

■ Public Relations and Professional Training

In order to encourage young journalists and trainees to report on complex informatics topics, Schloss Dagstuhl offers an annual workshop on science journalism. In 2013, the workshop took place from June 9–12 in parallel to Dagstuhl Seminar 13241, “Virtual Realities.” Trainers included Tim Schröder from Oldenburg (scientific writer and media trainer) and Gordon Bolduan from Saarland University (press relations officer at the Cluster of Excellence “Multimodal Computing and Interaction”). Participants as well as trainers and referees were very satisfied with the workshop.

Schloss Dagstuhl also holds an annual teacher training workshop specifically designed for teachers of secondary students working in the Saarland or the Rhineland Palatinate. The workshop is organized together with the Landesinstitut Pädagogik und Medien (LPM), Saarland, and the Pädagogisches Landesinstitut Rheinland-Pfalz (PL). Interest in the workshop has risen steadily since the program began in 1991 and the 23th annual Dagstuhl teacher training workshop, held at Schloss Dagstuhl on December 11–13, 2013, attracted more participants than ever before. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, Schloss Dagstuhl does receive requests for participation from teachers of other federal states.

The center was also involved in supervising a student project at Saarland University within the context of Prof. Andreas Zeller’s (Winter Semester 2013/14) Software Engineering course in the Department of Computer Science. The project aimed to technically integrate the dblp literature database into the authoring workflow as it occurs in the daily routine of the Dagstuhl Publishing department.

Further details about public relations and professional training at Schloss Dagstuhl can be found in Chapter 5.

■ Community Outreach

In parallel to its scientific program, Schloss Dagstuhl often opens its doors to special visits from community representatives and groups. One 2013 highlight in this area was the three-day “HBKsaar Student Research Project,” part of an interdisciplinary project series from the Saarbrücken University of Art and Design. On February 10–14, the center hosted a group of ten photography and fine art students from the HBKsaar together with their instructor, Professor Gabriele Langendorf. The students, who had already participated in similar projects with the German Radio Philharmonic Orchestra and Saarland state parliament, were instructed to hang around the Schloss and portray their impressions of Dagstuhl Seminars 13071 and 13072 without disturbing seminar guests. Artists and seminar participants alike later reported that the experience had noticeably amplified the creative synergy of both groups.

Reporter des SR Fernsehen dokumentierten die Veranstaltung im Nachrichtenmagazin „Aktueller Bericht“ vom 14. Februar 2013. Die Fernsehreportage zeigt Nahaufnahmen der von den Studenten angefertigten Zeichnungen und Fotos, Interviews mit den Künstlern und ein Gespräch mit Reinhard Wilhelm. Er gab einen kurzen Überblick zu dem Kunstprogramm von Schloss Dagstuhl und betonte die Gemeinsamkeit des „abstrakten Denkens“, die Informatiker und Künstler verbindet.

Ebenfalls interessiert an der Arbeit von Schloss Dagstuhl und der Internationalität der Dagstuhl-Seminare war der Chef der Staatskanzlei des Saarlandes, Jürgen Lennartz, der Schloss Dagstuhl zusammen mit seinen Mitarbeitern Christian Mees und Jochen Krämer am 6. August 2013 besucht hat. Dr. Roswitha Bardohl und Dr. Christian Lindig haben über die Arbeit von Schloss Dagstuhl berichtet und die Delegation durch das Haus geführt.

Der Besuch des Rotary-Club St. Wendel am Abend des 6. Mai 2013 hatte einen eher regionalen Charakter. Dr. Roswitha Bardohl begrüßte die 15 Clubmitglieder, gab den Interessenten einen Überblick über die Arbeit von Schloss Dagstuhl und führte sie danach durch das Schloss. Der Besuch wurde mit einem kleinen Imbiss abgeschlossen.

Natürlich fand auch wieder die jährliche Feuerwehrrübung auf Schloss Dagstuhl statt. Das Zentrum eignet sich aufgrund seiner Bauweise hervorragend dafür, so dass einmal jährlich eine Übung stattfindet. In 2013 wurde die Feuerwehrrübung im Juli durchgeführt. Mit ca. 60 Feuerwehrmännern und -frauen aus den Löschbezirken Wadern, Noswendel, Losheim, Hermeskeil und dem Rettungsdienst der Rettungswache Wadern fand sie eine größere Beteiligung als jemals zuvor. Für die Übung fanden sich zehn Freiwillige, die aus den schlecht zugänglichen Räumen im 2. Stock des Schlosses über zwei Drehleitern gerettet wurden. Parallel zu der Rettung bekämpften die Feuerwehrmänner einen simulierten Brand.

■ Spender und Förderer von Schloss Dagstuhl

Schloss Dagstuhl ist seinen wissenschaftlichen Gästen, Institutionen und Firmen dankbar, die durch großzügige Spenden das Zentrum und seine Bibliothek unterstützen.

Neben zahlreichen Buchspenden durch Autoren und Gäste, erhielt die Bibliothek auch 2013 wieder eine großzügige Spende des Springer Verlags. Insgesamt erhielt die Bibliothek von ihm 877 Monographien, die meisten aus der Serie *Lecture Notes in Computer Science*, im Wert von 60 000 €.

Wie in den Jahren 2011 und 2012 förderte die Klaus Tschira Stiftung 2013 erneut das „LZI+DBLP“-Projekt mit einer Spende von 60 000 €. Dieses Projekt wurde bis Juli 2013 durch eine Förderung des Senatsausschuss Wettbewerb (SAW) der Leibniz-Gemeinschaft in Höhe von 167 094 € finanziert (SAW-2011-LZI-3). Die Förderung belief sich auf insgesamt 336 824 € und erstreckte sich über einen Zeitraum von zwei Jahren, beginnend Mitte 2011.

Schloss Dagstuhl Publishing – vor allem das Open-Access Journal *Leibniz Transactions on Embedded Systems*

The “Aktueller Bericht” news program of SR TV featured the event on February 14, 2013, giving close-ups of the students’ drawings and photos, artist interviews, and a discussion with Reinhard Wilhelm. He gave a brief overview of Dagstuhl’s art program and stressed the abstract thinking that connects computer scientists and artists.

Also interested in the work of Schloss Dagstuhl and the internationality of the Dagstuhl Seminar series was Jürgen Lennartz, head of the State Chancellery of the German State of Saarland. Mr. Lennartz visited the center together with his colleagues Christian Mees and Jochen Krämer on August 6, 2013. Dr. Roswitha Bardohl and Dr. Christian Lindig, the Technical-Administrative Director of Schloss Dagstuhl, gave the delegation a report on the activities at Schloss Dagstuhl and led them on a guided tour of our house.

Of a more local nature was the visit made by the St. Wendel Rotary Club on the evening of May 6, 2013. Scientific staff member Dr. Roswitha Bardohl welcomed 15 club members and gave them an overview of the center, followed by a guided tour. The visit concluded with a snack.

The year would not have been complete without the annual fire rescue operation at Schloss Dagstuhl. Because of its unusual construction features, Schloss Dagstuhl is particularly suitable for fire drills and offers its grounds for this purpose once a year. In 2013, the drill was conducted on July 12 and involved about 60 firemen and firewomen from the Wadern, Noswendel, Losheim, and Hermeskeil voluntary fire brigades, plus the emergency rescue service of the Wadern ambulance station – more participants than ever. During the exercise, ten volunteers were rescued from the poorly accessible rooms on the third floor of the Schloss via two ladders. In parallel, firefighters fought a simulated fire.

■ Sponsors and Donors of the Center

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generously donating funds to support its core scientific work and books to its research library.

In addition to numerous private book donations from guests, the center’s scientific library received 877 monographs worth 60,000 € in 2013 from the German publisher Springer. A large share of the donated monographs belong to the series *Lecture Notes in Computer Science*.

With respect to project grants, the “LZI+DBLP” project received a grant of 60,000 € from Klaus Tschira Foundation, which had already generously donated the project in 2011 and 2012. The project was further financed until July 2013 by a grant of 167,094 € from the Senatsausschuss Wettbewerb (SAW) of the Leibniz Association (SAW-2011-LZI-3). The grant provided a total of 336,824 € for a two-year period starting in mid-2011.

Schloss Dagstuhl Publishing – specifically the open-access journal *Leibniz Transactions on Embedded Systems*

(LITES) – erhielt 2013 substantielle Unterstützung durch Google Deutschland und die Klaus Tschira Stiftung in Höhe von 10 000 US-Dollar (ca. 7 250 €) bzw. 5 000 €. Weitere Informationen über LITES finden sich auf der LITES Webseite².

Neben diesen Spenden für das wissenschaftliche Angebot wird ebenso die Kunstsammlung von Schloss Dagstuhl seit dem Bestehen des Zentrums von Kunstmäzenen kontinuierlich durch Sach- und Geldspenden gefördert. Auch 2013 erhielt Schloss Dagstuhl dankenswerterweise mehrere solcher Spenden.

„Informatiker als künstlerisches Thema“ hätte der Titel eines ungewöhnlichen Geschenks der Studenten von Professor Gabriele Langendorf an der Hochschule für Bildende Künste Saar (HBK Saar) sein können. Die Gruppe stiftete im April 2013 zwei gerahmte Zusammenstellungen von Skizzen und Fotos, die die Teilnehmer der Seminare 13071 und 13072 aus der Sicht der Kunststudenten während ihres dreitägigen Aufenthalts im Schloss Dagstuhl im Februar 2013 zeigen.

Zum Thema „Informatiker als Künstler“ erhielt Schloss Dagstuhl im Juli 2013 als Leihgabe ein Gemälde des deutschen Informatikers Konrad Zuse (1910–1995). Im Laufe seines Lebens malte Konrad Zuse zahlreiche abstrakte Werke und Porträts. Das Gemälde wurde freundlicherweise von Dr. Jürgen Alex zur Verfügung gestellt und ist zur Zeit im Vorraum der Cafeteria des Zentrums zu sehen.

Weiterhin spendete Professor Ben Shneiderman vom Institute of Advanced Computer Studies, University of Maryland sein Werk „Blooming Businesses“. Das Bild wurde auf Bitte Shneidermans persönlich von Wissenschaftlichen Direktor Prof. Reinhard Wilhelm ausgewählt. Es entstammt einer Serie von 12 Bildern, die mit Hilfe des Treemap-Algorithmus zur Datenvisualisierung geschaffen wurden. Sowohl die Werke der Kunststudenten als auch das Werk von Ben Shneiderman werden derzeit im Gästehaus ausgestellt.

Für seine historische Kunstaustellung erhielt Schloss Dagstuhl als Spende von Manfred Stein eine fotografische Reproduktion eines Portraits von Octavie de Lasalle von Louisenthal (1811–1890), deren leidenschaftliches und künstlerisches Leben sie zu einer der faszinierendsten Personen des historischen Dagstuhls macht. Das Foto, das schon seit den frühesten Anfängen Teil der Sammlung von Schloss Dagstuhl war, wurde im Juni 2013 durch eine neue Reproduktion ersetzt und mit einem Rahmen versehen, der zu einer weiteren Reproduktion eines Selbstporträts der jungen Gräfin passt. Diese Reproduktion hatte uns Herr Stein bereits 2012 geschenkt.

■ Änderungen im Gästeservice

Schloss Dagstuhl bietet allen Teilnehmern von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops, die mit Kindern anreisen, eine qualifizierte Kinderbetreu-

(LITES) – also received strong external support in 2013 in the form of 10,000 US-Dollar (approx. 7,250 €) and 5,000 € from Google Germany and the Klaus Tschira Foundation, respectively. See the LITES website² for further details.

In parallel to these scientific donations, the center is fortunate to count as friends several patrons of the arts who have, over the years, greatly enriched our art collection through personal and financial gifts. In 2013 Schloss Dagstuhl was pleased to accept several such donations.

“The computer scientist as artistic subject” might have been the title of the unusual gift presented by the students of Professor Gabriele Langendorf of the Saarbrücken University of Art and Design (HBK Saar) in April, 2013. The group donated two framed collections of sketches and photos featuring the participants in Dagstuhl Seminars 13071 and 13072, as depicted by the students during their three-day stay at Schloss Dagstuhl in February of 2013.

On the topic of “computer scientist as artist,” in July of 2013 Schloss Dagstuhl received a picture by German computer scientist Konrad Zuse (1910–1995). Zuse’s passion for art led him to paint numerous abstract works and portraits throughout his lifetime. The painting was loaned to Dagstuhl courtesy of Dr. Jürgen Alex and is currently on display in the foyer of the Schloss café.

In addition, Professor Ben Shneiderman of the Institute of Advanced Computer Studies at the University of Maryland donated the image “Blooming Businesses” to Dagstuhl as a gift in November of 2013. The image was personally selected at Prof. Shneiderman’s request by Dagstuhl’s Scientific Director, Reinhard Wilhelm, and is one of 12 such images created using the tree map algorithm for data visualization. Both the student art pieces and the tree map image are currently on display in Schloss Dagstuhl’s new guest house.

For its historical art collection, the center also received as a gift from Manfred Stein a framed photographic reproduction of a portrait of Octavie de Lasalle von Louisenthal (1811–1890), whose artistic and impassioned life make her one of Dagstuhl’s most fascinating historical personages. The photograph, which had been a part of Schloss Dagstuhl’s collection since its earliest days, was reprinted and reframed in June 2013 to match another photographic reproduction of the countess in her younger days. The latter had been donated to the center by Mr. Stein in 2012.

■ Changes in Dagstuhl Services and Facilities

Schloss Dagstuhl gladly offers to organize qualified child care for participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops who need to visit our

² <http://www.dagstuhl.de/lites/>



Fig. 1.2
Student artist Jennifer Lubahn sketches a Dagstuhl Seminar participant during the HBK Saar Student Research Project. Photograph courtesy of Ingeborg Knigge.



Fig. 1.3
The annual fire rescue drill at Schloss Dagstuhl.

ung für ihre Kinder an. Seit März 2013 hat Schloss Dagstuhl die Bedingungen zur Kinderbetreuung angepasst, um allen Eltern weiterhin eine zuverlässige Kinderbetreuung anzubieten. Eltern können bis zu einem Jahr im voraus eine Kinderbetreuung buchen. Zusammen mit der Buchung ist ein Eigenanteil von 100€ zu zahlen. Die Betreuung kann unter Erstattung dieses Anteils bis zu einigen Wochen vor dem Seminar storniert werden. Alternativ können Eltern auch eine Begleitperson zur Betreuung der Kinder mitbringen. Ebenso wie für die Kinder übernimmt Schloss Dagstuhl die Aufenthaltskosten einer Begleitperson zur Kinderbetreuung. In 2013 besuchten 21 Kinder Schloss Dagstuhl. Davon wurden 14 Kinder durch einen Tagesmutter und sieben weitere durch Verwandte betreut.

Seit 2012 haben wir die Kinderbetreuung auf Schloss Dagstuhl kontinuierlich verbessert. Unser Anliegen für 2013 war es, den Aufenthaltsraum als Spielraum noch attraktiver zu gestalten und zusammen mit unserer Tagesmutter die Kommunikation mit den Eltern und Seminarteilnehmern, die die Kinderbetreuung in Anspruch nehmen, zu verbessern. Während den Tagungszeiten haben die Kinder die Dagstuhl-Lounge mit Kindertisch und -stühlen ganz für sich. Viele unserer jüngsten Gäste haben Spaß daran, Zeichnungen, Kommentare und Fotos in Dagstuhls „Kindergästebuch“ zu hinterlassen.

Als Ergebnis eines zweiwöchigen Schülerpraktikums haben alle Gäste nun die Möglichkeit, sich auf einer Dagstuhl-eigenen Webseite³ über das Wetter in Dagstuhl zu informieren. Die von den vier Schülern Felix Schonarth, Philipp Gehlen, Daniel Berwanger und Yannick Schillo der Gymnasien in Wadern und Hermeskeil realisierte Webseite zeigt die aktuelle Temperatur und Luftfeuchtigkeit an, sowie ein aktuelles Bild des Schloss-Gartens. Aus diesen Bildern wird für die vergangenen sieben Tage ein Zeitraffer berechnet, der als Video angeboten wird. Zusätzlich zeigen Kurven den Verlauf und die Temperatur der Vergangenheit an. Die Seite ist auf der Website von Schloss Dagstuhl verlinkt und hilft Gästen vor ihrer Anreise beim Packen der richtigen Kleidung.

Schloss Dagstuhl hatte bereits 2012 die Dagstuhl Seminar Reports der Jahre 1990–1993 digitalisiert und auf der Werbeseite öffentlich zur Verfügung gestellt. Als Fortsetzung der Bemühung, unsere früheren Dokumente auch elektronisch der Allgemeinheit zur Verfügung zu stellen, haben wir 2013 die Tätigkeitsberichte der Jahre 1990 bis 1999 digitalisiert und auf unseren Webseiten veröffentlicht. Die Berichte sind überwiegend in Deutsch, englische Versionen stehen zur Verfügung, wenn sie vorhanden waren. Alle Jahresberichte können auf unserem Webserver⁴ abgerufen werden.

Verglichen mit der Neueröffnung unseres Gästehaus in 2012 gab es in 2013 nur vergleichsweise kleine Änderungen in Dagstuhls Ausstattung. Im Sommer wurde die Bruchsteinmauer, die das Hauptgebäude mit dem neuen Gästehaus verbindet, von der Firma Ditandy fertiggestellt. Eine moderne Gartentür der Firma Arendt aus senkrechten Holzplatten erlaubt den Zugang zu dem inneren Bereich und nimmt Bezug auf die senkrechten Holzplatten an der Fassade des Gästehaus.

center with young children. In order to ensure a more reliable service and allow parents to book child care up to one year in advance of a seminar or workshop, in March of 2013 Dagstuhl began implementing a child care reservation co-payment fee of 100€ per week. Parents pay the fee when they reserve the service and can receive a refund up to several weeks before the seminar if they need to cancel. Parents also have the option to bring along their own “nanny,” usually a spouse or relative, whose room and board costs are gladly absorbed by the center just as they are for children. In 2013, Dagstuhl hosted 21 children, 14 of whom were cared for by a nanny on site and seven by relatives.

Since 2012, Schloss Dagstuhl has been continually improving its child care program. In 2013 the center focused on enhancing the lounge as a play space and providing clearer information beforehand to parents and seminar-goers who use the service. Children now enjoy exclusive use of the lounge during seminar meeting times and the use of a child-sized table and chairs set. Our youngest guests also have fun writing notes and inserting drawings or photographs into their very own “Dagstuhl children’s guest book.”

Guests also now have access to a weather information web page,³ fruit of a two-week student internship in the IT department. Four high school students from Wadern and Hermeskeil – Felix Schonarth, Philipp Gehlen, Daniel Berwanger and Yannick Schillo – designed the page to display the current and past temperature and humidity at Schloss Dagstuhl, a real-time digital video image of the Schloss garden, and a graph of tracing the evolution of weather conditions during the last 24 hours. The site is featured on Schloss Dagstuhl’s homepage, where it serves to help guests pack the right clothing for their trip.

The center’s website also gained some new texts for its digital repository during the year under review. In 2012 Schloss Dagstuhl had digitalized the Dagstuhl Seminar Reports corresponding to the years 1990-1993 and made them publicly available on our web pages. As a continuation of this effort to offer free electronic versions of core documents from its archive, in 2013 the center retro-digitalized the Dagstuhl annual reports published from 1990 to 1999. The reports are in German and, when available, English. They can be read or downloaded from the Dagstuhl web page.⁴

Schloss Dagstuhl made only minor changes in its facilities in 2013, following the large-scale expansion of the year before. Summertime offered the opportunity to finish building the natural stone wall connecting the guest house to the main building and install a modern wooden garden gate. The Ditandy and Arendt companies installed the wall and garden gate, respectively. The vertical slats on the garden gate offer a pleasing visual reference to the vertical slats on the facade of the guest house.

Leisure facilities also received a facelift with the redesigning of the fitness room located in the basement of the new building. The old linoleum floor was replaced by a sports floor, the previously conspicuous installations below the ceiling were covered by a net fabric, and lights on the

Auch das Erholungsangebot von Schloss Dagstuhl wurde mit der Neugestaltung des Fitness-Raums im Keller des Neubaus verfeinert. Der Linoleum-Fußboden wurde durch einen Sport-Boden ersetzt und die zuvor auffälligen Installationen unter der Decke durch ein Netzgewebe verkleidet. Neue Lampen an den Unterzügen sorgen für viel Licht über der Tischtennisplatte. Die Ausstattung wurde durch Sportbänke, Hocker, einen Kühlschrank für Getränke und ein Regal für Sportgeräte ergänzt.

Für seine wechselnden Kunstausstellungen hat Schloss Dagstuhl 2013 weitere Teile seines Beleuchtungssystems im sogenannten Kreuzgang des Neubaus erneuert. Im Mittelflur und im Flur auf der Nordseite wurde das alte Beleuchtungssystem von 1993 von der Haustechnik durch ein Niedervolt-System der Firma Buschfeld ersetzt. Eine schmale Schiene dient gleichzeitig zur Aufhängung von Bildern und versorgt frei positionierbare Leuchten mit Strom. Die Leuchten sind zur Zeit mit Halogen-Glühhirnen ausgestattet und können später LED-Leuchten aufnehmen.

Heizkörper in ausgewählten öffentlichen Räumen im Altbau von Schloss Dagstuhl – Weinkeller, Cafeteria und Foyer vor der Cafeteria – wurden im Herbst 2013 mit drahtlos steuerbaren Heizungsventilen ausgestattet. Mit den Komponenten der Firma Homatic kann die Haustechnik zentral die aktuellen Temperaturen ablesen und tageszeitabhängig Vorgaben für die Temperatur machen. Eine intelligente Steuerung hilft, Komfort und Energiesparen zu verbinden und mögliche Probleme schnell zu erkennen.

Der Speisesaal im Haupttrakt des Zentrums erhielt im Berichtsjahr 105 neue Stühle „Vitra Hal Wood“, die die alten Stühle von 1990, die 1997 einmal neu gepolstert worden waren, ersetzen.

Seit Juli 2013 steht im Raum „Kaiserslautern“ allen Teilnehmern ein Apple TV zur Verfügung. Dieses erlaubt von kompatiblen Geräten drahtlos auf den Beamer zuzugreifen und macht es dadurch Seminarteilnehmern einfacher, ihre Präsentationen vorzuführen oder Materialien zu teilen.

beams installed to provide plenty of light for the ping pong table area. Fitness benches, stools, a refrigerator and a shelf for exercise equipment complement the new installations.

At a more basic level, the old 1993 lighting system that ran along the central corridor and hallway on the north side of the new building was replaced by a low-voltage system. A narrow rail now supports picture hangings and features position-adjustable lights outfitted for both halogen bulbs and LED lamps. The system serves to illuminate the artworks displayed in the temporary exhibits at Schloss Dagstuhl, in the so-called “cloister walk” area.

The heating system in various key parts of the main building – the oldest of the three – was also retrofitted during the year under review with wireless heating controls from the Homatic company. The system combines energy savings with the benefits of centralized meter readings that allows for fast identification of potential problems. The wireless components of Homatic are also interesting because of their low cost and easy integration into an IT infrastructure for buildings

Schloss Dagstuhl’s main building received some new moveable furnishings in 2013 as well. A total of 105 new chairs of the model “Vitra Hal Wood” were purchased to replace the old chairs, which dated back to 1990 and had already been reupholstered once in 1997.

Apple TV was also installed in the “Kaiserslautern” seminar room in July of 2013. The service wirelessly connects compatible devices to the projector in the room, making it easier to share and present materials during a seminar.

³ <https://portal.dagstuhl.de/weather/>

⁴ http://drops.dagstuhl.de/portals/dagstuhl_annual_reports/

2 **Das wissenschaftliche Programm 2013** *Scientific Program 2013*

Dagstuhl-Seminare

2.1

Dagstuhl Seminars

Die Dagstuhl-Seminare haben als wesentliches Instrument der Forschungsförderung Priorität bei der Gestaltung des Jahresprogramms. Hauptziel der Seminare ist die Unterstützung der Kommunikation und des Dialogs zwischen Wissenschaftlern, die in Randgebieten von miteinander verknüpften Forschungsfeldern in der Informatik arbeiten. Die Seminare ermöglichen die Vorstellung neuer Ideen, die Diskussion von aktuellen Problemen sowie die Weichenstellung für zukünftige Entwicklungen. Sie bieten außerdem die Möglichkeit zum Austausch zwischen vielversprechenden Nachwuchswissenschaftlern und internationalen Spitzenforschern in einem speziellen Forschungsgebiet.

Die Teilnahme an den üblicherweise einwöchigen Dagstuhl-Seminaren ist nur auf persönliche Einladung durch Schloss Dagstuhl möglich. Das Zentrum übernimmt einen Teil der Kosten, sodass die besten Wissenschaftler einschließlich junger Forscher und Doktoranden teilnehmen können. Zu den ehemaligen Gästen zählen 22 Preisträger des Turing-Awards, der höchsten Auszeichnung, die in der internationalen Informatik-Community verliehen wird.

Charakteristisch für Dagstuhl ist die Etablierung von richtungsweisenden sowie gebietsübergreifenden Seminaren. Manche Themen, die ausgiebig in Dagstuhl diskutiert wurden, entwickelten sich anschließend zu sehr aktiven Forschungsbereichen, die teilweise zu DFG-Schwerpunkten und anderen Förderprogrammen führten. Bei einer Reihe von Forschungsgebieten wurden durch Dagstuhl-Seminare Gruppen zusammengeführt, die zwar an verwandten Problemen und Verfahren forschen, denen aber bisher keine gemeinsame Diskussionsplattform zur Verfügung stand. Dies gilt insbesondere auch für Disziplinen, die nicht zur Informatik gehören. Wichtige Forschungsgebiete, für die in Dagstuhl bereits mehrfach eine intensive Zusammenarbeit mit der Informatik erschlossen und vertieft wurde, sind Biologie (seit 1992) und Sport (seit 2006). Die Themen der Dagstuhl-Seminare bieten eine hervorragende wenn nicht sogar erschöpfende Übersicht über die Gebiete der Informatik, die derzeit weltweit diskutiert werden.

Für jedes Dagstuhl-Seminar soll ein Bericht, genannt Dagstuhl Report, erstellt werden, der eine Zusammenfassung des Seminarverlaufs, eine Kurzübersicht über die gehaltenen Vorträge und eine Zusammenfassung grundsätzlicher Ergebnisse enthält. Der Bericht sollte in Artikelform verfasst werden, sodass Transparenz und zeitnahe Kommunikation der Ergebnisse gewährleistet sind. Die Zeitschrift *Dagstuhl Reports* wird jährlich in einem Band mit zwölf Ausgaben veröffentlicht. Jede Ausgabe dokumentiert jeweils die Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops eines Monats. Die Dagstuhl Reports sind frei zugänglich und können jederzeit von der Dagstuhl-Website⁵ heruntergeladen werden.

In den Kapiteln 4 und 14 finden sich Listen aller Veranstaltungen, die 2013 auf Schloss Dagstuhl stattfanden.

Dagstuhl Seminars, the center's key instrument for promoting research, are accorded top priority in its annual program. The central goal of the Dagstuhl Seminar program is to stimulate new research by fostering communication and dialogue between scientists working on the frontiers of knowledge in interconnected fields related to informatics. New ideas are showcased, topical problems are discussed, and the course is set for future development in the field. The seminars also provide a unique opportunity for promising young scientists to discuss their views and research findings with the international elite of their field in a specific cutting-edge field of informatics.

Participation in these events – which generally last one week – is possible only by way of personal invitation from Schloss Dagstuhl. The center assumes part of the associated costs in order to enable the world's most qualified scientists, including young researchers and doctoral students, to participate. Among Dagstuhl's alumni are 22 Turing Award laureates, the highest achievable award within the international computer science community.

Dagstuhl's distinguished accomplishment is to have established pioneering, interdisciplinary seminars that have virtually become institutions unto themselves. Many of the topics addressed in-depth at Dagstuhl subsequently develop into highly active research fields, resulting in some cases in DFG priority programs and other grant and funding programs. Dagstuhl Seminars often succeed in bringing together scientists from a range of research areas and disciplines whose work overlaps with respect to issues, methods and/or techniques, but who had never previously entered into constructive dialogue with one another. This especially applies to disciplines outside of the field of informatics. Key research areas for which in-depth collaboration with informatics specialists was initiated and consolidated at Dagstuhl include biology (since 1992) and sports (since 2006). The spectrum of seminar topics provides an excellent if not comprehensive view of the areas currently under discussion in the international informatics arena.

Each Dagstuhl Seminar is asked to contribute a record of the seminar proceedings in the form of a Dagstuhl Report. The report gives an overview of the seminar program, talks, and results obtained in a journal-like manner to allow for a high visibility and timely communication of its outcome. The periodical *Dagstuhl Reports* is published in one volume with twelve issues per year; each issue documents the Dagstuhl Seminars and Dagstuhl Perspectives Workshops of a given month. Dagstuhl Reports are open-access and can be downloaded at any time from the Dagstuhl webpages.⁵

Chapters 4 and 14 provide a comprehensive list of all events that took place at Schloss Dagstuhl during the year under review and summaries of the 2013 Seminars and

den, sowie Zusammenfassungen der Seminare und Perspektiven-Workshops. Auf der Dagstuhl-Website ist ein tagesaktuelles Programm für die kommenden 24 Monate verfügbar.

Dagstuhl-Perspektiven-Workshops

2.2

In Ergänzung zu den Dagstuhl-Seminaren werden seit 2004 Dagstuhl-Perspektiven-Workshops veranstaltet. An den Workshops nehmen meist 25–30 ausgewiesene Wissenschaftler teil, die ein bereits fest etabliertes Forschungsgebiet betreffende Tendenzen und neue Perspektiven der weiteren Entwicklung dieses Gebietes diskutieren. Im Gegensatz zu Dagstuhl-Seminaren werden statt aktueller Forschungsergebnisse im Wesentlichen Positionspapiere vorgetragen, welche den aktuellen Stand des Gebietes, offene Probleme, Defizite und vielversprechende Richtungen beschreiben. Der Fokus in den Workshops liegt auf Teilgebieten oder mehreren Gebieten der Informatik. Jeder Workshop hat zum Ziel

- den Stand eines Gebietes zu analysieren,
- Potenziale und Entwicklungsperspektiven bestehender Forschungsfelder zu erschließen,
- Defizite und problematische Entwicklungen insbesondere in der deutschen Forschungslandschaft aufzudecken,
- Forschungsrichtungen aufzuzeigen und
- Innovationsprozesse anzustoßen.

Die Dagstuhl-Perspektiven-Workshops, die 2013 statt fanden, sind in Fig. 2.1 aufgelistet.

Die Ergebnisse der intensiven Diskussionen werden in einem Manifest zusammengefasst, welches die offenen Probleme und die möglichen Forschungsperspektiven für die nächsten 5–10 Jahre aufzeigt. Dagstuhl koordiniert die gezielte Weitergabe dieses Manifests, um forschungsspezifische Impulse an deutsche und europäische Institutionen der Forschungsförderung zu geben (EU, BMBF, DFG, etc.). Kurzfassungen der Manifeste werden regelmäßig im Forum des *Informatik Spektrum* (Springer-Verlag) vorgestellt. Die vollständigen Manifeste werden in unserer Fachzeitschrift *Dagstuhl Manifestos* veröffentlicht.

Sowohl eine Liste der vergangenen und kommenden Dagstuhl-Perspektiven-Workshops⁶ als auch eine Liste der bereits erschienenen Dagstuhl Manifeste⁷ werden auf unseren Webseiten veröffentlicht.

⁵ <http://www.dagstuhl.de/dagrep/>

⁶ <http://www.dagstuhl.de/pw-list>

⁷ <http://www.dagstuhl.de/dagman>

Perspectives Workshops. An up-to-the-minute program covering the coming 24 months is available on the Dagstuhl website.

Dagstuhl Perspectives Workshops

In addition to the traditional Dagstuhl Seminars, since 2004 the center has organized Dagstuhl Perspectives Workshops. Perspectives Workshops are oriented towards a small group of 25–30 internationally renowned senior scientists who wish to discuss strategic trends in a key research area that is already well established and to develop new perspectives for its future evolution. In contrast to Dagstuhl Seminars, Perspectives Workshops do not address current research results but reflect the overall state of a field, identifying strengths and weaknesses, determining promising new developments, and detecting emergent problems and synergies. The workshops tend to focus on subfields or are interdisciplinary in nature, thus covering more than one informatics field. Each workshop aims to:

- contribute to an analysis of the present status of a field
- tap into potentials and development perspectives of existing fields of research
- detect shortcomings and problematic developments, particularly in the German research landscape
- show research directions
- trigger innovation processes

Dagstuhl Perspectives Workshops held in 2013 are listed in Fig. 2.1.

The results of the in-depth discussions of each workshop are presented in a manifesto detailing open issues and possible research perspectives in that specific field for the coming 5–10 years. Schloss Dagstuhl coordinates the targeted dissemination of this manifesto as research policy impulses to German and other European research donors and sponsors (EU, German Federal Ministry of Education and Research, DFG, etc.). Short versions of the manifestos are regularly presented in a forum of the *Informatik Spektrum* journal (published by Springer); full versions of the manifestos are published in our periodical *Dagstuhl Manifestos*.

Past and upcoming Dagstuhl Perspectives Workshops⁶ and published Dagstuhl Manifestos⁷ can be found on our web site.

Einreichung der Anträge und Begutachtungsverfahren

2.3

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch Auswahl der Anträge gewährleistet, aus denen am Wahrscheinlichsten erfolgreiche Seminare resultieren. Das Zentrum erbittet zweimal im Jahr Themenvorschläge von führenden Wissenschaftlerinnen und Wissenschaftlern aus der ganzen Welt, die ihre Seminaranträge zusammen mit einer vorläufigen Teilnehmerliste einreichen. Die Anträge und vorläufigen Gästelisten werden von mehreren Mitgliedern eines speziell hierfür beauftragten akademischen Gremiums, dem Wissenschaftlichen Direktorium von Schloss Dagstuhl, begutachtet.

Damit ein Antrag angenommen wird, muss er verschiedenen, vom wissenschaftlichen Direktorium festgelegten Kriterien genügen. Das Wissenschaftliche Direktorium stellt sicher, dass jedes Dagstuhl-Seminar durch ein starkes Organistorenteam unterstützt wird, ein für die Informatik-Community relevantes Thema anspricht, ein kohärentes und gut strukturiertes wissenschaftliches Programm präsentiert und eine Gruppe von geeigneten Teilnehmerinnen und Teilnehmern zusammenbringt, deren kollektive Fachkenntnis einen bedeutenden Durchbruch in dem betreffenden Forschungsfeld ermöglichen kann.

Im Begutachtungsprozess liegt der Fokus auf einer ausgeglichenen Repräsentation wissenschaftlicher Gemeinden, geographischer Regionen und besonders auf der Miteinbeziehung junger und weiblicher Wissenschaftler. Sowohl das Wissenschaftliche Direktorium als auch der wissenschaftliche Stab von Schloss Dagstuhl sind bemüht, unterrepräsentierte Gruppen in das Seminarprogramm aufzunehmen, indem die Organisatoren in einigen Fällen aufgefordert werden, die vorläufige Teilnehmerliste zugunsten größerer Ausgeglichenheit zu ändern.

Nähere Betrachtung des Dagstuhl-Seminarprogramms 2013

2.4

Die internationalen Informatikforscher zeigten 2013 wieder ihr hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 106 Anträgen auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops in den Antragsrunden im Januar und Juni 2013. Etwas mehr als 67 % der eingereichten Anträge wurden genehmigt, worin sich die außerordentliche Qualität der Anträge widerspiegelt. In den vergangenen 6 Jahren variierte die Rate der angenommenen Anträge zwischen 67 % und 77 % (siehe Fig. 2.2).

Wie bereits 2012 wurden auch 2013 mehr kleine Seminare als große Seminare genehmigt. Dies liegt daran, dass seit Sommer 2012 im Allgemeinen jede Woche ein großes und ein kleines Seminar parallel geplant werden. Um die

Proposal Submission and Review Process

Schloss Dagstuhl maintains the high quality of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop series by identifying the proposals that are most likely to result in successful seminars. The center solicits topics for new seminars and workshops twice a year from leading researchers worldwide, who submit their proposals together with a list of potential scientists to be invited. The proposals and suggested invitee lists are then reviewed by multiple members of a specifically-appointed academic board, the Dagstuhl Scientific Directorate.

To be successful, a proposal for a Dagstuhl Seminar or Dagstuhl Perspectives Workshop must meet specific quality criteria defined by the Directorate. It is the task of the Directorate to ensure that every accepted seminar is backed by a strong team of organizers, addresses a topic of relevance to the computer science community, presents a coherent and well-structured scientific agenda, and brings together the right group of participants whose collective expertise can lead to a significant breakthrough in the area to be addressed.

The review process places a strong emphasis on a balanced representation of communities, geographical regions, and especially on the inclusion of junior and female researchers. Both the Schloss Dagstuhl Scientific Directorate and the Dagstuhl scientific support staff proactively seek to include underrepresented groups in the Dagstuhl Seminar program by encouraging and sometimes requiring organizers to modify their proposed invitation lists in order to achieve a better balance.

A Closer Look at the Dagstuhl Seminar Program in 2013

The international scientific community expressed a lively interest in organizing seminars and workshops at Schloss Dagstuhl in 2013, submitting a record 106 proposals for new Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2013 and June 2013 submission rounds. The quality of the proposals was excellent, resulting in a 67 % acceptance rate by Dagstuhl's Scientific Directorate. Since 2008, proposal acceptance rates have tended to range between 67 % and 77 % (see Fig. 2.2).

Again in 2013, small seminars predominated over larger seminars in terms of accepted seminars. The reason for this stems from the center's expanded seminar schedule featuring two parallel seminars per week – one small, one large – since summer of 2012. In order to allocate its free (and generally small) short-term program slots to Dagstuhl

noch kurzfristig verfügbaren Termine flexibel zu belegen, fördert Schloss Dagstuhl besonders kleine Seminare. Dadurch wurden 2012 und 2013 mehr solcher Seminare in das wissenschaftliche Programm aufgenommen als es in den Jahren zuvor der Fall war.

Von den 72 in 2013 neu genehmigten Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops waren 38 klein (vgl. Fig. 2.3). Viele dieser Seminare werden Ende 2013 oder Anfang 2014 ausgerichtet, so dass weiterhin die durchschnittliche Vorlaufzeit der genehmigten Seminare und Workshops verkürzt werden konnte.

Die Anzahl der Wochen, in denen zwei Seminare parallel stattfanden, ist 2013 erneut stark gestiegen. Insgesamt fanden im Berichtsjahr in 30 von 48 Wochen Seminare parallel statt, was 62,5 % der verfügbaren Wochen entspricht. Bereits ein starker Kontrast zu 35 % in 2012 und eine dramatische Erhöhung verglichen mit den 12 % der verfügbaren Wochen in 2011.

Mehr Wochen mit parallelen Seminaren bedeuten ebenfalls mehr Seminare insgesamt: In 2013 waren 75 von 113 Veranstaltungen – und somit etwa 2/3 aller Veranstaltungen – entweder Dagstuhl-Seminare oder Dagstuhl-Perspektiven-Workshops (vgl. Fig. 13.5 im Kapitel 13). Noch deutlicher ist der Fokus an der Anzahl der Teilnehmer abzulesen: 2 639 der 3 503 Gästen haben 2013 Seminare oder Workshops besucht (vgl. 2.4). Dies entspricht einem Anteil von etwa 75 %.

Zusammenfassend gesagt wurde auch in 2013 der Trend zu mehr Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops, verglichen mit den anderen Veranstaltungstypen, fortgesetzt: Es gab mehr Seminare, mehr Wochen mit parallelen Seminaren, ein wenig mehr kurze Seminare, die genehmigt wurden, und einen höheren Anteil von Teilnehmern, die die Seminare und Workshops besuchten.

Seminars, the center actively promoted smaller seminars and included more of these in its 2012–2013 scientific program than in previous years.

Of 72 new Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted by the Directorate in 2013, 38 were small seminars (see Fig. 2.3). Many of these small seminars could be scheduled already in late 2013 or early 2014 – thereby reducing the overall lead time of the accepted seminars and workshops.

In keeping with this trend, the proportion of program space reserved for parallel seminars rose sharply. In total, 62.5 % of Dagstuhl’s program space featured parallel seminars during the year under review – a striking contrast even with respect to 2012 (35 %) and a dramatic one with respect to 2011 (only 12 %). Thirty out of 48 weeks had parallel seminars in 2013.

More parallel seminars also meant more seminars and workshops overall: approximately two-thirds of the center’s 2013 scientific program (75 out of 113 total events, cf. Fig. 13.5 in Chapter 13) was devoted to Dagstuhl Seminars and Dagstuhl Perspectives Workshop, and over 75 % of those who visited our center in 2013 (2,639 guests out of 3,503) did so in order to participate in one of our seminars or workshops (see Fig. 2.4).

In summary, the expansion of Dagstuhl’s seminar and workshop program in 2013, even more than in 2012, meant more Dagstuhl Seminars overall, more weeks with parallel seminars, a slight predominance of smaller seminars over larger ones with respect to accepted proposals, and a greater proportion of guests visited Dagstuhl was in order to attend a Dagstuhl Seminar or Dagstuhl Perspectives Workshop.

Year	Proposals		Accepted		Rejected	
	#	#	%	#	%	
2008	83	60	72.3	23	27.7	
2009	95	68	71.6	27	28.4	
2010	94	65	69.1	29	30.9	
2011	80	54	67.5	26	32.5	
2012	90	69	76.7	21	23.3	
2013	107	72	67.3	35	32.7	

Fig. 2.2

Dagstuhl Seminar proposals and acceptance rates.

	small	large
short	13	1
long	25	33

Fig. 2.3

Small vs. large and short vs. long Dagstuhl Seminars and Dagstuhl Perspectives Workshops approved in 2013. Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

Angaben zu Teilnehmern und Organisatoren

2.5

Participant and Organizer Data

Die Teilnehmer der Seminare kommen aus aller Welt und eine erhebliche Anzahl besucht Dagstuhl mehrmals. Nichtsdestotrotz zieht das Zentrum jedes Jahr auch neue Gesichter an, was den ständigen Wandel in der internationalen Informatikforschung und auch auf Schloss Dagstuhl widerspiegelt. Wie bereits 2012 besuchte die Mehrheit aller Teilnehmer Schloss Dagstuhl zum ersten Mal. So nahmen 2013 etwa 55 % der Gäste, 1 444 von 2 639, das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während fast 22 % der Teilnehmer nur an einem vorherigen Seminar teilgenommen hatten (siehe Fig. 2.5).

Ein solider Anteil der Gäste besteht aus jungen Wissenschaftlern am Anfang ihrer Karriere, die unter Umständen ein Leben lang von dem Dagstuhl-Erlebnis zehren. Etwa 29 % der Gäste der Seminare und Workshops in 2013, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stufen sich selbst als Nachwuchswissenschaftler ein, 53 % als erfahrene Forscher (siehe Fig. 2.6). Diese ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern ist im Laufe der Jahre relativ konstant geblieben, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlern und ihren erfahrenen Kollegen zeigt.

Mit 75 % war der Anteil an Gästen aus dem Ausland 2013 wieder sehr hoch, sogar höher als in 2012. Das Diagramm in Fig. 2.7 zeigt die regionale Verteilung der Gäste bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops 2013. Eine detaillierte Aufstellung der fast 50 Herkunftsländer aller Teilnehmer bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops und anderer Veranstaltungen kann Fig. 13.1a in Kapitel 13 entnommen werden.

Ziel der Programmplanung ist es, ein ausgewogenes wissenschaftliches Seminarprogramm anzubieten. Antragsteller werden angehalten, qualifizierte weibliche Kollegen in das Organisatorenteam und die Gästeliste aufzunehmen. In 2013 waren fast die Hälfte aller Organisatorenteam des wissenschaftlichen Programms hinsichtlich des

Participants in Dagstuhl Seminars come from all over the world and a significant number of them choose to repeat the experience. Nevertheless, we see many fresh new faces every year, reflecting the changing face of informatics research across the globe and at Schloss Dagstuhl itself. As in 2012, the majority (1,444 of 2,639, or nearly 55 %) of Dagstuhl Seminar participants in 2013 were first-time visitors to Dagstuhl, followed by 572 participants (nearly 22 %) who had already attended one previous seminar at the center (see Figure 2.5).

A healthy number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 29 % of 2013 seminar and workshop survey respondents self-classified as junior and nearly 53 % as senior (see Fig. 2.6). This proportion of junior to senior researchers has remained relatively constant over the years, reflecting the center’s determined effort to maintain the “Dagstuhl connection” between brilliant junior scientists and their senior colleagues.


At over 75 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2013 – even higher than in 2012. The chart in Fig. 2.7 shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2013. For a detailed breakdown of the nearly 50 countries of origin for all participants in Dagstuhl Seminars, Dagstuhl Perspectives Workshops, and other events at our center, please refer to Fig. 13.1a in Chapter 13.

The Dagstuhl Seminar program strives to promote a balanced scientific program that proactively encourages applicants to include qualified female colleagues in their organizer teams and invitee lists. In 2013, nearly half of all organizer teams in our scientific program were mixed with respect to gender, a proportion that has remained relatively unchanged in comparison to most previous years since 2008 (see Fig. 2.8). The percentage of female seminar participants was also high both in total and relative terms, at 15.2 % (see Fig. 2.9).

Year	DS		PW		GI		EDU		OE		Total #
	#	%	#	%	#	%	#	%	#	%	
2008	1622	55.7	179	6.1	32	1.1	166	5.7	912	31.3	2911
2009	1983	65.9	185	6.1	26	0.9	131	4.4	686	22.8	3011
2010	1950	64.7	103	3.4	25	0.8	192	6.4	743	24.7	3013
2011	1894	70.2	64	2.4	0	0.0	103	3.8	637	23.6	2698
2012	2226	64.4	120	3.5	48	1.4	144	4.2	916	26.5	3454
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503

Fig. 2.4

Number of participants by event type and year. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl-Seminar, EDU = educational event, OE = other event.



(a) Graphical distribution

Previous visits	Participants	
	#	%
2008–2013		
0	1444	54.7
1	572	21.7
2	316	12
3	140	5.3
4	68	2.6
5	57	2.2
≥ 6	42	1.6

(b) Distribution of previous visits

Fig. 2.5
Previous visits of Dagstuhl participants in 2013. The number of Dagstuhl Seminars or Dagstuhl Perspectives Workshops our 2013 participants attended between 2008 and 2013.

Geschlechts gemischt, wobei dieses Verhältnis im Vergleich zu den meisten Vorjahren seit 2008 relativ konstant geblieben ist (siehe Fig. 2.8). Der prozentuale Anteil an weiblichen Seminarteilnehmern war mit 15,2% wieder erfreulich hoch (siehe Fig. 2.9).

Themen und Forschungsgebiete

2.6

Topics and Research Areas


Bis auf vereinzelte Ausnahmen stammen die Themen und thematischen Schwerpunkte der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von den Organisatoren, die diese beantragen und leiten. Der breite Zuspruch, den Schloss Dagstuhl erhält, zeigt sich auch in der thematisch großen Bandbreite von Seminaren, die beantragt und genehmigt wurden. Im Folgenden sind einige thematische Schwerpunkte und beispielhaft einige Seminare in diesen aufgeführt. Weder die Aufzählung der Themen noch die der Seminare und Workshops ist erschöpfend. Kapitel 4 bietet mit den Kurzzusammenfassungen der Seminar- und Workshopergebnisse einen vollständigen Überblick über das wissenschaftliche Programm 2013 von Schloss Dagstuhl.

Seminare mit dem Thema der automatischen Verifikation von Programmen, Systemen und Protokollen wurden 2013 im Vergleich zu den Vorjahren vermehrt beantragt und fanden vermehrt statt (13091 – *Analysis, Test and Verification in The Presence of Variability*; 13141 – *Formal Verification of Distributed Algorithms*; 13181 – *VaToMAS – Verification and Testing of Multi-Agent Systems*). Hierbei wurden sowohl die praktischen Aspekte als auch die formalen Grundlagen angesprochen. Methodisch von einer anderen Seite her kommend wurden aber auch robuste Systeme, Fehlererkennung und Korrektur diskutiert (13022 – *Engineering Resilient Systems: Models, Methods and Tools*; 13061 – *Fault Prediction, Localization, and Repair*). Beide Ansätze zielen aber dennoch darauf, Systeme zu entwerfen, um den Einfluss von Fehlern auszuschalten.

With few exceptions, Dagstuhl Seminar and Dagstuhl Perspectives Workshop topics and their topical focus are given by the organizer teams who propose and lead these seminars. The broad spectrum of seminar topics proposed and approved attest to the wide popularity of Schloss Dagstuhl throughout the computer science community. The following gives some topical focal points and seminars representative of each topic in 2013. Neither the list of focal points nor the list of seminars and workshops is exhaustive. The seminar summaries in Chapter 4 provide a full overview of Schloss Dagstuhl's scientific program during the year under review.

In comparison with previous years, the center hosted and received as proposal submissions more seminars in the area of automatic verification of programs, systems and protocols in 2013 than it had in previous years (e.g. 13091 – *Analysis, Test and Verification in The Presence of Variability*; 13141 – *Formal Verification of Distributed Algorithms*; 13181 – *VaToMAS – Verification and Testing of Multi-Agent Systems*). In these seminars, practical topics and theoretical foundations were discussed. Another group of seminars focused on a totally different approach – that of resilient systems, fault detection and correction (e.g. 13022 – *Engineering Resilient Systems: Models, Methods and Tools*; 13061 – *Fault Prediction, Localization, and Repair*). Although in different ways, both approaches seek to achieve systems that can run efficiently unaffected by faults.

Dagstuhl Seminars on security focused mainly on special application areas in 2013 (13062 – *Decentralized*




Year	Junior (%)	Senior (%)	Neither (%)
2008	23.8	55.5	20.7
2009	25.2	52.9	21.9
2010	28.9	49.7	21.4
2011	27.9	51.2	20.9
2012	27.6	52.1	20.3
2013	28.9	52.8	18.3

(a) Graphical distribution

(b) Detailed numbers

Fig. 2.6
Self-assigned seniority of Dagstuhl Seminar participants.




Region	Participants (#)	Participants (%)
Europe (w/o Germany)	1170	44
Germany	653	25
North America	576	22
Asia	156	6
Australia	54	2
South America	24	1
Africa	6	0

(a) Graphical distribution

(b) Detailed numbers

Fig. 2.7
Origin of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2013.



Year	Teams #	Organizers #	Mixed Teams		Women	
			#	%	#	%
2008	52	200	27	51.9	31	15.5
2009	60	228	20	33.3	20	8.8
2010	59	233	32	54.2	34	14.6
2011	55	213	27	49.1	31	14.6
2012	64	256	32	50.0	39	15.2
2013	75	282	36	48.0	43	15.2

(a) Graphical distribution

(b) Detailed numbers

Fig. 2.8
Dagstuhl Seminars with mixed-gender organizer teams. About 50% of the seminars have a mixed-gender organizer team.

Year	Participants #	Female Participants	
		#	%
2008	1801	244	13.5
2009	2168	296	13.7
2010	2053	293	14.3
2011	1958	294	15.0
2012	2346	378	16.1
2013	2639	401	15.2

Fig. 2.9
Female participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops by year.

Im Bereich Sicherheit und Datenschutz war ein Fokus auf spezielle Anwendungsszenarien zu erkennen. (13062 – *Decentralized Systems for Privacy Preservation*; 13412 – *Genomic Privacy*).

Wie in den Jahren zuvor fanden 2013 auch wieder mehrere Seminare statt, die der Kerninformatik im Bereich Datenstrukturen, Algorithmen und Komplexität gewidmet waren (13232 – *Indexes and Computation over Compressed Structured Data*; 13042 – *Epidemic Algorithms and Processes: From Theory to Applications*; 13331 – *Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time*). Viele Seminare dieses Themenbereichs begleiten das wissenschaftliche Programm von Schloss Dagstuhl seit den frühesten Anfängen und bieten erfreulicherweise dennoch jedes Jahr wieder eine Plattform für die Diskussion neuer Schwerpunkte und aktueller Erkenntnisse (13101 – *Computational Geometry*; 13111 – *Scheduling*).

Im Bereich der angewandten Informatik ist eine Tendenz zu mehr Diversität in den Themen zu erkennen. Auffallend viele Seminare in diesem Bereich sind interdisziplinär und bringen Forscher verschiedener Fachbereiche auch aus den Randgebieten der Informatik sowie Nicht-Informatiker aus den Kernbereichen der Anwendungen zusammen (13451 – *Computational Audio Analysis*; 13272 – *Computer Science in High Performance Sport – Applications and Implications for Professional Coaching*; 13212 – *Computational Methods Aiding Early-Stage Drug Design*). Eines der ersten Seminare des Jahres diskutierte aus zahlreiche Perspektiven den Einsatz der Informatik in der zivilen Krisenintervention (13041 – *Civilian Crisis Response Models*) und stand damit ganz im Zeichen der großen Natur- und Umweltkatastrophen der letzten Jahre.

Generell umfasste das Seminar-Programm 2013 wie üblich eine breite Palette von Forschungsgebieten, wobei eine Vielzahl von Wissenschaftsgemeinschaften und Disziplinen involviert war.

Systems for Privacy Preservation; 13412 – *Genomic Privacy*).

As always, we had again in 2013 several core informatics seminars in the areas of algorithms, data structures, and complexity (13232 – *Indexes and Computation over Compressed Structured Data*; 13042 – *Epidemic Algorithms and Processes: From Theory to Applications*; 13331 – *Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time*). Many of these seminars form part of long-running series, and some date back right to Schloss Dagstuhl's earliest days. They continue to offer a platform for discussing new focal points and recent developments (13101 – *Computational Geometry*; 13111 – *Scheduling*).

In the area of applied computer sciences, a tendency towards more topical diversity could be observed. There were a striking number of interdisciplinary seminars that brought together researchers from different areas, including those on the periphery of computer science, and included non-computer scientists from the core areas of application discussed (13451 – *Computational Audio Analysis*; 13272 – *Computer Science in High Performance Sport – Applications and Implications for Professional Coaching*; 13212 – *Computational Methods Aiding Early-Stage Drug Design*). One of the first seminars in 2013 (13041 – *Civilian Crisis Response Models*) discussed how computer science could be applied to the field of civil crisis intervention, in view of the large natural and environmental disasters of the recent years.

In general, Dagstuhl's 2013 seminar program included the usual broad range of research areas, involving a wide variety of communities and disciplines.

Weitere Veranstaltungstypen

2.7

Neben den Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops finden noch weitere Veranstaltungen im Zentrum statt. Zu diesen Veranstaltungen gehören:

- GI-Dagstuhl-Seminare, die den wissenschaftlichen Nachwuchs zu einem bestimmten Thema zusammenführen und in Kooperation mit der GI durchgeführt und von der GI sowie von Dagstuhl gefördert werden
- Sommerschulen, Weiterbildungsveranstaltungen der GI, Lehrerfortbildungen, Ausbildung von jungen Journalisten und Volontären
- Klausurtagungen von Graduiertenkollegs, GI-Fachgruppen und anderen akademischen und industriellen Arbeitsgruppen
- in geringem Umfang internationale Informatik-Fachtagungen
- Forschungsaufenthalte

Further Event Types

In addition to Dagstuhl Seminars and Dagstuhl Perspectives Workshops, Schloss Dagstuhl hosts a number of further events, including:

- GI-Dagstuhl seminars, sponsored by the German Informatics Society (GI) in association with Schloss Dagstuhl, that bring young scholars together to discuss and learn about a specific topic
- summer schools, continuing education courses sponsored by the German Informatics Society (GI), vocational training for teachers and instructors, and educational and training workshops for young journalists and trainees
- departmental conferences of graduate colleges, GI specialist groups and other academic and industrial working groups
- a small number of international informatics conferences
- research stays

Das Angebot, Dagstuhl zu einem wissenschaftlichen Forschungsaufenthalt zu besuchen, wird regelmäßig genutzt. In den meisten Fällen sind es Einzelpersonen, die sich für eine oder mehrere Wochen für intensive Studien nach Dagstuhl in Klausur zurückziehen.

People regularly take advantage of Dagstuhl's offer to use the center for research stays. In most cases these are individuals who wish to use the center as a retreat for several weeks in order to devote themselves to their studies undisturbed.

Qualitätssicherung

2.8

Quality Assurance

Schloss Dagstuhl befragt die Teilnehmer der Dagstuhl-Seminare und der Dagstuhl-Perspektiven-Workshops mit Hilfe eines Fragebogens nach ihrer Zufriedenheit mit inhaltlichen und organisatorischen Aspekten ihres Besuchs. Die Ergebnisse jedes Fragebogens werden im Haus wöchentlich allen Abteilungen zugänglich gemacht, um eine schnelle Reaktion auf Probleme und Wünsche zu erreichen. Gleichzeitig werden die anonymisierten Ergebnisse von inhaltlichen Fragen den Teilnehmern eines Seminars per E-Mail mitgeteilt, typischerweise in der Woche nach ihrem Aufenthalt. So erhalten insbesondere Organisatoren Rückmeldungen über den Verlauf des Seminars und Hinweise für die Organisation von zukünftigen Seminaren. Seit 2013 werden diese statistischen Ergebnisse mit Hilfe von aussagekräftigen Diagrammen aufbereitet und als PDF-Dokumente zur Verfügung gestellt. Der Student Dominik Michels realisierte die Erstellung dieser verbesserten Zusammenfassungen während eines Praktikums im Sommer 2013.

The center conducts surveys of the participants of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop, the questionnaire containing questions about their satisfaction with the content of the event and the organization of their visit. The results of each questionnaire are made available to all of the center's departments every week, thus enabling a quick response to issues and requests. At the same time the anonymized results of the content questions are made available to the seminar participants via e-mail, typically in the week following their stay at the center. This enables the organizers to receive feedback on how the seminar went and tips for organizing future seminars. In 2013, Schloss Dagstuhl began sending the report as a PDF attachment with an enhanced visual layout. Student intern Dominik Michels carried out the improvements.

Fig. 2.10 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2013 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1485 Fragebögen, welche die Meinung von etwa 56 % der 2639 Teilnehmer repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

Fig. 2.10 shows the satisfaction of responding participants in 2013 with regard to selected aspects of their stay. The results were compiled from 1,485 questionnaires, representing the responses of about 56 % of all participants (2,639). These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

Seit 2013 bietet Schloss Dagstuhl allen Organisatoren den direkten Zugriff auf den Status der eingeladenen Gästen bezüglich Zu- oder Absage. Die Webseite mit täglich aktualisierten Daten bietet den Organisatoren einen transparenteren Überblick über die administrative Organisation ihrer Seminare und stieß auf positive Resonanz bei ihnen.

In 2013 Schloss Dagstuhl also began systematically offering all Dagstuhl Seminar organizers a more transparent invitation process by giving them direct access to the status of invitee replies via a dedicated webpage. The page is available 24/7 and has met with very positive feedback from organizers

Auslastung des Zentrums

2.9

Utilization of the Center

Dank der Eröffnung des neuen Gästehauses im Jahr 2012 konnte Schloss Dagstuhl sein Seminarprogramm ausbauen. Dadurch gab es 2013 mit 14097 Übernachtungen insgesamt und 11612 Übernachtungen von Dagstuhl-Seminar und Dagstuhl-Perspektiven-Workshop Gästen mehr Übernachtungen als jemals zuvor. Bezogen auf die Seminar- und Workshopgäste bedeutet dies ein Wachstum von 19 % verglichen mit den Übernachtungen in 2012 und ein Wachstum von sogar 38 % bezogen auf das Jahr 2011. Dies zeigt den Trend zu einer höheren Auslastung, jedoch erwarten wir, dass sie sich im Jahr 2014 stabilisieren wird. Es fanden im Berichtsjahr 113 Veranstaltungen mit insgesamt 3503 Gästen statt. Weitere Details können Kapitel 13 entnommen werden.

Thanks in part to the intensification of the Dagstuhl Seminar program following the opening of the new guest house the previous year, in 2013 Schloss Dagstuhl had more overnight stays (14,097) and more overnight stays in seminars and workshops (11,612) than ever before. The latter represented a 19 % increase over 2012 and a dramatic 38 % increase over 2011, tracing a steady trend towards higher utilization that is expected to stabilize in 2014. The center hosted a total of 113 events with 3,503 guests in 2013. See Chapter 13 for further details.

Weekends were kept free in 2013, as well as two weeks in August and at the end of the year, this time being required for maintenance work to building facilities and administrative work. Apart from a few isolated periods

Die Wochenenden blieben 2013 ebenso unbelegt wie jeweils zwei Wochen im August und am Jahresende. Diese wurden zu Instandhaltungs- und Verwaltungsarbeiten benötigt. Abgesehen von vereinzelt Zeiträumen und einigen bisher noch nicht belegten Plätzen für Gruppen mit bis zu 20 Teilnehmern ist das Zentrum zur Zeit für fünftägige Seminare mit 45 Teilnehmern bis einschließlich Juni 2015 und bis einschließlich Februar 2015 für drei- und fünftägige Seminare mit 30 Teilnehmern ausgebucht.⁸

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2013 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminaren und Veranstaltungen wie Tagungen und Sommerschulen, bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite kann unser Kalender⁹ mit anstehenden Veranstaltungen eingesehen werden, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

and a series of as yet unbooked parallel event slots for groups of up to 20 participants, the center is currently fully booked up through through June of 2015 for 5-day, 45-person seminars and through February of 2015 for 3- and 5-day, 30-person seminars⁸.

A comprehensive listing of all events at Schloss Dagstuhl in 2013, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar⁹ of upcoming events and further information and material on all events past, present and future, e.g. aims and scope, participant list, and concluding report.

⁸ Stand: April, 2014
As of April, 2014

⁹ http://www.dagstuhl.de/no_cache/programm/kalender/




Fig. 2.10 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2013. According to survey results.



Fig. 2.11

The Türkismühle train station, a familiar landmark for many of our guests. Photograph reprinted with permission of 2013 Dagstuhl Seminar participant Robert Kosara, who discussed the experience on his blog. See <http://eagereyes.org/blog/2013/schloss-dagstuhl>.

3 Resonanz *Feedback*

Resonanz von Seminarteilnehmern

3.1

Feedback from Seminar Participants

Schloss Dagstuhl bekommt viel Lob von seinen Gästen, meistens in mündlicher Form, wenn die Gäste nach einer intensiven Seminarwoche das Schloss verlassen. Manche Gäste nehmen sich jedoch auch die Zeit, uns nachträglich zu schreiben und ihre Eindrücke mit uns zu teilen. Im folgenden haben wir mit freundlicher Genehmigung der Autoren einen Auszug aus unserer großen Sammlung an Dankeschön-Nachrichten zusammengestellt.

Schloss Dagstuhl receives a lot of positive feedback, typically verbally when our guests are checking out after an intense seminar. However, many guests take the time to write to us about their impressions. What follows is an excerpt from our large thank-you collection, cited here with the authors' appreciated permission.

Alexander Wolff (Universität Würzburg, DE)

12261 – Putting Data on the Map | Dagstuhl Seminar | <http://www.dagstuhl.de/12261>

In summary, it is our impression that the (56!) participants enjoyed the great scientific atmosphere offered by Schloss Dagstuhl and profited from the scientific program. We are grateful for having had the opportunity to organize this seminar.

Marcus Gallagher (The University of Queensland – Brisbane, AU)

13271 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | <http://www.dagstuhl.de/13271>

I would have essentially given the seminar (and Dagstuhl) maximum ratings for everything on the survey. This was my 4th time visiting Dagstuhl (and this seminar series) and I think it is an exceptional venue and seminar. Everything at Dagstuhl runs to perfection: I can't think of a better research meeting place. It is something that everyone involved should be very proud of. Travelling from Australia to attend a workshop is a significant time and financial commitment. But to attend Dagstuhl (and this seminar) I do this without hesitation. Hopefully I will be back at Dagstuhl again in time for the next seminar in the series!

Colin Snook (University of Southampton – Southampton, UK)

13372 – Integration of Tools for Rigorous Software Construction and Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13372>

I think this was by far the best seminar or workshop I have ever attended. The format: - a smaller number of participants, splitting into sub-groups to work in parallel on problems/case studies, reporting back to the main group. With at most a small number of pre-prepared presentations of really essential background work. This worked really well and ensured that some real work and progress was actually made during the seminar week. As a result of these sessions I will start a new collaboration with Uni Milan to try to adapt our UML-B diagram tools to a new formalism, ASM. I would urge Dagstuhl to encourage this format.

Anna Esposito (The Second University of Naples – Caserta, IT)

13451 – Computational Audio Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13451>

The seminar was a unique and brilliant experience. Let me express my deep appreciation for allowing me to join it, thank to the support. I will take care of disseminate all the results I will get for this unforgettable experience and let you know.

Jianmin Wang (Tsinghua University – Beijing, CN)

13481 – Unleashing Operational Process Mining | Dagstuhl Seminar | <http://www.dagstuhl.de/13481>

As a participant, I would like to thank your excellent organization of this event. It gave me a valuable chance to learn about so much insightful knowledge from so many outstanding participants.

Wil van der Aalst (Eindhoven University of Technology – Eindhoven, NL)13481 – Unleashing Operational Process Mining | Dagstuhl Seminar | <http://www.dagstuhl.de/13481>

The seminar showed that that process mining is a new/exciting research discipline with a huge potential and many interesting challenges. The discussion sessions and presentations illustrated the breadth of process mining as field. It was also nice to see that PM academics are genuinely interested in real-life applications (this is not always the case in other fields) and that the field is maturing (cf. benchmark-related discussions). I noted many people working together on concrete topics, so I'm confident that the seminar triggered new collaborations, joint papers, joint project proposals, etc.

Roly Perera (University of Edinburgh – Edinburgh, UK)13382 – Collaboration and learning through live coding | Dagstuhl Seminar | <http://www.dagstuhl.de/13382>

Guess my main goal when I came here was having my brain melted into new shapes—that definitely happened.

Resonanz unserer Seminarorganisatoren

3.2

Feedback from Seminar Organizers

Der Erfolg von Schloss Dagstuhl hängt im wesentlichen Maße auch von den Seminarorganisatoren ab, die interessante und neue Themen vorschlagen. Wir sind hoch erfreut, dass die Seminarorganisatoren selber, die Angebote und die Umgebung, die wir zur Verfügung stellen, schätzen. Im folgende geben mit freundlicher Genehmigung der Autoren einige der Kommentare unsere Seminarorganisatoren wieder.

The success of Schloss Dagstuhl depends to a large extent on our outstanding seminar organizers, who continually enrich the scientific program with a range of interesting and new topics. We are very glad to be able to provide services and an environment that organizers appreciate. The following comments from organizers are excerpted from the Dagstuhl Report for each of the cited seminars.

Organizers of Dagstuhl Seminar 1310113101 – Computational Geometry | Dagstuhl Seminar | <http://www.dagstuhl.de/13101>

Dagstuhl itself is a great strength of the seminar. Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, we warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! We believe that the lottery created space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar was much „younger“ than in the past, and certainly more „family-friendly.“ Five young children roaming the premises created an even cosier atmosphere than we are used in Dagstuhl. Without decreasing the quality of the seminar, we had a more balanced attendance than in the past. Feedback from both seminar participants and from researchers who were not selected was uniformly positive.

Organizers of Dagstuhl Seminar 1332213322 – The Critical Internet Infrastructure | Dagstuhl Seminar | <http://www.dagstuhl.de/13322>

The editors of this report would like to thank all participants for very fruitful and open-minded discussions! In particular, we thank the operators for sharing practical insights. We gratefully acknowledge the Dagstuhl staff for helping on all administrative coordination, for their patience, and most importantly for providing an extremely inspiring environment.

Organizers of Dagstuhl Seminar 1319213192 – Tree Transducers and Formal Methods | Dagstuhl Seminar | <http://www.dagstuhl.de/13192>

We thank Schloss Dagstuhl for the professional and inspiring atmosphere it provides. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers' needs. For instance, elaborate research discussions in the evening were followed by musical intermezzi of playing piano trios by Beethoven and Mozart, or by table tennis matches and sauna sessions.

Organizers of Dagstuhl Seminar 1323213232 – Indexes and Computation over Compressed Structured Data | Dagstuhl Seminar | <http://www.dagstuhl.de/13232>

We thank Schloss Dagstuhl for the professional and inspiring atmosphere. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers' needs. For instance, elaborate research discussions in the evening were followed by local wine tasting or by heated sauna sessions.

Organizers of Dagstuhl Seminar 1345113451 – Optimality and tight results in parameterized complexity | Dagstuhl Seminar | <http://www.dagstuhl.de/13451>

Beside the scientific aspect, the social aspect of our seminar was just as important. We had an interdisciplinary, international, and very interactive group of researchers, consisting of leaders and future leaders in our field. Most of our participants visited Dagstuhl for the first time and enthusiastically praised the open and inspiring atmosphere. The group dynamics were excellent with many personal exchanges and common activities. Some scientists mentioned their appreciation of having the opportunity for prolonged discussions with researchers from neighboring research fields—something which is often impossible during conference-like events. In conclusion, our expectations of the seminar were not only met but exceeded, in particular with respect to networking and community building. Last but not least, we heartily thank the Dagstuhl board for allowing us to organize this seminar, the Dagstuhl office for their great support in the organization process, and the entire Dagstuhl staff for their excellent services during the seminar.

Organizers of Dagstuhl Seminar 1338113381 – Algorithms and Scheduling Techniques for Exascale Systems | Dagstuhl Seminar | <http://www.dagstuhl.de/13381>

The feedback provided by the participants show that the goals of the seminar, namely to circulate new ideas and create new collaborations, were met to a large extent.

The organizers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive event.

Organizers of Dagstuhl Seminar 1339113391 – Algorithm Engineering | Dagstuhl Seminar | <http://www.dagstuhl.de/13391>

Schloss Dagstuhl and its staff provided a very convenient and stimulating environment. The seminar participants appreciated the cordial atmosphere which improved mutual understanding and inspiration. The organizers of this seminar wish to thank all those who helped make the workshop a fruitful research experience.

Organizers of Dagstuhl Seminar 1341213412 – Genomic Privacy | Dagstuhl Seminar | <http://www.dagstuhl.de/13412>

The organizers, together with the participants, agreed that this problem should be addressed in a sequel Dagstuhl-seminar. Hence, they set up a future work agenda in order to organize again such a fruitful gathering. We thank Schloss Dagstuhl for the professional and inspiring atmosphere it provides. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers' needs.

Öffentliche Resonanz im Web

3.3

Feedback in Social Media

3

Mehr und mehr Gäste nutzen die Möglichkeiten des Webs wie Blogs, etc., über ihre positiven Erfahrungen in Dagstuhl zu berichten. Wir geben hier einige Referenzen.

More and more of our guests are using social media such as blogs, Twitter, etc. to share their positive experiences of Dagstuhl with others. Below are some selected excerpts.

David Griffiths (FoAM – Kernow, UK)

13382 – Collaboration and learning through live coding | Dagstuhl Seminar | <http://www.pawfal.org/dave/blog/2013/09/dagstuhl>

[...] Our seminar was called ‘Collaboration and learning through live coding’, organised by Alan Blackwell, Alex McLean, James Noble and Julian Rohrhuber and included people from the fields of Software Engineering, Computer Science Education as well as plenty of practising livecoders and multidisciplinary researchers. [...] So Dagstuhl’s music room was immediately useful in providing a more ‘normal’ livecoding situation. It was of course more stressful than usual, knowing that you were being critically appraised in this way by world experts in related fields! However it paid off hugely as we had some wonderful interpretations from these different viewpoints. [...]

Mark Guzdial (Georgia Institute of Technology – Atlanta, US)

13382 – Collaboration and learning through live coding | Dagstuhl Seminar | <https://twitter.com/guzdial/status/380041095025266688>

Jam session at #dagstuhl: One Live-Coder, one guitarist, and one pianist.

Robert Kosara (Tableau Software – Seattle, US)

13201 – Information Visualization - Towards Multivariate Network Visualization | Dagstuhl Seminar | <http://eagereyes.org/blog/2013/schloss-dagstuhl>

For many computer science researchers, the name Dagstuhl rings a bell. Anybody who has been there has fond memories of interesting talks, great conversations, and lots of social interaction (lubricated by the abundantly available wine and beer). But what is Dagstuhl? [...]

Mark Guzdial (Georgia Institute of Technology – Atlanta, US)

13382 – Collaboration and learning through live coding | Dagstuhl Seminar | <http://cacm.acm.org/blogs/blog-cacm/168153>

I spent last week at the remarkable Schloss-Dagstuhl in a seminar [...] It was a terrific event that has me thinking about a whole new set of research questions[...] Most of the attendees were live coders, but there were a number of us others who helped explore the boundary disciplines for live coding [...] I was there to connect live coding to computing education. I learned the connections from the seminar – I hadn’t really seen them before I got there. (*Content reprinted with kind permission from the ACM*).

Resonanz im Fragebogen

3.4

Seminar Survey Feedback

Jeder Seminarteilnehmer erhält von uns einen Fragebogen zur Evaluation der Seminare. Durch dieser anonymen Befragung erhalten wir ebenfalls eine menge positiver Kommentar. Im folgenden zitieren wir hier einige von diesen.

Every seminar participant has the opportunity to fill out a questionnaire about the seminar for evaluation purposes. Below are some excerpts from the many positive comments we received through this anonymous survey in 2013.

13241 – Virtual Realities | Dagstuhl Seminar | <http://www.dagstuhl.de/13241>

Best: friendly, informal atmosphere. Worst: I can’t come here very often.

13072 – Mechanisms of Ongoing Development in Cognitive Robotics | Dagstuhl Seminar | <http://www.dagstuhl.de/13072>

The best part was the small number of highly relevant people that enabled a lot of very productive discussions. I especially liked the relaxed and somewhat secluded nature of the meeting, which made interaction and focus much easier (and more fun!).

13272 – Computer Science in High Performance Sport - Applications and Implications for Professional Coaching | Dagstuhl Seminar | <http://www.dagstuhl.de/13272>

The best aspect of the seminar is the informal occasion to discuss openly the works and experiences with other colleagues and not to worry about the time.

13271 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | <http://www.dagstuhl.de/13271>

Traditionally, each time, it is a great opportunity for spending a week 100% in science. More specifically for this time, the best is a joint work started with a Japanese and two UK researchers. The random choice of seats at lunch is just great for meeting everyone.[...]

13312 – "My Life, Shared"– Trust and Privacy in the Age of Ubiquitous Experience Sharing | Dagstuhl Seminar | <http://www.dagstuhl.de/13312>

The seminar brought together an excellent group of people from different backgrounds who were able to interact in ways that would not have been possible otherwise.

13352 – Interaction with Information for Visual Reasoning | Dagstuhl Seminar | <http://www.dagstuhl.de/13352>

The interaction and sharing of ideas across disciplines has been fabulous and inspirational.

13372 – Integration of Tools for Rigorous Software Construction and Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13372>

In my opinion the best part was the open and constructive atmosphere allowing for intensive and fruitful discussions, a milestone for the communities to further grow together and develop a shared vision for future collaborations on integration of Formal Methods at multiple levels (methodologies, tool development and modelling approaches), effectively taking advantage of the broader landscape of methods and tools by viewing them as complementary rather than as exclusive choices.

13312 – "My Life, Shared"– Trust and Privacy in the Age of Ubiquitous Experience Sharing | Dagstuhl Seminar | <http://www.dagstuhl.de/13312>

This is one of my best academic experiences ever, especially given the relaxed and open atmosphere to share ideas on similar research interests. Really unique!

13451 – Computational Audio Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13451>

A lot room for social events and discussions.
Extremely relaxed and open atmosphere.

13471 – Synchronous Programming | Dagstuhl Seminar | <http://www.dagstuhl.de/13471>

In computer science, probably the best place to work and spend a week in interesting discussions.

13502 – Approaches and Applications of Inductive Programming | Dagstuhl Seminar | <http://www.dagstuhl.de/13502>

The best thing about the seminar was that it was small and focused, leading to high quality discussions about specific technical approaches. The worst thing about the seminar was its brevity.

13081 – Consistency In Distributed Systems | Dagstuhl Seminar | <http://www.dagstuhl.de/13081>

A big plus for Dagstuhl is the integration of industry and academia in a more personal and interactive setting. Keep up the good work. I appreciate the ready availability of wine, cheese, coffee, etc., and the meeting places (the rooms) which seem to me to work very well.

13452 – Proxemics in Human-Computer Interaction | Dagstuhl Seminar | <http://www.dagstuhl.de/13452>

Very nice composition of the group. People who have worked in relaxed areas of the field have been carefully selected and invited. All participants were very well prepared for the seminar. Everyone was really openly discussing the challenges and also the issues of the research field.

13481 – Unleashing Operational Process Mining | Dagstuhl Seminar | <http://www.dagstuhl.de/13481>

I do both scientific research and also strive to have impact on products. This seminar had an ideal mix of both.

13252 – Interoperation in Complex Information Ecosystems | Dagstuhl Seminar | <http://www.dagstuhl.de/13252>

Dagstuhl staff has lots of experience. Just keep level, no need for more!

13372 – Integration of Tools for Rigorous Software Construction and Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13372>

The seminar program was not determined in all details before the seminar but left room for dynamic adjustments (made at the end of each seminar day for the next day) in response to feedback from the participants, scientific progress toward the defined goals and practical needs coming out of each individual day. This approach proved very effective and more flexible than sticking to a relatively inflexible program and schedule decided prior to the seminar.

13452 – Proxemics in Human-Computer Interaction | Dagstuhl Seminar | <http://www.dagstuhl.de/13452>

The seminar was extremely well organized. Ahead of the seminar emails were asking the participants to prepare short presentations. The organizers set a very clear frame for the full week. It was a very good mix of short talks and breakout sessions. The timing was perfect and we had a good balance of the sessions and the breaks. It was a very inspiring seminar with a great atmosphere for research discussion but also on research and life in general.

13041 – Civilian Crisis Response Models | Dagstuhl Seminar | <http://www.dagstuhl.de/13041>

It was my first experience. I would like to congratulate the Dagstuhl team for your high quality standards and thank you for the friendly atmosphere.

13051 – Software Certification: Methods and Tools | Dagstuhl Seminar | <http://www.dagstuhl.de/13051>

I have thoroughly enjoyed the seminar. Were getting here not so expensive, I would be very interested in coming again, but the travel costs mean that it is unlikely that I will be able to attend another seminar any time soon. If the money was available within my organization to enable me to attend, I would do so.

13082 – Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices | Dagstuhl Seminar | <http://www.dagstuhl.de/13082>

I hope that this is not a “once in a lifetime“ event. But if it is. I have had a better experience than most others.

13091 – Analysis, Test and Verification in The Presence of Variability | Dagstuhl Seminar | <http://www.dagstuhl.de/13091>

The Dagstuhl concept is unparalleled, and it performs an invaluable service to the CS research community. The CS community can be very happy to have such a venue. The value Dagstuhl has for researchers can be gauged from the fact that people come here for a whole week without getting any publications out of it.

13111 – Scheduling | Dagstuhl Seminar | <http://www.dagstuhl.de/13111>

To me, visits to Dagstuhl have become one of the main venues for communicating with other researchers in my field, in spite of geographical distance (I live in the US), and I know that it plays a similar role for some of my colleagues. Most conferences don't provide enough free time for interaction, other than attending the talks. I actually don't think any improvement is needed; I do hope that Dagstuhl will continue its operation in its present form.

13151 – Drawing Graphs and Maps with Curves | Dagstuhl Seminar | <http://www.dagstuhl.de/13151>

Dagstuhl is a unique place to have a seminar. Every person in the staff, at each position, is doing such a wonderful job. As usual... Thanks!

13182 – Meta-Modeling Model-Based Engineering Tools | Dagstuhl Seminar | <http://www.dagstuhl.de/13182>

I always look forward to Dagstuhl seminars as the venue is outstanding, promoting communication and collaboration. I always gain a lot from interacting with other researchers. The quality of organisation varies with the organisers but even with suboptimal organisers, so far a Dagstuhl event has always been worth attending.

13182 – Meta-Modeling Model-Based Engineering Tools | Dagstuhl Seminar | <http://www.dagstuhl.de/13182>

Thank you for a wonderful experience! Please keep Dagstuhl as a place for fostering international collaborations.

13251 – Parallel Data Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13251>

Excellent forum! I wish the USA had some venue like this.

13251 – Parallel Data Analysis | Dagstuhl Seminar | <http://www.dagstuhl.de/13251>

It was my first time at Dagstuhl and I enjoyed it a lot. Thanks to the people I met and the ideas we exchanged, I hope my research will improve, and I already better understand some aspects of my field thanks to meeting with industrials. I'm grateful to seminar organizers and Dagstuhl staff for all of that!

13271 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | <http://www.dagstuhl.de/13271>

These seminars are a true inspiration for work in the field. Having one of these seminars in every two years would be highly desirable.

13271 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | <http://www.dagstuhl.de/13271>

Facilities at Dagstuhl are perfect really. The only improvement that I can think of is only indirectly related. If the departure of the buses from St. Wendel on Sundays could be coordinated with the arrival of trains from Frankfurt that'd be great.

13311 – Duality in Computer Science | Dagstuhl Seminar | <http://www.dagstuhl.de/13311>

Dagstuhl is simply wonderful, thanks for making it work in such a splendid way!

13321 – Reinforcement Learning | Dagstuhl Seminar | <http://www.dagstuhl.de/13321>

Dagstuhl is a very special place. I have learned important things at every meeting.

13321 – Reinforcement Learning | Dagstuhl Seminar | <http://www.dagstuhl.de/13321>

We had an excellent seminar. Congratulations and many thanks to the organizers, as well as the Dagstuhl staff and organization.

13441 – Evaluation Methodologies in Information Retrieval | Dagstuhl Seminar | <http://www.dagstuhl.de/13441>

I'm very impressed with the seminar and Dagstuhl so I only have positive things to say. Keep up the good work!

13441 – Evaluation Methodologies in Information Retrieval | Dagstuhl Seminar | <http://www.dagstuhl.de/13441>

i very much enjoyed my stay. i'm extremely impressed with the venue and concept of dagstuhl. i wish we had a comp. sci retreat in the UK like this, where people can meet, or individuals can retreat to work on proposals/books/grants/etc. please remain.

13452 – Proxemics in Human-Computer Interaction | Dagstuhl Seminar | <http://www.dagstuhl.de/13452>

Day care for the kids is and outstanding support and the support for room, food, day care and everything for the kids was so great, so caring, so friendly that my kids will really miss it.

13471 – Synchronous Programming | Dagstuhl Seminar | <http://www.dagstuhl.de/13471>

Dagstuhl is a great place, truly unique world-wide. It has established itself as an important part of Germany's visibility and reputation for research in Computer Science.

13471 – Synchronous Programming | Dagstuhl Seminar | <http://www.dagstuhl.de/13471>

Such a center is unfortunately unique and lacks in France, for example. Thanks.

13402 – Physical-Cyber-Social Computing | Dagstuhl Seminar | <http://www.dagstuhl.de/13402>

Given the trends of workshops and conferences, I find Dagstuhl even more appealing and refreshing – esp for the quality of discussions leading to new ideas rather than just information sharing.

Resonanz zu Dagstuhl Publishing

3.5

Im Prozess der Veröffentlichung von Konferenz-Proceedings, Zeitschriften-Artikeln und Büchern stehen wir in engem Kontakt mit den Herausgebern und Autoren.

Feedback on Dagstuhl Publishing

We are in close contact with editors and authors as part of the publishing procedures for conference proceedings, journal articles, and books.

Anil Seth (Indian Institute of Technology Kanpur, IN) and Nisheeth K. Vishnoi (Microsoft Research – Bangalore, IN)

LIPICs, Vol. 24, FSTTCS'13 – Preface | <http://dx.doi.org/10.4230/LIPICs.FSTTCS.2013.i>

Finally, we thank the Dagstuhl LIPICs staff for their coordination in production of this proceedings, particularly Marc Herbstritt who was very prompt and helpful in answering our questions.

Femke van Raamsdonk (VU University Amsterdam, NL)

LIPICs, Vol. 21, RTA'13 – Preface | <http://dx.doi.org/10.4230/LIPICs.RTA.2013.i>

I wish to thank the editorial board of LIPICs for agreeing to publish these proceedings, and the team of the LIPICs editorial office for their help in the preparation of these proceedings.

Simon M. Lucas (University of Essex, GB), Michael Mateas (University of California – Santa Cruz, US), Mike Preuss (Universität Münster, DE), Pieter Spronck (Tilburg University, NL), Julian Togelius (IT University of Copenhagen, DK)

Dagstuhl Follow-Ups, Vol. 6, Artificial and Computational Intelligence in Games | <http://dx.doi.org/10.4230/DFU.Vol6.12191.i>

The 2012 gathering at Schloss Dagstuhl was deemed a great success by all participants, and it drew a large part of this strength out of the agile and very adaptive style it was held in, with several unforeseen developments in themes and results. This follow-up volume exemplifies the high level of the scientific discussions and the strong focus on scientific progress of the seminar as a whole.

Resonanz zur Bibliographiedatenbank dblp

3.6

Feedback on the dblp Computer Science Bibliography

Die Bibliographiedatenbank dblp wird von zahlreichen internationalen Wissenschaftlern hoch geschätzt und erhält viel Lob. Feedback erhalten wir per Mail, durch Gespräche mit Forschern vor Ort in Dagstuhl, oder über einem Jahr auch durch anonyme Feedback-Fragebögen von Dagstuhl-Seminarteilnehmern.

The dblp computer science bibliography is internationally well known and appreciated. We receive a lot of feedback via mail, through discussions with researchers at Schloss Dagstuhl, and since late 2012 also via anonymous survey feedback from Dagstuhl seminar participants. Below are some excerpts.

Guillaume Cabanac (University of Toulouse, FR)

email feedback | dblp | <http://dblp.dagstuhl.de/>

Many thanks for providing such a brilliant service as the DBLP. I use it daily for my research, and had several papers published that make use of the DBLP in XML format.

anonymous survey feedback | dblp | <http://dblp.dagstuhl.de/>

DBLP is an indispensable resource for CS researchers.
Many thanks for this great service to the community

anonymous survey feedback | dblp | <http://dblp.dagstuhl.de/>

DBLP is most useful for me. I use it regularly preparing bibliography for articles.

anonymous survey feedback | dblp | <http://dblp.dagstuhl.de/>

DBLP ist ein fantastisches und wichtiges Projekt fuer Forschende in Informatik! Ich moechte mich bei den Initiatoren und Betreibern ausdruecklich fuer ihre Arbeit und ihr Engagement bedanken!

anonymous survey feedback | dblp | <http://dblp.dagstuhl.de/>

Die DBLP ist eine einzigartige Quelle und gehört zum täglichen Arbeiten wie Papier und Bleistift

4 Die Seminare in 2013

The 2013 Seminars

<p>Data Structures, Algorithms, Complexity</p> <ul style="list-style-type: none"> ■ Computational Counting (13031) ■ Epidemic Algorithms and Processes: From Theory to Applications (13042) ■ Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices (13082) ■ Computational Geometry (13101) ■ Scheduling (13111) ■ Bidimensional Structures: Algorithms, Combinatorics and Logic (13121) ■ Theory of Evolutionary Algorithms (13271) ■ Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time (13331) ■ Coding Theory (13351) ■ Algorithm Engineering (13391) ■ Algorithms for Optimization Problems in Planar Graphs (13421) 	<p>Verification, Logic, Formal Methods, Semantics</p> <ul style="list-style-type: none"> ■ Symbolic Methods in Testing (13021) ■ Software Certification: Methods and Tools (13051) ■ Dependence Logic: Theory and Applications (13071) ■ Analysis, Test and Verification in The Presence of Variability (13091) ■ Formal Verification of Distributed Algorithms (13141) ■ Pointer Analysis (13162) ■ VaToMAS – Verification and Testing of Multi-Agent Systems (13181) ■ Tree Transducers and Formal Methods (13192) ■ Duality in Computer Science (13311) ■ Integration of Tools for Rigorous Software Construction and Analysis (13372) ■ Deduction and Arithmetic (13411) ■ Nominal Computation Theory (13422) 	<p>Applications, Interdisciplinary Work</p> <ul style="list-style-type: none"> ■ Civilian Crisis Response Models (13041) ■ Interface of Computation, Game Theory, and Economics (13161) ■ Computational Methods Aiding Early-Stage Drug Design (13212) ■ Computer Science in High Performance Sport – Applications and Implications for Professional Coaching (13272) ■ ICT Strategies for Bridging Biology and Precision Medicine (13342) ■ Collaboration and learning through live coding (13382) ■ Physical-Cyber-Social Computing (13402) ■ Computational Audio Analysis (13451) ■ Proxemics in Human-Computer Interaction (13452) ■ Electronic Markets and Auctions (13461) ■ Forensic Computing (13482) ■ Computational Mass Spectrometry (13491) ■ Social Issues in Computational Transportation Science (13512)
<p>Artificial Intelligence, Computational Linguistics</p> <ul style="list-style-type: none"> ■ Mechanisms of Ongoing Development in Cognitive Robotics (13072) ■ Belief Change and Argumentation in Multi-Agent Scenarios (13231) ■ Reinforcement Learning (13321) ■ Computational Models of Language Meaning in Context (13462) ■ Approaches and Applications of Inductive Programming (13502) 	<p>Software Technology</p> <ul style="list-style-type: none"> ■ Engineering Resilient Systems: Models, Methods and Tools (13022) ■ Fault Prediction, Localization, and Repair (13061) ■ Customizing Service Platforms (13171) ■ Meta-Modeling Model-Based Engineering Tools (13182) ■ Automated Reasoning on Conceptual Schemas (13211) ■ Crowdsourcing: From Theory to Practice and Long-Term Perspectives (13361) ■ Cloud-based Software Crowdsourcing (13362) ■ Software Engineering for Self-Adaptive Systems: Assurances (13511) 	<p>Distributed Computation, Networks, Architecture, Systems</p> <ul style="list-style-type: none"> ■ Multicore Enablement for Embedded and Cyber Physical Systems (13052) ■ Consistency In Distributed Systems (13081) ■ Future Internet (13131) ■ The Critical Internet Infrastructure (13322) ■ Algorithms and Scheduling Techniques for Exascale Systems (13381) ■ Inter-Vehicular Communication – Quo Vadis (13392) ■ Automatic Application Tuning for HPC Architectures (13401) ■ Synchronous Programming (13471) ■ Global Measurement Framework (13472) ■ Geosensor Networks: Bridging Algorithms and Applications (13492)
<p>Cryptography, Security, Privacy</p> <ul style="list-style-type: none"> ■ Decentralized Systems for Privacy Preservation (13062) ■ “My Life, Shared” – Trust and Privacy in the Age of Ubiquitous Experience Sharing (13312) ■ Verifiably Secure Process-Aware Information Systems (13341) ■ Quantum Cryptanalysis (13371) ■ Genomic Privacy (13412) 	<p>Data Bases, Information Retrieval, Data Mining</p> <ul style="list-style-type: none"> ■ Indexes and Computation over Compressed Structured Data (13232) ■ Parallel Data Analysis (13251) ■ Interoperation in Complex Information Ecosystems (13252) ■ Evaluation Methodologies in Information Retrieval (13441) ■ Unleashing Operational Process Mining (13481) 	<p>Geometry, Image Processing, Graphics, Visualization</p> <ul style="list-style-type: none"> ■ Drawing Graphs and Maps with Curves (13151) ■ Information Visualization – Towards Multivariate Network Visualization (13201) ■ Virtual Realities (13241) ■ Interaction with Information for Visual Reasoning (13352) ■ Real-World Visual Computing (13431)

4.1 Symbolic Methods in Testing

Organizers: Thierry Jéron, Margus Veanes, and Burkhart Wolff
Seminar No. 13021

Date: January 6–11, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.1.1
 © Creative Commons BY 3.0 Unported license
 © Thierry Jéron, Margus Veanes, and Burkhart Wolff

Participants: Sébastien Bardin, Axel Belinfante, Nikolaj Bjørner, Jasmin Christian Blanchette, Achim D. Brucker, Lukas A. Brügger, Cristian Cadar, Maria Christakis, Sylvain Conchon, Wilkerson de Lucena Andrade, Catherine Dubois, Juhan Ernits, Abderrahmane Feliachi, Christophe Gaston, Arnaud Gotlieb, Wolfgang Grieskamp, Robert M. Hierons, Thierry Jéron, René Just, Marko Kääramees, Pascale Le Gall, Martin Leucker, Delphine Longuet, Dominique Méry, David Molnar, Brian Nielsen, Grgur Petric Maretic, Frank Rogin, Michel Rueher, Nikolai Tillmann, Jan Tretmans, Jaco van de Pol, Margus Veanes, Luca Vigano, Sabrina von Styp, Hélène Waeselynck, Burkhart Wolff, Fatiha Zaidi



Recent breakthroughs in deductive techniques such as satisfiability modulo theories (SMT), abstract interpretation, model-checking, and interactive theorem proving, have paved the way for new and practically effective techniques in the area of software testing and analysis. It is common to these techniques that statespaces, model-elements, program-fragments or automata are represented symbolically making systems amenable to analysis that have formerly been out of reach. Several research communities apply similar techniques to attack the classical problem of state space explosion by using symbolic representation and symbolic execution: parametrized unit testing, fuzz testing, model-based testing, theoremprover based test case generation techniques, and real-time system testing. Moreover, several areas where symbolic methods are used in testing, are often considered more closely related to verification and end up in conferences specialized on those topics rather than at testing conferences. There is little synergy between the different communities although many of them use similar underlying symbolic techniques.

In the following areas, symbolic analysis techniques have recently had significant impact, both industrially as well as in academia. The following areas capture some topics of interest for the proposed seminar, assuming focus on the use of symbolic techniques in each area: Unit Testing, Symbolic Automata Theory in Testing, Model Based Testing, Fuzz Testing, Security Testing, Real-time System Testing, Theorem-Prover-based Test-Case Generation, Hybrid System Testing, and Mutation Testing.

4.2 Engineering Resilient Systems: Models, Methods and Tools

Organizers: Nicolau Guelfi, Maritta Heisel, Mohamed Kaaniche, Alexander Romanovsky, and Elena Troubitsyna

Seminar No. 13022

Date: January 6–11, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.30

© Creative Commons BY 3.0 Unported license

© Maritta Heisel, Mohamed Kaaniche, Alexander Romanovsky, and Elena Troubitsyna



Participants: Antonia Bertolino, Felicita Di Giandomenico, Giovanna Di Marzo Serugendo, Peter H. Feiler, Stefania Gnesi, Vincenzo Grassi, Denis Hatebur, Maritta Heisel, Mohamed Kaaniche, Linas Laibinis, Paolo Masci, Henry Muccini, Andras Pataricza, Patrizio Pelliccione, Matteo Risoldi, Alexander Romanovsky, Thomas Santen, Rolf Schumacher, Janos Sztipanovits, Anton Tarasyuk, Elena Troubitsyna, Marco Vieira

The Dagstuhl Seminar 13022 – Engineering Resilient Systems: Models, Methods and Tools has brought together prominent researchers from different fields to discuss the problems of engineering resilient systems. The seminar was run in a highly interactive manner. The discussions were centered around the following topics:

- defining resilience
- resilience in modelling languages for requirement analysis and system design
- resilience in implementation languages and frameworks
- verifying resilience using testing, model checking and static analysis
- assessing resilience using probabilist models
- resilience mechanisms at architectural and implementation level

The concept of resilience has been introduced to capture the move towards a greater adaptability and flexibility. However, the notion of resilience is still a subject of debates. The seminar has discussed various proposed definitions and converged to defining resilience as *dependability in presence of changes*.

Over the last decades a remarkable progress has been achieved in engineering of highly dependable systems, i.e., the systems that can be justifiably trusted to provide critical services to a society. However, novel computing paradigms pose new scientific and technological challenges to the dependability field. To deliver critical services in a dependable way, the systems should smoothly adapt to changes. At the seminar, we had a dedicated session discussing the nature of changes. Among the proposed categories were

- evolving user requirements
- changing operating environment
- unforeseen failure modes
- scalability challenge

Resilience is strongly linked with the entire life-cycle of a system. Engineering of resilient systems should empower the systems with capabilities to cope with changes in a predictable way, cater for evolution and ensure robust behavior in spite of faults. These require novel techniques that explicitly address resilience through the entire system development cycle. Our seminar has explored challenges in formal modelling and verification of resilient systems. At the seminar we discussed suitable models for resilience, resilience-explicit development methods and verification techniques enabling both quantitative and qualitative resilience evaluation.

Modelling is the primarily vehicle driving development of resilient systems. However, system modelling area is still highly fragmented. The most acute problems are caused by

- the gap between the requirements and models and
- heterogeneity of models used to represent different aspects of system behaviour

Indeed, over the last few years the problem of poor flow-down of system requirements to software requirements has started to receive a proper attention. The vast majority of development relate the severe design problems with the flawed requirements and misunderstandings about what the software should do. Requirements tend to focus on describing nominal behaviour while omitting or poorly describing off-nominal conditions, safety constraints and fault tolerance mechanisms.

During the seminar we have brainstormed the examples of requirements that would be specific to resilient systems and tried to link them with the modelling techniques.

While developing resilient systems the designers use dedicated models to reason about different (often antagonistic) aspects of system behaviour. Hence, the design space is inherently heterogeneous. On the one hand, specialised models provide the designers with expressive and powerful techniques to analyse

various aspects of system behaviour. On the other hand, it becomes hard to obtain a holistic view on the system characteristics and analyse trade-offs between several potentially conflicting goals, define the mechanisms for adapting to volatile operating conditions and devise appropriate mechanisms for proactive fault tolerance.

We have discussed the advances in formal modelling of resilient systems and in particular proactive fault tolerance and adaptive fault tolerance mechanisms at various frameworks. We have reviewed the advances achieved in the area of formal modelling of resilient systems and brain-stormed the techniques leveraging an integration of various models to facilitate emergence of integrated modelling approaches.

Essentially, any design flow can be seen as a set of well-defined abstraction levels. The design flow should allow the designer to optimize design decision at each level and move freely between abstraction layers. At our seminar we discussed the principles of mapping abstract models onto architectural models and design implementation. We addressed the problem of

achieving architectural plasticity and brain-stormed architectural patterns supporting adaptation as well as mechanisms guaranteeing adequate predictable system reaction on changes. A significant attention has also been paid to the methods and tools for resilience assessment.

■ Open Problems

Engineering resilient systems is a young research area. The participants of the seminar have agreed that often it is hard to distinguish a traditional dependability research from the resilience research. We have converged to the view that the system ability to scale, cope with changes and evolve emphasizes the resilience aspect.

It was also noted that the area of resilience engineering lacks a comprehensive reference guide that would allow the designers of resilient systems understand how various proposed methods and tools can facilitate design of resilient systems. The participants of the seminar has decided to work on such a book.

4.3 Computational Counting

Organizers: Peter Bürgisser, Leslie Ann Goldberg, Mark Jerrum, and Pascal Koiran
Seminar No. 13031

Date: January 13–18, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.47

© Creative Commons BY 3.0 Unported license

© Peter Bürgisser, Leslie Ann Goldberg, Mark Jerrum, and Pascal Koiran



Participants: Markus Bläser, Magnus Bordewich, Irénée Briquel, Peter Bürgisser, Andrei A. Bulatov, Xi Chen, Radu Curticapean, Holger Dell, Arnaud Durand, Martin Dyer, Jo Ellis-Monaghan, Hervé Fournier, Andreas Goebel, Leslie Ann Goldberg, Bruno Grenet, Heng Guo, Miki Hermann, Thore Husfeldt, Mark Jerrum, Marek Karpinski, Pascal Koiran, Meena Mahajan, Johann A. Makowsky, Guillaume Malod, Colin McQuillan, Kitty Meeks, Stefan Mengel, Mike S. Paterson, Natacha Portier, David Richerby, Peter Scheiblechner, Uwe Schöning, Daniel Stefankovic, Yann Strozecki, He Sun, Sebastien Tavenas, Leslie Valiant, Pascal Vontobel, William Whistler, Tyson Williams, Mingji Xia, Tomoyuki Yamakami

■ Introduction

Computational complexity is typically concerned with decision problems, but this is a historical accident, arising from the origins of theoretical computer science within logic. Computing applications, on the other hand, typically involve the computation of numerical quantities. These applications broadly fall into two types: optimisation problems and counting problems. We are interested in the latter, broadly interpreted: computing sums, weighted sums, and integrals including, for example, the expectation of a random variable or the probability of an event. The seminar covered all aspects of computational counting, including applications, algorithmic techniques and complexity. Computational counting offers a coherent set of problems and techniques which is different in flavour from other algorithmic branches of computer science and is less well-studied than its optimisation counterpart.

Specific topics covered by the meeting include

- Techniques for exact counting, including moderately exponential algorithms for intractable problems, fixed parameter tractability, and holographic algorithms and reductions;
- techniques for approximate counting including Markov Chain Monte Carlo (MCMC);
- computational complexity of counting, including complexity in algebraic models; and
- applications, for example to models in statistical physics, and to constraint satisfaction.

The questions addressed include: What algorithmic techniques are effective for exact counting and approximate counting? Do these techniques remain effective in the presence of weights (including negative and complex weights)? What inherent limitations arise from computational complexity? Are there inherent limitations for specific techniques such as MCMC? Our

nominated application areas prompted many of those questions and hopefully will benefit from the answers.

Although each of these topics is important in its own right, the real goal of this seminar was to bring them together to allow cross-fertilisation. Here is an example. A key issue for MCMC is the rate at which a Markov chain converges to equilibrium, which determines the length of simulation needed to get a good estimate. An important insight has been that this mixing rate is connected to the phenomenon of phase transitions in statistical physics. But it also seems likely that phase transitions are connected with computational intractability more generally, i.e., resistance to all efficient approximation algorithms, not just those based on MCMC. A further example is provided by the way algebra pervades several of our topics – holographic algorithms, complexity of counting, and constraint satisfaction – and yet the connections between these are only now being explored. For example, algebraic methods permit semi-automatic generation of reductions between counting problems, and open up the speculative possibility of resolving the P versus NP question positively through “accidental algorithms”.

We are interested in the complexity of counting in different models of computation. Counting in models of arithmetic circuits is intimately connected with the permanent versus determinant problem. The latter has recently triggered the study of several specific counting problems such as the computation of Littlewood-Richardson coefficients. Another direction of research that is relevant to the meeting is the classification of counting problems in computational algebraic geometry (counting irreducible factors, connected components, etc).

Two key applications areas, statistical physics and constraint satisfaction, have a central role. The problem of computing and approximating weighted sums already arises frequently in statistical physics, where such sums are referred to as partition functions. Constraint Satisfaction is a wide class of problems

which arose in the context of AI – many computer science problems can be cast in this framework. Weights are not traditionally considered in CSP, but with this addition, many applications can be viewed in terms of counting CSPs.

■ Participation and Programme

The seminar brought together 43 researchers from Canada, China, Europe, India, Israel, Japan and the United States with interests and expertise in different aspects of computational counting. Among them there was a good mix of senior participants, postdoctoral researchers and PhD students. Altogether, there were 32 talks over the week.

If the spread of talks at the meeting is a reliable guide, the most active topics in the field at the moment are: algorithms and complexity in algebraic models, the complexity of Counting CSPs (Constraint Satisfaction Problems), and holographic algorithms and the holant framework. Other topics covered included: graph polynomials, MCMC (Markov Chain Monte Carlo) algorithms, parameterised complexity, phase transitions/decay of correlation and its relation to computational complexity, streaming algorithms, and exponential-time exact algorithms. In addition to

the technical presentations listed in the online programme, there were tutorial-style talks on topics featured in the Seminar. On Monday afternoon, Tyson Williams introduced the audience to holant problems and holographic transformations, and on Tuesday, Thore Husfeldt provided a similar service for newcomers to ETH and #ETH (the “Exponential Time Hypothesis” and its counting analogue).

One of the main aims of the seminar was to bring together researchers from different, but related fields, covering all aspects of computational counting with the goal of fostering the exchange of knowledge and to stimulate new research. This goal was fully achieved according to our opinion and the participant’s feedback. The programme was as usual a compromise between allowing sufficient time for participants to present their work, while also providing unstructured periods for informal discussions. New contacts and maybe even friendships were made.

Snow and an early sunset did not prevent the traditional Wednesday “hike” from taking place, though they did curtail it somewhat. The scenery was enhanced by the recent snowfall.

The organisers and participants thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive meeting.

4.4 Civilian Crisis Response Models

Organizers: Ozlem Ergun, Bernhard Katzy, Ulrike Lechner, and Luk van Wassenhove
Seminar No. 13041

Date: January 20–25, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.67

© Creative Commons BY 3.0 Unported license

© Ulrike Lechner and Bernhard Katzy



Participants: Edwin Bakker, Oliver Block, Kateryna Bondar, Matthias Brechmann, Ozgur Dedehayir, Simon French, Hanno Friedrich, Ivo Häring, Bernd Hellgrath, Erich Heumüller, Nico Kaptein, Bernhard Katzy, Nicole Krämer, Erik Kropat, Dietmar Kühne, Ulrike Lechner, Francesc Miralles, Stefan Pickl, Wolfgang Reinhardt, Jens Schwarter, Gideon Shimshon, Christina Weber, Nils B. Weidmann, Heiko Werner, Volker Wulf

The vulnerability of modern societies to the threats of man made and natural disaster increases and scale and number of disasters are expected to rise. The earthquakes of Haiti with its subsequent Cholera epidemics, the natural disasters in Pakistan as well as the ongoing situation in Japan illustrate the need for effective and efficient crisis and disaster response organizations as well as humanitarian aid organizations in developing and first world countries. Disaster preparedness is a key to effectiveness and efficiency in case of crisis or disaster – but we observe that natural and human disasters are too often beyond what is being planned for.

There is a need for new and better approaches in disaster and crises response and humanitarian aid. Think of IT-systems and how well designed systems can help or think of what science can contribute in terms of models, methods, instruments and tools for analysis and decision making. This Dagstuhl Seminar is motivated by the fact that computer science is an enabler for the changes and should contribute to the body of scientific knowledge and instruments and tools alike. This seminar on crisis response Models aims to make a contribution to the systematic development of a body of scientific knowledge for crisis and disaster response and Humanitarian Aid organizations. We invite researchers and practitioners in the field of humanitarian aid and crisis and disaster response as well as researchers in computer science and related disciplines to this Dagstuhl Seminar on Civilian Crisis Response Models.

We address with this seminar on crisis response models questions concerning the design of systems in crisis and disaster response and humanitarian aid. Currently, there is a window of opportunity for redesigning the crisis response system as the proliferation of mobile phones, smart phones and social software facilitate novel services and new C2 systems allows for new designs. Many examples demonstrate the increasing use of social media in emergencies: For human and man-made

disasters websites and Internet services are created to support the inflicted population as well as the aid organizations. A popular and successful example is Ushahidi, a NGO developed platform in response to civil war in Kenya 2008 mapping incidents of violence. In the ongoing crisis in Japan, Twitter and Facebook messages were compiled to provide guidance of what kind of help is needed. Web services are used widespread to locate missing persons. “Google Crisis” provides its set of services to be deployed in case via the Google website.

These systems, many of which have been created ad-hoc by volunteers illustrate the feasibility of better information systems in crisis response management. In many cases, they turned out to be efficient, precise and easy to operate. From these services, evaluation towards a permanently information system is needed. These novel systems illustrate the need for good governance and the need to analyze and reconsider the whole disaster response system with its information flows. What is the impact of the use of such systems in case of a disaster on communication, logistics, the behavior of the population and the aid organizations? Again, scientific methods eventually might be useful to build new systems and develop new processes and strategies.

With this Dagstuhl Seminar on Civilian Crisis Response Models we go beyond the design of technology and aims at contributing to the scientific body of knowledge of crisis and disaster response and Humanitarian aid. Disaster preparedness is the area in the field of crisis and disaster management Civilian Crisis Response Models that requires well developed, evidence-based quantitative models and theories to feed and guide the simulations, optimizations, serious games, analytical methods, architectures and process models, creative techniques or case studies. Disaster preparedness requires its body of scientific knowledge to be used for exploring disaster preparedness, for building IT-systems, for assessing humanitarian aid and disaster response organizations and for guiding the necessary changes in the crisis response

system to adopt it to new threats and new scenarios. Methods and models are crucial for making better decisions in tight financial situations.

The research leading to these results has received funding from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme FP7/2007-2013/ under REA grant agreement n° 317382, NITIMesr.

4.5 Epidemic Algorithms and Processes: From Theory to Applications

Organizers: Benjamin Doerr, Robert Elsässer, Pierre Fraigniaud, and Rachid Guerraoui
Seminar No. 13042

Date: January 20–25, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.94

© Creative Commons BY 3.0 Unported license

© Benjamin Doerr, Robert Elsässer, and Pierre Fraigniaud



Participants: Chen Avin, Keren Censor-Hillel, Pierluigi Crescenzi, Oksana Denysyuk, Benjamin Doerr, Carola Doerr, Robert Elsässer, Pascal Felber, Pierre Fraigniaud, Davide Frey, Tobias Friedrich, Chryssis Georgiou, George Giakkoupis, Bernard Haeupler, Hovhannes A. Harutyunyan, Anna Huber, Amos Korman, Silvio Lattanzi, Andrea Marino, Gerhard Niederbrucker, Adrian Ogierman, Konstantinos Panagiotou, Alessandro Panconesi, Francesco Pasquale, Luis Rodrigues, Stefan Rührup, Thomas Sauerwald, Christian Schindelhauer, He Sun, Philipp Woelfel

The Dagstuhl seminar 13042 “Epidemic Algorithms and Processes: From Theory to Applications” took place from January 20 to 25, 2013, and the main goal of the seminar was to fertilize interaction between theory and applications in this emerging research area. Especially in the algorithmic community several fundamentally new ideas have been developed in recent years. At our Dagstuhl seminar, we explored them further, by mixing various ideas coming from experts working on different fields. Theoretical computer scientists presented their results and methods, in order to disseminate them to a wider community. Researchers from application areas presented their current findings and new challenging research directions, in order to influence (theoretical) research toward real-world applications. The interaction between the seminar participants led to ample discussions and further research collaborations between different domains.

Epidemic algorithms provide a powerful paradigm for distributed computing. Some of the most interesting application areas are the efficient dissemination of updates in replicated data-bases, as well as data dissemination in peer-to-peer systems or wireless sensor networks. By contacting random neighbors in parallel, and making them join forces, an epidemic like progress can be achieved. Furthermore, epidemic processes inherently possess a high level of simplicity and robustness, and therefore the corresponding algorithms can easily deal with the dynamically changing structure of the networks mentioned before.

Theoretical Computer Science makes these useful observations precise and provides certain performance guarantees. One of the well-known algorithms is the so called *randomized rumor spreading*, which disseminates a piece of information in a network to all nodes in a number of communication rounds. In the corresponding communication model, in each round every informed node (i.e., a node which possesses the message) passes/retrieves the information to/from a randomly

chosen neighbor. Since 2008, epidemic algorithms received an increased attention by the theory community, leading to a series of new developments such as the development of new analysis techniques for e.g. the bit-complexity of random phone call algorithms, flooding protocols for dynamic graphs, or relating the performance of an epidemic algorithm to the conductance of the network. On the other side, new algorithm design principles have been introduced, which allow the nodes to remember (and avoid) a certain number of previously contacted neighbors, or the use of intentionally dependent randomized decisions. The first modification resulted in an exponential improvement in the number of message transmissions, and lead to the remarkable result that in social networks information can be spread in sublogarithmic time. The second idea gave rise to a number of high-quality papers ranging from, e.g., a theoretical analysis of the amount of randomness needed to the design of the first epidemic rumor spreading algorithm having a safe termination criterion.

One of the main goals of the seminar was to intensify the collaboration between theory and application fields on epidemic algorithms and processes. We mainly concentrated on two major applications. The first one focuses on the construction and maintenance of peer-to-peer networks in a highly dynamic scenario. Since the epidemic algorithms described above are scalable, robust against edge or node failures, and only require a small amount of message transmissions, they can successfully deal with the challenges imposed in a peer-to-peer environment.

The second focus was on the generation of personalized connections in social networks by using epidemic algorithms. Personalization is applied to fundamental processes such as dissemination, search, and navigation, in order to improve the benefits of social networking. The generated views give rise to certain clusters within the network, and the gossip algorithm for communicating profiles and broadcasting messages distinguishes then between intra-cluster and inter-cluster connections.



Fig. 4.1

Juliana Hümpfner – Untitled. Part of the Dagstuhl art collection and donated by Saarland Sportfoto GmbH.

4.6 Software Certification: Methods and Tools

Organizers: Darren Cofer, John Hatcliff, Michaela Huhn, and Mark Lawford
Seminar No. 13051

Date: January 27 to February 1, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.111

© Creative Commons BY 3.0 Unported license

© Darren Cofer, John Hatcliff, Michaela Huhn, and Mark Lawford



Participants: Dominique Blouin, Darren Cofer, Cyrille Comar, Mirko Conrad, John S. Fitzgerald, Kim R. Fowler, Hubert Garavel, Janusz Górski, Arie Gurfinkel, John Hatcliff, Mats P. E. Heimdahl, Constance L. Heitmeyer, Michael Holloway, Jozef Hooman, Jérôme Hugues, Michaela Huhn, Hardi Hungar, Daniel Kästner, Peter Karpati, Vikash Katta, Tim Kelly, Andrew King, John C. Knight, Brian Larson, Mark Lawford, Dominik Mader, Tom S. Maibaum, John McDermid, Dominique Méry, Frank Ortmeier, Richard F. Paige, Andras Pataricza, Jan Philipps, Robby, Julia Rubin, John Rushby, Bernhard Schätz, David von Oheimb, Alan Wassying, Jens H. Weber, Virginie Wiels

■ Context

An increasingly important requirement for success in many domains is the ability to cost-effectively develop and certify software for critical systems (e.g. pacemakers, health monitoring equipment, core banking applications, financial reporting, nuclear reactors, rail automation and active safety in vehicles etc.). Software errors in each of these domains continue to lead to catastrophic system failures, sometimes resulting in loss of life. A recent report by the U.S. National Academy of Sciences [1], concludes that “new techniques and methods will be required in order to build future software systems to the level of dependability that will be required...In the future, more pervasive deployment of software...could lead to more catastrophic failures unless improvements are made.” Thus, society is increasingly demanding that software used in critical systems must meet minimum safety, security and reliability standards. Manufacturers of these systems are in the unenviable position of not having consistent and effective guidelines as to what constitutes acceptable evidence of software quality, and how to achieve it. This drives up the cost of producing these systems without producing a commensurate improvement in dependability.

Multiple trends and activities (a) point to the changing nature of development of certified systems and (b) indicate the need for community-wide efforts to assess and form a vision of the future for development of certified systems.

New and Evolving Standards. To adapt to the significant changes in the role of software in dependable systems and to improve current industrial practice in software engineering, international standards like the IEC 61508 are currently under revision. DO-178C governing certification of software in commercial aircraft has recently been revised to accommodate the use of software technologies such as formal methods and model-based

development processes. In several other software-intensive domains new domain-specific standards are being developed.

Process- vs. Product-oriented Certification. In practice, current certification of software-intensive systems is primarily process based. A reliance on process oriented standards has established a certification practice that is dominated by assessing process-related documents and marking off checklists that are derived from the recommendation annexes of the standards or so-called “approved practice in use”. Thorough evaluation of the product itself or the adequacy, coverage and maturity of design and quality assurance methods are sometimes neglected because there is currently no fundamental agreement on software engineering principles and product qualities to achieve demonstrably dependable software. An alternative to process-oriented certification regimes is “safety and assurance cases” [7]. In Europe, and particularly in the UK, assurance cases have been adopted as a product centric alternative approach to certification and are widely used in practice already. Recently the U.S. FDA has issued guidance documents recommending the use of assurance cases in submissions for approval of infusion pumps. However, while assurance cases offer some product oriented focus to certification, the lack of standardization of safety and assurance case arguments has its own pitfalls [8].

Advances in Formal Methods. In academia, research on formal methods has made substantial progress with respect to scalability and coverage recently, e.g. in tool-supported model-based design and code generation, but also in the area of software model checking or timing analysis [2]. Thus, formally assuring safety requirements has become feasible at least on the level of components. Nevertheless, research usually focuses on specific techniques, thereby often neglecting the cross-cutting nature of dependability and the need of providing traceable evidence.

Software Development Trends. Two trends relevant in industrial software development for critical systems are the success of model-based design environments that support automated code generation and the need to integrate pre-developed or Commercial Off The Shelf (COTS) software components: (1) Model based tools facilitate rapid prototyping and validation and verification in earlier design phases than traditional processes, but with a price of higher effort in the design phases performed by well-trained and experienced personnel. Software quality will only benefit from these approaches, if certification procedures are adapted towards a cost-effective assessment on the level of models wherever it is adequate. For instance, if model based tools are supported by V & V tools that perform some verification at design time, how does this affect certification standards that require independent design and V & V teams? (2) Evidence based upon prior usage and operating history are typically key components in making decisions in industry about the “fitness for use” of a pre-developed software application or component. However, platform-specific and environmental constraints on the usage are sometimes not specified in detail which has lead to catastrophic failures in the past.

Community-building Activities. Various community-building organizations are being formed drive research, education, and cross-domain coordination in the area of software certification. For instance, the Software Certification Consortium (SCC) was formed in 2007 as a North American initiative to promote product based software certification. Its members are drawn from regulators, industry and academia. SCC has been successful in highlighting shortcomings in current certification regimes and in providing challenge problems and example certification artifacts to the broader community.

■ Seminar Topics and Goals

The Dagstuhl Seminar 13051 *Software Certification: Methods and Tools* brought together experts for the purpose of assessing the current state of practice, identifying challenges, promising techniques/methods, and for creating a road map for future research, education, and standards development in the area of certification of software and systems.

The seminar addressed the following topics:

- Identification of the challenges, regulatory bodies, primary certification standards, typical development and certification processes in variety of safety-critical domains including avionics, automotive, medical systems, and rail, as well as cross-cutting aspects of security certification.
- Developing a *rational basis* for the primary activities in certification. This included work on the interrelation between i) how we develop software in a way that facilitates certification; and, ii) how we collect and use evidence about software products to evaluate whether they should or should not be certified for use, and iii) cost-benefits issues in certification.
- Pros and cons of assurance-cases in regulatory regimes, assessing the confidence given by assurance cases, new techniques for presenting assurance case arguments, tools for managing the collection of evidence and organization of arguments for assurance cases, and the relationship between assurance cases and software certification standards such as DO-178C.
- The use of tools and open source infrastructure in certification, along with new approaches and guidelines for qualifying tools for use in development of certified systems.
- The latest advances in relevant formal methods for software verification, and integrating formal method with other quality assurance techniques such as testing in the context of certified system development.

- The increasing use of “systems of systems” in safety-critical domains, and the need for new approaches supporting compositional certification and reuse of components in the context of certified systems.
- The structure, nature, use, of current certification standards, current business models and organizational principles for developing standards, and how these aspects might be evolved to better address the needs of the community.
- Strategies for managing the complexity of software intensive systems, including model based development, refinement-based methodologies, and generative techniques.
- Challenges problems, infrastructure, and pedagogical resources to support research and education for both academia and industry in the area of certified system development.

■ Seminar Participants and Activities

41 researchers participated in the “Software Certification” Seminar, 21 academic researchers, 10 are affiliated to research institutes and 10 experts from industries proving the strategic relevance of the subject to both, research and practice. With about 40% the portion of North American participants was remarkable high.

The seminar started with an introductory session on Monday morning at which the organizers recapitulated the outline, the objectives, and goals of the seminar. Each participant shortly introduced him/herself, his/her scientific background and personal goals for the seminar week. Then senior experts gave an overview on software certification in different domains, namely the avionic, nuclear, medical devices, automotive, and the rail domain. Monday afternoon ended with a discussion on the major differences and similarities between software assessment in the domains and cross-domain challenges. From Tuesday to Thursday experts presented their work. Panel discussions, challenge problem advertisements as well as working group sessions took place in the afternoons and evenings. A wide range of topics was covered including assurance cases and the fundamentals of how to achieve evidence, tool support to software assessment in the certification process, experience reports and new methodologies for the medical device domain, model based design methods appropriate to certification, issues in cloud security and security certification, tools and methods for static analysis, formal verification and testing. On Thursday evening we had a fruitful discussion with the participants of the Dagstuhl seminar on “Multicore Enablement for Embedded and Cyber-Physical Systems” organized by Andreas Herkersdorf, Michael Hinchey, and Michael Paulitsch that was held in parallel. Among others the following questions were discussed: What are the requests on predictability that have to be satisfied by multicore architectures to be well suited for dependable systems? What are the compelling cyberphysical dependable applications that need multicore architectures? What mechanisms known from dependable software development may be transferred to multicore architecture design and vice versa? Friday was dedicated to working groups as well as outlining and scheduling post seminar proceedings in which we plan to summarize the state of the art in software certification and the results of the seminar. The areas identified by the plenum to be most relevant for further progress on software certification are:

- Fundamentals on confidence and evidence
- Compositional certification
- Education on dependable systems and certification
- Tool qualification
- Security
- Methods for the development of certifiable software and methods supporting certification

■ References

- 1 D. Jackson, M. Thomas, L. Millett. “Software for Dependable Systems: Sufficient Evidence?” Committee on Certifiably Dependable Software Systems, National Research Council, National Academies Press, 2007.
- 2 M. Huhn, H. Hungar. UML for software safety and certification – Model-based development of safety-critical software-intensive systems. In H. Giese, G. Karsai, E. Lee, B. Rumpe, and B. Schätz (Eds.): *Model-Based Engineering of Embedded Real-Time Systems – Int’l Dagstuhl Workshop*, Dagstuhl Castle, Germany, November 4–9, 2007. Revised Selected Papers, LNCS 6100, Springer, pp. 203–240. 2011. DOI: 10.1007/978-3-642-16277-0_8
- 3 M. Huhn, A. Zechner. Analysing Dependability Case Arguments Using Quality Models. In B. Buth, G. Rabe, and T. Seyfarth (Eds.): *28th Int’l. Conf. on Computer Safety, Reliability, and Security (SAFE-COMP)*, LNCS 5775, Springer, pp. 118–131, 2009. DOI: 10.1007/978-3-642-04468-7_11
- 4 A. Wass yng, T.S.E. Maibaum, M. Lawford, On Software Certification: We Need Product-Focused Approaches. C. Choppy and O. Sokolsky (Eds.): *Monterey Workshop 2008*, LNCS Vol. 6028, Springer, pp. 250–274, 2010. DOI: 10.1007/978-3-642-12566-9_13
- 5 J. Hatcliff, M.P.E. Heimdahl, M. Lawford, T.S.E. Maibaum, A. Wass yng, F.L. Wurden. A Software Certification Consortium and its Top 9 Hurdles. In *Proc. of the First Workshop on Certification of Safety-Critical Software Controlled Systems (SafeCert 2008)*, ENTCS, Vol. 238, No. 4, pp. 11–17, 2009. DOI: 10.1016/j.entcs.2009.09.002
- 6 FDA, “FDA Launches Initiative to Reduce Infusion Pump Risks,” News Release, April 23, 2010 (see: <http://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm209042.htm>)
- 7 R. Bloomfield and P. Bishop. Safety and assurance cases: Past, present and possible future – an Adelard perspective. In C. Dale, T. Anderson (Eds.): *Making Systems Safer*, Proc. of the Eighteenth Safety-Critical Systems Symp., Bristol, UK (February 2010), pp. 51–67.
- 8 A. Wass yng, T.S.E. Maibaum, M. Lawford and H. Behr. Is there a case against safety cases? Submitted to post-proceedings volume of Monterey 2010 Workshop, to be published in LNCS.

4.7 Multicore Enablement for Embedded and Cyber Physical Systems

Organizers: Andreas Herkersdorf, Michael G. Hinchey, and Michael Paulitsch

Seminar No. 13052

Date: January 27 to February 1, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.1.149

© Creative Commons BY 3.0 Unported license

© Andreas Herkersdorf, Michael G. Hinchey, and Michael Paulitsch

Participants: Michael Deubzer, Christian El Salloum, Rolf Ernst, Glenn Farrall, Christian Ferdinand, Massimo Ferraguto, Steffen Görzig, René Graf, David Gregg, Geoff Hamilton, Andreas Herkersdorf, Johan Lilius, Enno Lübbers, Roman Obermaisser, Sri Parameswaran, Michael Paulitsch, Stefan M. Petters, Matthias Pruksch, Georg Sigl, Claus Stellwag, Jürgen Teich, Christian Thiel, Lothar Thiele, Sergey Tverdyshev, Theo Ungerer, Stefan Wallentowitz, Alexander Weiss, Thomas Wild, Reinhard Wilhelm



Multicore processors are a key enabling technology for solving grand societal challenges of the coming decades. Secure and ecological mobility, geographic coverage of high-tech healthcare, sustainable energy generation, distribution and management, and in general the development of our digitized society impose compute performance requirements on distributed embedded and cyber physical IT equipment which makes multicore technology indispensable. All leading processor vendors – ARM, Freescale, IBM, Infineon, Intel, MIPS, Nvidia – follow a strictly multicore-oriented strategy. Due to the paradigm shift from exploiting instruction level to process level parallelism, multicore processors are superior over single-core representatives with respect to computing performance and energy efficiency. Prerequisite is, processes can be balanced among parallel cores such that the nominally available computing performance can be utilized effectively, and cores can be set into sleep mode or power gated when not busy. As of today, the ability to efficiently utilize the available resources depends to a large extent on the aptitude of experienced programmers and the inherent ability of being able to parallelize the computing problem.

Embedded and Cyber Physical Systems exhibit demands for “non-functional requirements”, such as low(est) power and energy dissipation, reliability, availability and security, real-time and cost constraints, which are typically not found to the same extent in general purpose computing applications. The enablement of multicore technology for embedded and cyber-physical markets imposes serious challenges to industry and academia which can easily overwhelm the capabilities and capacities of individual corporations or even consortia. Industry and university research in Europe recognized early and invested significantly into the establishment of multicore know-how and competences. Examples of related projects at EU level and in Germany are: RECOMP – Reduced Certification Costs Using Trusted Multicore Platforms, ACROSS – ARTEMIS CROSS-Domain Architecture, SPES 2020 – Software Plattform Embedded Systems 2020, Cesar – Cost-

efficient methods and processes for safety relevant embedded systems, MERASA – Multicore Execution of Hard Real-Time Applications Supporting Analysability (see Relationship to other seminars and projects for a more complete listing), and ARAMiS – Automotive, Railway and Avionics Multicore Systems.

The seminar brought together leading industry and university research groups from different fields of embedded system design and application development, multicore architecture and hardware/software design methodology & tools. The main objective of the seminar was on reporting experiences and discussing challenges of reusable and transferable multicore technologies among participants representing different application markets and scientific backgrounds. The technical focus of the agenda was on:

- Generic hardware/software building blocks for real-time performance, dependability, functional safety and security for embedded systems built around enhanced standard multicore solutions.
- System modeling, design and validation methods and tools for such platforms.

The seminar established new and strengthened existing ties between players and networks in the area of multicore and embedded technologies. Topical working groups were formed on the following topics:

- Specification & Interference
- Industrial Perspective on MultiCore Motivations and Challenges
- Certification of Safety-Critical Multicore Systems: Challenges and Solutions
- Network-on-Chip – Dependability and Security Aspects
- Multicore Ecosystem
- Secure Elements in future embedded multicore systems

The working groups compiled summaries reflecting the status and outlook on the respective topic.

4.8 Fault Prediction, Localization, and Repair

Organizers: Mary Jean Harrold, Friedrich Steimann, Frank Tip, and Andreas Zeller
Seminar No. 13061

Date: February 3–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.1

© Creative Commons BY 3.0 Unported license

© Mary Jean Harrold, Friedrich Steimann, Frank Tip, and Andreas Zeller



Participants: Rui Abreu, Shay Artzi, George K. Baah, Benoit Baudry, Margaret M. Burnett, Satish Chandra, Jake Cobb, Julian Dolby, Marcus Frenkel, Vijay Ganesh, Milos Gligoric, Alessandra Gorla, Mangala Gowri Nanda, Christian Hammer, Mary Jean Harrold, Jens Krinke, Ben Liblit, Rupak Majumdar, Martin Monperrus, Alessandro Orso, Marco Pistoia, Andy H. Podgurski, Jeremias Rößler, Abhik Roychoudhury, Barbara G. Ryder, Hesam Samimi, Friedrich Steimann, Lin Tan, Frank Tip, Emina Torlak, Cemal Yilmaz, Andreas Zeller, Xiangyu Zhang, Thomas Zimmermann

Even today, an unpredictable part of the total effort devoted to software development is spent on debugging, i.e., on finding and fixing bugs. This is despite the fact that powerful static checkers are routinely employed, finding many bugs before a program is first executed, and also despite the fact that modern software is often assembled from pieces (libraries, frameworks, etc.) that have already stood the test of time. In fact, while experienced developers are usually quick at finding and fixing their own bugs, they too spend too much time with fixing the interplay of components that have never been used in combination before, or just debugging the code of others. Better automated support for predicting, locating, and repairing bugs is therefore still required.

Due to the omnipresence of bugs on the one side and the vastly varying nature of bugs on the other, the problems of fault prediction, localization, and repair have attracted research from many different communities, each relying on their individual strengths. However, often enough localizing a bug resembles the solution of a criminal case in that no single procedure or evidence is sufficient to identify the culprit unambiguously. It is therefore reasonable to expect that the best result can only be obtained from the combination of (insufficient) evidence obtained by different, and ideally independent, procedures. One main goal of this seminar is therefore to connect the many different strands of research on fault prediction, localization, and repair.

For researchers it is not always obvious how debugging is embedded in the software production process. For instance, while ranking suspicious program statements according to the likelihood of their faultiness may seem like a sensible thing to do from a research perspective, programmers may not be willing to look at more than a handful of such locations when they have their own inkling of where a bug might be located. On the other hand, commercial programmers may not be aware of the inefficiency of their own approaches to debugging, for which promising

alternatives have been developed by academics. Bringing together these two different perspectives is another goal of this seminar.

Last but not least, the growing body of open source software, and with it the public availability of large regression test suites, provide unprecedented possibilities for researchers to evaluate their approaches on industrial-quality benchmarks. In fact, while standard benchmarks such as the so-called Siemens test suite still pervade the scientific literature on debugging, generalization of experimental results obtained on such a small basis is more than questionable. Other disciplines, such as the model checking or the theorem proving communities, have long established competitions based on open benchmarks to which anyone can submit their problems. Based on such benchmarks, progress would be objectively measurable, and advances in research would be better visible. It is another goal of this seminar to establish a common understanding for the need of such benchmarks, and also to initiate the standards necessary for installing them.

4.9 Decentralized Systems for Privacy Preservation

Organizers: Sonja Buchegger, Jon Crowcroft, Balachander Krishnamurthy, and Thorsten Strufe

Seminar No. 13062

Date: February 3–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.22

© Creative Commons BY 3.0 Unported license

© Sonja Buchegger, Jon Crowcroft, Balachander Krishnamurthy, and Thorsten Strufe

Participants: Jonathan Anderson, N. Asokan, Rainer Böhme, Nikita Borisov, Sonja Buchegger, Ramon Caceres, Jan Camenisch, Jon Crowcroft, George Danezis, Claudia Diaz, Vijay Erramilli, Simone Fischer-Hübner, Paul Francis, Ian Goldberg, Artur Hecker, Urs Hengartner, Jaeyeon Jung, Mohamed Ali Kaafar, Gunnar Kreitz, Balachander Krishnamurthy, Leonardo A. Martucci, Bart Preneel, Stefanie Roos, Krzysztof Rzadca, Hervais-Clemence Simo Fhom, Thorsten Strufe, Paul Syverson, Claire Vishik, Marcel Waldvogel



Distributed and decentralized systems offer more potential resilience to various failures, and, on paper, higher aggregate availability than centralized systems. Centralized management repositories lead to potential risks to users' privacy and the temptation to monetize processing of large aggregates of such data, as seen in systems such as webmail, search and online social networks. Recent years have seen the emergence of projects building prototypes with varying levels of decentralization to reduce these risks. Such systems have not seen great success in contrast to large cloud services. This seminar brought together diverse groups to tackle a series of questions to attempt to answer what may be the root causes of the logjam preventing success of these alternative approaches. There appears to be some consensus amongst at least some groups that there are good reasons for these alternatives. We present here the output of our group working sessions on these questions. We also provide the reasoning and outcomes of the discussions along with an evaluation of the effectiveness of our mode of working in this seminar.

4.10 Dependence Logic: Theory and Applications

Organizers: Samson Abramsky, Juha Kontinen, Jouko Vaananen, and Heribert Vollmer
Seminar No. 13071

Date: February 10–15, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.45

© Creative Commons BY 3.0 Unported license

© Samson Abramsky, Juha Kontinen, Jouko Väänänen, and Heribert Vollmer



Participants: Samson Abramsky, Dietmar Berwanger, Olaf Beyersdorff, Andreas R. Blass, Julian Bradfield, Panayiota Constantinou, Nadia Creignou, Anuj Dawar, A. Philip Dawid, Arnaud Durand, Johannes Ebbing, Uwe Egly, Fredrik Engström, Pietro Galliani, Georg Gottlob, Erich Grädel, Miika Hannula, Lauri Hella, Asa Hirvonen, Wilfrid Hodges, Theo Janssen, Phokion G. Kolaitis, Juha Kontinen, Antti Kuusisto, Pierfrancesco La Mura, Sebastian Link, Allen L. Mann, Arne Meier, Eric Pacuit, Tero Tulenheimo, Jouko Väänänen, Jonni Virtema, Heribert Vollmer, Dag Westerstahl, Fan Yang

■ Brief Introduction to the Topic

Dependence Logic is a new tool for modeling dependencies and interaction in dynamical scenarios. Reflecting this, it has higher expressive power and complexity than classical logics used for these purposes previously. Algorithmically, first-order dependence logic corresponds exactly to the complexity class NP and to the so-called existential fragment of second-order logic.

Since the introduction of dependence logic in 2007, the framework has been generalized, e. g., to the contexts of modal, intuitionistic and probabilistic logic. Moreover, interesting connections have been found to complexity theory and database theory, and dependence logic has been applied in areas such as linguistics, social choice theory, and physics. Although significant progress has been made in understanding the computational side of these formalisms, still many central questions remain unsolved so far.

The notions of logical dependence and independence are pervasive, and occur in many areas of science. The development of logical and semantical structures for these notions provides an opportunity for a systematic approach, which can expose surprising connections between different areas (e. g., quantum mechanics, social choice theory, and many more), and may lead to useful general results.

One of the main aims of this Dagstuhl Seminar was to bring together, for the first time, researchers working in this area so that they can communicate state-of-the-art advances and embark on a systematic interaction. In particular, bringing together researchers from areas of theoretical studies with the application areas will enhance the synergy between the different communities working on dependence logic.

■ Organization of the Seminar and Activities

The workshop brought together 35 researchers from mathematics, theoretical physics, statistics, social choice theory, and theoretical computer science. The participants consisted of both senior and junior researchers, including a number of postdocs and a few advanced graduate students.

Participants were invited to present their work and to communicate state-of-the-art advances. Seventeen talks of various lengths took place over the five days of the workshop. Introductory and tutorial talks of 90-60 minutes were scheduled prior to workshop. Most of the remaining slots were filled, mostly with shorter talks, as the workshop commenced. The organizers considered it important to leave ample free time for discussion.

The tutorial talks were scheduled during the beginning of the week in order to establish a common background for the different communities that came together for the workshop. The presenters and topics were:

- Jouko Väänänen, Dependence Logic
- Erich Grädel, Logics with team semantics and second-order reachability games
- Philip Dawid, Conditional Independence and Irrelevance
- Pietro Galliani, Definability Issues in Team Semantics
- Phokion Kolaitis, Foundations and Applications of Schema Mappings
- Samson Abramsky, From Quantum Mechanics to Logic, Databases, Constraints, Complexity and Beyond
- Sebastian Link, Dependence, Independence, Logic
- Wilfrid Hodges, Compositionality: Its history and formalism
- Eric Pacuit, Dependence and Independence in Social Choice Theory

There were additionally 8 other talks with a more focused and technical topic.

1. Georg Gottlob, From Local Hidden Variables in Quantum Mechanics to Robust Colorability and Satisfiability
2. Panayiota Constantinou, Extended Conditional Independence
3. Fan Yang, Uniform definability in propositional dependence logic
4. Pierfrancesco La Mura, A double-slit experiment for non-classical interference effects in decision-making
5. Julian Bradfield, Concurrency, causality and dependency
6. Miika Hannula, Axiomatizing first-order consequences in independence logic
7. Andreas R. Blass, Introduction to Secret Sharing
8. Arnaud Durand, Complexity issues in dependence logic

The workshop achieved its aim of bringing together researchers from various related communities to share state-of-the-art research. The organizers left ample time outside

of this schedule of talks and many fruitful discussions between participants took place throughout the afternoons and evenings.

■ Concluding Remarks and Future Plans

The organizers regard the workshop as a great success. Bringing together researchers from different areas fostered valuable interactions and led to fruitful discussions. Feedback from the participants was very positive as well. Many attendants expressed their wish for a continuation and stated that this seminar was among the most fruitful Dagstuhl seminars they attended.

Finally, the organizers wish to express their gratitude toward the Scientific Directorate of the Center for its support of this workshop, and hope to establish a series of workshops on *Dependence Logic: Theory and Applications* in the future.

4.11 Mechanisms of Ongoing Development in Cognitive Robotics

Organizers: Jacqueline Fagard, Roderic A. Grupen, Frank Guerin, and Norbert Krüger
Seminar No. 13072

Date: February 10–15, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.55

© Creative Commons BY 3.0 Unported license

© Jacqueline Fagard, Roderic A. Grupen, Frank Guerin, and Norbert Krüger



Participants: John Alexander, Gianluca Baldassare, Emily W. Bushnell, Paul R. Cohen, Rana Esseily, Jacqueline Fagard, Severin Fichtl, Roderic A. Grupen, Beata Joanna Grzyb, Frank Guerin, Verena V. Hafner, Matej Hoffmann, Bipin Indurkha, Sinan Kalkan, George Konidakis, Norbert Krüger, Benjamin Kuipers, Ales Leonardis, Honghai Liu, Jeffrey J. Lockman, Bärbel Mertsching, J. Kevin O'Regan, Mohamed Oubbati, Pierre-Yves Oudeyer, Justus Piater, Lauriane Rat-Fischer, Helge Ritter, Aaron Sloman, Georgi Stojanov, Alexander Stoytchev, Emre Ugur, Markus Vincze, Florentin Wörgötter

In cognitive robotics “ongoing development” refers to the ability to continuously build on what the system already knows, in an ongoing process, which acquires new skills and knowledge, and achieves more sophisticated levels of behaviour. Human infants are possibly the best known demonstrators of this ability; developmental psychology has many results documenting what infants can and cannot do at various ages, however we know very little about the mechanisms underlying the development. On the robotics side, creating a computational system which displays ongoing development is still an unsolved problem. There are major unsolved questions regarding the mechanisms of ongoing development, in both biological and artificial systems; for example: how to transfer existing skills to a new context, how to build on existing skills, and how to represent knowledge (or skills).

The primary aim of the seminar was to bring together researchers from two communities (developmental robotics and infant developmental psychology) in order to spawn new collaborative research projects which will advance our scientific understanding of the mechanisms underlying ongoing development (whether in infants or robots). We especially focused on perception, understanding and manipulation skills relating to physical objects in the world, and the skills which infants acquire in approximately the 4–24 months period.

Working groups were formed in the areas of (i) transfer of means/skills; (ii) motor skills/manipulation; (iii) concepts/representations; (iv) motivation; (v) visual perception. These discussed gaps between what infants and robots can do and what research might close the gap. In discussion groups the most significant issue that was raised (and discussed at length) was how to get psychologists and roboticists talking together and doing research together, as there seems to exist a wide gap between the communities. It was concluded that there was a need for psychologists to become computer scientists and computer

scientists to become psychologists; i.e. that the meeting of the two fields would not happen simply by people getting together in a room, but that the meeting must happen inside individual heads. Furthermore challenge problems were posed by each of the two respective communities; challenges which they would like the other community to work on.



Fig. 4.2
Petra Stoeppel – Musik in der Stadt. Part of the Dagstuhl art collection and donated by: Chris Hankin, Harald Ganzinger, Ulrike Lechner, Flemming Nielson, Hanne Riis Nielson, Ute Vollmar, Roland Vollmar, Reinhard Wilhelm, and participants in Dagstuhl Seminar 99151.

4.12 Consistency in Distributed Systems

Organizers: Bettina Kemme, Ganesan Ramalingam, André Schiper, Marc Shapiro, and Kapil Vaswani

Seminar No. 13081

Date: February 17–22, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.92

© Creative Commons BY 3.0 Unported license

© Bettina Kemme, Ganesan Ramalingam, André Schiper, and Marc Shapiro



Participants: Marcos K. Aguilera, Hagit Attiya, Carlos Baquero, Annette Bieniusa, Alejandro P. Buchmann, Sebastian Burckhardt, Bernadette Charron-Bost, Allen Clement, Mike Dodds, Amr El-Abadi, Alan Fekete, Pascal Felber, Carla Ferreira, Alexey Gotsman, Maurice Herlihy, Ricardo Jimenez-Peris, Bettina Kemme, Petr Kuznetsov, David B. Lomet, Maged M. Michael, Achour Mostefaoui, Yiannis Nikolakopoulos, Fernando Pedone, Nuno Preguica, Vivien Quema, Kaushik Rajan, Ganesan Ramalingam, Robert Rehner, Noam Rinetzky, Luis Rodrigues, Rodrigo Rodrigues, Nicholas Rutherford, Mooly Sagiv, André Schiper, Marc Shapiro, Liuba Shrira, Alexander A. Shvartsman, Pierre Sutra, Douglas B. Terry, Peter Van Roy, Kapil Vaswani, Marko Vukolic, Jennifer L. Welch, Pawel T. Wojciechowski

In distributed systems, there exists a fundamental trade-off between data consistency, availability, and the ability to tolerate failures. This trade-off has significant implications on the design of the entire distributed computing infrastructure such as storage systems, compilers and runtimes, application development frameworks and programming languages. Unfortunately, it also has significant, and poorly understood, implications for the designers and developers of end applications. As distributed computing become mainstream, we need to enable programmers who are not experts to build and understand distributed applications.

A seminar on “Consistency in Distributed Systems” was held from 18th to 22nd, February, 2013 at Dagstuhl. This seminar brought together researchers and practitioners in the areas of distributed systems, programming languages, databases and concurrent programming, to make progress towards the above-mentioned goal. Specifically, the aim was to understand lessons learnt in building scalable and correct distributed systems, the design patterns that have emerged, and explore opportunities for distilling these into programming methodologies, programming tools, and languages to help make distributed computing easier and more accessible.

We may classify current approaches to deal with the challenges of building distributed applications into the following three categories:

- **Strong Consistency and Transactions:** Strong consistency means that shared state behaves like on a centralised system, and programs (and users) cannot observe any anomalies caused by concurrent execution, distribution, or failures. From a correctness perspective, this is a most desirable property. For instance, a database management system protects the integrity of shared state with transactions, which provide the so-called ACID guarantees: atomicity (all-or-nothing), consistency (no transaction in isolation violates database integrity), isolation (intermediate states of a transaction

cannot be observed by another one), and durability (a transaction’s effects are visible to all later ones).

- **Weak Consistency:** Unfortunately strong consistency severely impacts performance and availability [1, 2]. As applications executing in the cloud serve larger workloads, providing the abstraction of a single shared state becomes increasingly difficult. Scaling requires idioms such as replication and partitioning, for which strongly-consistent protocols such as 2-Phase Commit are expensive and hard to scale. Thus, contemporary cloud-based storage systems, such as Amazon’s Dynamo or Windows Azure Tables, provide only provide weak forms of consistency (such as eventual consistency) across replicas or partitions. Weakly consistent systems permit *anomalous* reads, which complicates reasoning about correctness. For example, application designers must now ascertain if the application can tolerate stale reads and/or delayed updates. More parallelism allows better performance at lower cost, but at the cost of high complexity for the application programmer.
- **Principled Approaches to Consistency:** A number of approaches and tools have been developed for reasoning about concurrently-accessed shared mutable data. The concept of linearizability [3] has become the central correctness notion for concurrent data structures and libraries. This has led to significant advances in verification, testing and debugging methodologies and tools. Transactional memory provides a higher-level, less error-prone programming paradigm [4].
- **Principles for weak consistency:** More recently, a number of principles have emerged for dealing with weak consistency. For example, if all operations in a program are monotonic, strong correctness guarantees can be provided without the use of expensive global synchronization. Similarly, certain data structures such as sets and sequences can be replicated in a correct way without synchronisation.

These developments illustrate the benefits of cross-fertilization of ideas between these different communities, focused on the topic of concurrency. We believe that such principled approaches will become increasingly critical to the design of scalable and correct distributed applications. The time is ripe for the development of new ideas by cross-fertilisation between the different research communities.

■ Goals

It is crucial for researchers from different communities working in this same space to meet and share ideas about what they believe are the right approaches to address these issues. The questions posed for the seminar include:

- Application writers are constantly having to make trade-offs between consistency and scalability. What kinds of tools and methodologies can we provide to help this decision? How does one understand the implications of a design choice?
- Weakly consistent systems are hard to design, test and debug. Do existing testing and debugging tools suffice for identifying and isolating bugs due to weak consistency?

- Can we formalize commonly desired (generic) correctness (or performance) properties?
- Can we build verification or testing tools to check that systems have these desired correctness properties?
- How do applications achieve the required properties, while ensuring adequate performance, in practice? What design patterns and idioms work well?
- To what degree can these properties be guaranteed by the platform (programming language, libraries, and runtime system)? What are the performance tradeoffs (when one moves the responsibility for correctness between the platform and application)?

In order to ensure a common understanding between the different research communities that the workshop brings together, the seminar started with a few tutorials from the perspective of each community. Other presentations presented a specific piece of research or a research question. Participants brain-stormed on a specific issue during each of the two break-out sessions.

The abstracts and slides of the different presentations are available at <http://www.dagstuhl.de/mat/index.en.phtml?13081>.

■ References

- 1 Daniel J. Abadi. Consistency tradeoffs in modern distributed database system design. *Computer*, 45(2):37–42, February 2012.
- 2 Eric Brewer. CAP twelve years later: How the “rules” have changed. *IEEE Computer*, 45(2):23–29, February 2012.
- 3 Maurice Herlihy and Jeannette Wing. Linearizability: a correctness condition for concurrent objects. *ACM Transactions on Programming Languages and Systems*, 12(3):463–492, July 1990.
- 4 Maurice Herlihy and J. Eliot B. Moss. Transactional memory: Architectural support for lock-free data structures. In *Int. Conf. on Comp. Arch. (ISCA)*, pages 289–300, San Diego CA, USA, May 1993.

4.13 Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices

Organizers: LeRoy B. Beasley, Hartmut Klauck, Troy Lee, and Dirk Oliver Theis
Seminar No. 13082

Date: February 17–22, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.2.127

© Creative Commons BY 3.0 Unported license

© LeRoy B. Beasley, Hartmut Klauck, Troy Lee, Dirk Oliver Theis



Participants: LeRoy B. Beasley, Hamza Fawzi, Samuel Fiorini, Anna Gál, Nicolas Gillis, Francois Glineur, Joao Gouveia, Alexander Guterman, Volker Kaibel, Stephen Kirkland, Hartmut Klauck, Raghav Kulkarni, Thomas Laffey, Troy Lee, Lek-Heng Lim, Nathan Linial, Pablo Parrilo, Kanstantsin Pashkovich, Sebastian Pokutta, Richard Robinson, Yaroslav Shitov, Adi Shraibman, Dirk Oliver Theis, Rekha R. Thomas, Hans Raj Tiwary, Stefan Weltge

The nonnegative rank is a measure of the complexity of a matrix that has applications ranging from Communication Complexity to Combinatorial Optimization. At the time of the proposal of the seminar, known lower bounds for the nonnegative rank were either trivial (rank lower bound) or known not to work in many important cases (bounding the nondeterministic communication complexity of the support of the matrix).

Over the past couple of years in Combinatorial Optimization, there has been a surge of interest in lower bounds on the sizes of Linear Programming formulations. A number of new methods have been developed, for example characterizing nonnegative rank as a variant of randomized communication complexity. The link between communication complexity and nonnegative rank was also instrumental recently in proving exponential lower bounds on the sizes of extended formulations of the Traveling Salesman polytope, answering a longstanding open problem.

This seminar brought together researchers from Matrix Theory, Combinatorial Optimization, and Communication Complexity to promote the transfer of tools and methods between these fields. The focus of the seminar was on discussions, open problems and talks surveying the basic tools and techniques from each area.

In the short time since the seminar, its participants have made progress on a number of open problems.

■ Program Overview

Background lectures on the connection between matrix factorizations to Communication Complexity and to Combinatorial Optimization were given by the organizers. More importantly, a number of participants contributed their latest research on factorization ranks. In this section, we summarize these talks.

■ Extended Formulations and Linear Optimization

Hamza Fawzi: Many lower bounds on the nonnegative rank only make use of the zero/nonzero pattern of the matrix. For certain applications, in particular for the extended formulation size lower bounds for approximation problems, nonnegative rank lower bounds need to be shown for matrices that are strictly positive. Hamza discussed an interesting approach to nonnegative rank lower bounds via conic programming that does not only rely on the zero/nonzero structure of the matrix. The bound is in many ways analogous to the trace norm lower bound for rank, but making use of the stronger fact that the factorization is nonnegative leads to a copositive program rather than a semidefinite one. For computing the bound in practice, Hamza discussed ways to approximate the bound by semidefinite programs, and examples of using this in practice.

Sam Fiorini: There is a rich theory on the hardness of approximating NP-optimization problems up to certain factors, given complexity assumptions like $P \neq NP$. Very recently a similar topic has emerged in the study of polytopes. Sam talked about tradeoffs between the approximation ratio and the size of linear formulations. One notable result in Sam's talk was that approximating CLIQUE to within $n^{1/2-\epsilon}$ requires extended formulations of exponential size.

■ Complexity

Nati Linial: On the first day, Nati Linial treated us to a survey of higher dimensional analogs of familiar combinatorial objects. For example, we are very familiar with permutation matrices, those matrices with entries from $\{0, 1\}$ with exactly one 1 in every row and column, and know that there are $n! = ((1 + o(1))n/e)^n$ many of them. What about 3-dimensional tensors with entries from $\{0, 1\}$ and exactly one 1 along every

row, column and shaft? Such 2-dimensional permutations turn out to coincide with latin squares and it is known that there are $((1 + o(1))n/e^2)^{n^2}$ many of them. This relies on some beautiful work on the minimum permanent of doubly stochastic matrices. Nati conjectures that the formula generalizes to count the number of d -dimensional permutations, described by a $d + 1$ -tensor with one 1 along every line. That is, that the number of d -dimensional permutations is $((1 + o(1))n/e^d)^{n^d}$. He is able to show such an upper bound, but the lower bound remains open.

Sebastian Pokutta: In order to prove that extended formulations for approximating optimization problems need to be large, communication and information complexity are important tools. In his talk Sebastian described a new approach on how to prove lower bounds on the nonnegative rank of matrices corresponding to the unique disjointness problem when perturbed. He gave tight lower bounds using a new information theoretic fooling set method.

Since the seminar, Sebastian and his co-author Gabór Braun have made available a preprint containing these results [1].

Hans Raj Tiwary: There are entire books of NP-complete problems and explicit reductions between them. For the extension complexity of the associated polytopes, however, this book is still slowly being written—usually by arguing that P is a projection of Q or finding P as a face of Q . Hans discussed the intriguing possibility of automatically turning an NP gadget reduction into a polytope reduction. While still not a general theory, Hans can currently do this for many NP-hard problems and their associated polytopes.

Nicolas Gillis: Nicolas spoke about the problem of actually computing a non-negative factorization of a nonnegative matrix. This talk was important to seminar participants on small matrices, allowing them to test the quality of their lower bounds against upper bounds. On small matrices, these upper bounds can be found computationally. The problem also has applications to compression of images, to identifying topics in documents, even to identifying the mineral composition of rocks from spectral data (hyper-spectral imaging). Nicolas discussed specifically the case of separable matrices. An n -by- n matrix M is r -separable if it has a factorization $M = WH$ where W is n -by- r , H is r -by- N and moreover W is a

subset of the columns of M . Such types of factorization can be more useful in practice. Nicolas talked about a linear programming approach to this problem that is polynomial time and moreover outperforms previous approaches in practice.

■ Matrix Theory

Alexander Guterman: Alexander Guterman gave a survey talk on various matrix ranks over semirings. A big focus was on tropical algebra over the real number with operations $a \oplus b = \max a, b$ and $a \otimes b = a + b$. Tropical algebra provides a way of formulating many hard combinatorial optimization problems (like scheduling problems) in terms of a very elegant linear algebraic type language. In tropical linear algebra there are varying notions of linear independence, for example Gondran-Minoux independence, weak linear independence, and strong linear independence. Each of these gives rise to a different notion of rank of a matrix and a hierarchy of these ranks is known.

Yaroslav Shitov: Yaroslav continued talking about tropical matrix rank, in particular the tropical factorization rank. This is defined as the minimum k such that $A = B \otimes C$ for a n -by- k matrix B and k -by- n matrix C . Note that in tropical matrix multiplication $(B \otimes C)(i, j) = \min_t B(i, t) + C(t, j)$. Yaroslav mentioned a very interesting application of the tropical factorization rank. Say that we are given an instance of the traveling salesman problem, with distances specified by a matrix A , and moreover we are given a tropical factorization $A = B \otimes C$ that witnesses that A has constant factorization rank. Then the resulting traveling salesman instance can be solved in polynomial time! This is a result of Barvinok, Johnson, Woeginger, and Woodroofe. Yaroslav also showed that the problem of detecting if the tropical factorization rank of a matrix is at most 8 is NP-hard.

Richard Robinson: In his talk, Richard Robinson gave a characterization, among all nonnegative matrices, of the extreme-ray / facet slack matrices of polyhedral cones, and vertex/facet slack matrices of polytopes. This characterization leads to an algorithm for deciding whether a given matrix is a vertex/facet slack matrix. The underlying decision problem is equivalent to the polyhedral verification problem whose complexity is unknown.

■ References

- 1 G. Braun and S. Pokutta. Common information and unique disjointness. Technical Report TR13-056, ECCC, 2013.

4.14 Analysis, Test and Verification in The Presence of Variability

Organizers: Paulo Borba, Myra B. Cohen, Axel Legay, and Andrzej Wąsowski
Seminar No. 13091

Date: February 24 to March 1, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.2.144

© Creative Commons BY 3.0 Unported license

© Paulo Borba, Myra B. Cohen, Axel Legay, and Andrzej Wąsowski



Participants: Vander Alves, Sven Apel, Joanne M. Atlee, Kacper Bąk, Don Batory, Thorsten Berger, Eric Bodden, Paulo Borba, Claus Brabrand, Dave Clarke, Andreas Classen, Roberta Coelho, Myra Cohen, Maxime Cordy, Krzysztof Czarnecki, Sabrina de Figueirêdo Souto, Martin Erwig, Alessandro Fantechi, Brady J. Garvin, Rohit Gheyi, Stefania Gnesi, Reiner Hähnle, Øystein Haugen, Martin Fagereng Johansen, Christian Kästner, Shriram Krishnamurthi, Kim Lauenroth, Axel Legay, Martin Leucker, Tiziana Margaria, Dusica Marijan, Jean-Vivien Millo, Gilles Perrouin, Márcio Ribeiro, Ina Schaefer, Holger Schlingloff, Sergio Segura, Vanessa Stricker, Leopoldo Teixeira, Maurice H. ter Beek, Thomas Thüm, Társis Tolêdo, Salvador Trujillo, Eric Walkingshaw, Andrzej Wąsowski, Cemal Yilmaz

The seminar “Analysis, Test and Verification in The Presence of Variability” that took place at Schloss Dagstuhl from February 24 to March 1, 2013, had the goal of consolidating and stimulating research on analysis of software models with variability, enabling the design of variability-aware tool chains. We brought together 46 key researchers from three continents, working on quality assurance challenges that arise from introducing variability, and some who do not work with variability, but that are experts in their respective areas in the broader domain of software analysis or testing research. The participants ranged from those in senior academic positions to successful graduate students. We also enjoyed the presence of several relevant experts from the software development industry.

The seminar included:

1. *Invited presentations on state of the art research in SPL testing and verification.* The presentations were delivered by experts in variability research. The topics included classifying and unifying product-line analyses, combinatorial interaction testing, model-based testing, analysis of programs with variability and model checking with variability. Material relevant to the topic of this Dagstuhl was organized in a recent classification by Thüm and coauthors [4]. The Dagstuhl seminar opened with a presentation of this classification, which created a common ontology for later presentations and discussions. This was very helpful for participants who had different areas of expertise.
2. *A keynote presentation on the Challenges and Science of Variability.*
 We organized a special keynote shared with the German FOSD meeting, that took place in parallel at the Schloss Dagstuhl facilities. The keynote speaker, Professor Don Batory, called for creating a simple meta-theory identifying and relating the core concepts and properties of variability science, i.e. the body of knowledge created by the community

of researchers studying engineering of highly configurable systems. During the workshop, several candidates for the starting point of such theory were mentioned, such as using simple models in constructive logic [2], choice calculus [3] or Clafer [1].

3. *A series of presentations on recent results in Variability Analysis.*

The bulk of the programme was filled with a mixture of research presentations about recent research advances in verification, analysis and test of software with variability. This function of the seminar was particularly important, as the usual dissemination outlets for these contributions are often disjoint – much of the work is normally presented in domain specific publication channels devoted to only test, verification or programming languages. For many participants the seminar created an opportunity to learn about advances at addressing similar problems in the neighboring research communities – an experience that is rarely possible outside of Dagstuhl.

4. *A session of student presentations.*

In order to enrich the presentations by senior researchers with a stream of fresh ideas, we organized a special session devoted to short student presentations. The presenters were selected from the participants of the German FOSD meeting. For many of the students it was a rare opportunity to share their ideas with international authorities in their work area. The topics of these lightning presentations were closely related to the seminar goals and included among others, discussions of experimental evaluation of product line analysis strategies, static analysis, type checking for variability, and performance prediction for configurable systems. The session enabled closer integration between the participants of the two events. Many discussions between the two groups continued throughout the week.

5. *Dynamically planned sessions on how to address the challenges, how to transfer knowledge, tools, and benchmarks between research areas.*

The first session (run by Professor Krzysztof Czarnecki) was devoted to extracting challenges for variability analysis out of industrial requirements. Participants from industry and participants from academia involved in industrial projects provided background on requirements known from projects in avionics, automotive and risk assessment domains. These were further discussed to identify research challenges for future work. The discussions were continued in a breakout session on product lines of safety critical systems. Other breakout sessions included dynamic product lines, generic representation of variability, and testing and modeling variability.

Overall, a core set of techniques were discussed at this seminar which include program analysis, model checking, type checking, and testing. We believe that the seminar fruitfully mixed computer science and software engineering researchers from several research sub-domains, allowing them to derive interesting basic research problems stemming from practical needs all related to how variability impacts their respective domains, with the sub-goal of inspiring the use of the latest research advances in software analysis technology to advance variability management tools.

■ Results

The different kinds of interactions offered by the seminar helped the participants to relate work covering different aspects in a number of dimensions such as:

1. An overall approach to thinking about variability, as defined by Thüm's classification [4] of analysis into product based, family based, feature based and hybrids;

2. Core techniques: testing, verification, refactoring, model checking, static analysis;
3. Mechanisms for representing variability: if-defs, deltas, generic representation, etc.;
4. Application domains;
5. The nature of variability: static product lines, dynamic product lines, configurable systems.

The seminar also produced a bibliography of core readings on the topic, that can enable new graduate students to engage more quickly in this area of research.

Trying to classify approaches with respect to these dimensions helped to identify similarities and differences among different techniques (static analysis, model checking, testing, and verification). This, in turn, might trigger new collaborations and research results. The presentations and the ad-hoc discussion sessions helped people to clarify differences and similarities among configurable systems and dynamic and static product lines, with similar consequences to the ones described above. More generally, of course, the Dagstuhl provided the benefit of mixing young and experienced researchers, from different countries and research areas.

An informal survey among a handful of participants has shown that each of them have started 2-3 new collaborations as a result of the seminar. These collaborations took the form of initiated research papers, mutual research visits, or student exchanges. In one anecdotal case, a researcher started a collaboration with a colleague sitting in the same corridor at his home university— but apparently one had to meet in Dagstuhl to enable the exchange of ideas. We can thus expect a new wave of research results in this area to flourish about a year from the seminar time. Because of this success, we intend to organize a follow up event in several years, be it under the Schloss Dagstuhl programme or under some other appropriate venue.

■ References

- 1 Kacper Bąk, Krzysztof Czarnecki, and Andrzej Wąsowski. Feature and Meta-Models in Clafer: Mixed, Specialized, and Coupled. In *Proc. of the 3rd Int'l Conf. on Software Language Engineering (SLE'10)*, LNCS, Vol. 6563, pp. 102–122, Springer, 2011. DOI: 10.1007/978-3-642-19440-5_7.
- 2 Benjamin Delaware, William R. Cook, and Don S. Batory. Product lines of theorems. In *Proc. of the 2011 ACM Int'l Conf. on Object-oriented Programming Systems, Languages, and Applications (OOPSLA'11)*, pp. 595–608, ACM, 2011. DOI: 10.1145/2048066.2048113.
- 3 Martin Erwig and Eric Walkingshaw. The Choice Calculus: A Representation for Software Variation. *ACM Trans. Softw. Eng. Methodol.*, Vol. 21, Issue 1, pp. 6:1–6:27, 2011. DOI: 10.1145/2063239.2063245.
- 4 Thomas Thüm, Sven Apel, Christian Kästner, Martin Kuhlemann, Ina Schaefer, and Gunter Saake. Analysis Strategies for Software Product Lines. Technical Report FIN-004-2012, School of Computer Science, University of Magdeburg, April 2012. http://www.cs.uni-magdeburg.de/inf_media/downloads/forschung/technical_reports_und_preprints/2012/04_2012.pdf.

4.15 Computational Geometry

Organizers: Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud

Seminar No. 13101

Date: March 3–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.3.1

© Creative Commons BY 3.0 Unported license

© Otfried Cheong, Kurt Mehlhorn, and Monique Teillaud



Participants: Pankaj Kumar Agarwal, Boris Aronov, Dominique Attali, Sang Won Bae, Mikhail Belkin, Eric Berberich, Kevin Buchin, Maïke Buchin, Siu-Wing Cheng, Otfried Cheong, Kenneth L. Clarkson, Tamal K. Dey, Anne Driemel, Alon Efrat, Jeff Erickson, Sándor Fekete, Joachim Giesen, Marc Glisse, Xavier Goaoc, Joachim Gudmundsson, Herman J. Haverkort, Menelaos Karavelas, Matthew J. Katz, Michael Kerber, Rolf Klein, Maarten Löffler, Kurt Mehlhorn, Guillaume Moroz, Dmitry Morozov, David M. Mount, Yoshio Okamoto, Marcel J. M. Roeloffzen, Günter Rote, Jörg-Rüdiger Sack, Ludmila Scharf, Lena Schlipf, Raimund Seidel, Micha Sharir, Bettina Speckmann, Fabian Stehn, Monique Teillaud, Hans Raj Tiwary, Suresh Venkatasubramanian, Antoine Vigneron, Yusu Wang, Carola Wenk, Nicola Wolpert

■ Computational Geometry and its Evolution

The field of computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric and topological 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 and topological algorithms. Several software libraries for geometric computation (e.g. LEDA, CGAL, CORE) have been developed. Remarkably, this emphasis on applications and implementations has emerged from the originally theoretically oriented computational geometry community itself, so many researchers are concerned now with theoretical foundations as well as implementations.

■ Seminar Topics

The seminar presented recent developments in the field and identified new challenges for computational geometry. Below we list some of the most interesting subareas of the field at this stage, covering both theoretical and practical issues in computational geometry.

- *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.

- *Geometric Computing* has 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.
- *Computational topology* concentrates on the properties of geometric objects that go beyond metric representation: modeling and reconstruction of surfaces, shape similarity and classification, and persistence are key concepts with applications in molecular biology, computer vision, and geometric databases.
- In its early years, computational geometry concentrated on low dimensions. *High-dimensional data* has become very important recently, in particular, in work related to machine learning and data analysis. Standard solutions suffer from the curse of dimensionality. This has led to extensive work on dimension-reduction and embedding techniques.
- 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, Voronoi diagrams and nearest-neighbor data structures under various metrics have turned out to be useful for many applications and are being investigated intensively.
- *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.

■ Participants

47 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive.

Dagstuhl seminars on computational geometry have been organized in a two year rhythm since a start in 1990. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together, fostering collaboration, and exposing young talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry.

No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at Dagstuhl. To accommodate

new, younger researchers, the organizers held a *lottery* for the first time this year. From an initial list of selected researchers, we randomly selected a certain number of senior, young, and female participants. Researchers on the initial list who were not selected by the lottery were notified by us separately per email, so that they knew that they were not forgotten, and to reassure them that—with better luck—they will have another chance in future seminars.

We believe that the lottery created space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar was much “younger” than in the past, and certainly more “family-friendly.” Five young children roaming the premises created an even cosier atmosphere than we are used in Dagstuhl. Without decreasing the quality of the seminar, we had a more balanced attendance than in the past. Feedback from both seminar participants and from researchers who were not selected was uniformly positive.

Dagstuhl itself is a great strength of the seminar. Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, we warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl!

4.16 Scheduling

Organizers: Susanne Albers, Onno J. Boxma, and Kirk Pruhs
Seminar No. 13111

Date: March 10–15, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.3.111
 © Creative Commons BY 3.0 Unported license
 © Susanne Albers, Onno J. Boxma, and Kirk Pruhs



Participants: Samuli Aalto, Susanne Albers, Konstantin Avrachenkov, Urtzi Ayesta, Yossi Azar, Nikhil Bansal, Vincenzo Bonifaci, Sem C. Borst, Onno J. Boxma, Marek Chrobak, Christoph Dürr, Fritz Eisenbrand, Thomas Erlebach, Naveen Garg, Anupam Gupta, John Hasenbein, Wiebke Höhn, Sungjin Im, Csana Imreh, Peter Jacko, Samir Khuller, Amit Kumar, Fei Li, Alberto Marchetti, Spaccamela, Monaldo Mastrolilli, Nicole Megow, Rolf H. Möhring, Benjamin J. Moseley, Kamesh Munagala, Viswanath Nagarajan, Seffi Naor, Sindo Nunez Queija, Kirk Pruhs, Rhonda Richter, Roman Rischke, Jiri Sgall, Hadas Shachnai, Devavrat Shah, René Sitters, Martin Skutella, Frits C. R. Spieksma, Mark S. Squillante, Clifford Stein, Sebastian Stiller, Alexander Stolyar, Leen Stougie, Maxim I. Sviridenko, Marc Uetz, Peter van de Ven, Suzanne van der Ster, Rob van Stee, Tjark Vredeveld, Gideon Weiss, Gerhard J. Woeginger, Prudence W. H. Wong

The primary objective of the seminar is to facilitate dialog and collaboration between researchers in two different mathematically-oriented scheduling research communities, the stochastic scheduling and queuing community, and the worst-case approximation scheduling community. To a large extent, the applications considered by the two communities are the same. The stochastic community considers questions related to determining stochastic information (like the expectation or tail bounds) about the performance of algorithms and systems from stochastic information about the input. The worst-case community considers questions related to determining the worst-case performance of algorithms and systems assuming no stochastic information about the input. Each community has developed its own set of mathematical techniques that are best suited to answer these different sorts of questions. While addressing similar problems, these communities tend to attend different conferences (e.g. SIGMETRICS vs. SODA/IPCO), and publish in different journals. Thus the organizers believed that each community would benefit from greater interaction with the other community, and this seminar was an opportunity to further such interaction. The seminar was attended by about 15 researchers from the stochastic community and 40 researchers from the worst-case community.

4.17 Bidimensional Structures: Algorithms, Combinatorics and Logic

Organizers: Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos

Seminar No. 13121

Date: March 17–22, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.3.51

© Creative Commons BY 3.0 Unported license

© Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos

Participants: Isolde Adler, Hans L. Bodlaender, Paul Bonsma, Jianer Chen, Rajesh Chitnis, Marek Cygan, Victor Dalmau, Anuj Dawar, Erik D. Demaine, Reinhard Diestel, Frederic Dorn, Fedor V. Fomin, Archontia Giannopoulou, Petr A. Golovach, Alexander Grigoriev, MohammadTaghi Hajiaghayi, Illya V. Hicks, Gwenaël Joret, Marcin Kaminski, Iyad A. Kanj, Mamadou Moustapha Kanté, Yusuke Kobayashi, Sudeshna Kolay, Guy Kortsarz, Stephan Kreutzer, O-joung Kwon, Daniel Lokshtanov, Johann A. Makowsky, Dániel Marx, Frédéric Mazoit, Bojan Mohar, Sang-il Oum, Geevarghese Philip, Marcin Pilipczuk, Michał Pilipczuk, Felix Reidl, Peter Rossmanith, Ignasi Sau Valls, Saket Saurabh, Konstantinos Stavropoulos, Blair D. Sullivan, Hisao Tamaki, Jan Arne Telle, Dimitrios M. Thilikos, Ioan Todinca, Erik Jan van Leeuwen, Yngve Villanger, Paul Wollan, Christian Wulff-Nilsen



The monumental Graph Minors project developed by Robertson and Seymour in the 1980s is one of the most fundamental achievements of Combinatorics. The project had several groundbreaking consequences for Theoretical Computer Science. However, the wide spread opinion in the algorithmic community, expressed by David S. Johnson in his NP-Completeness Column (J. Algorithms 1987), was that it is mainly of theoretical importance. It took some time to realize that the techniques developed in Graph Minors can be used in the design of efficient and generic algorithms. One of the main techniques extracted from Graph Minors is based on the structural results explaining the existence (or the absence) of certain grid-like or bidimensional structures in graphs. The usage of bidimensional structures and the related width parameters in many areas of Computer Science and Combinatorics makes such techniques ubiquitous.

Historically, the first applications of bidimensional structures are originated in Graph Minors of Robertson and Seymour, because of the structure of the graphs excluding some fixed some graph as a minor. There is still an on-going work in Combinatorics on obtaining new structural theorems. There are much more examples in Combinatorics, where bidimensional structures and width parameters play a crucial role like in obtaining Erdős-Pósa type of results. Reed used bidimensional structures to settle Erdős-Hajnal conjecture on near-bipartite graphs. Kawarabayashi and Reed used bidimensional structures to bound the size of a minimal counterexample to Hadwiger's conjecture. Demaine and Hajiaghayi optimized the original grid-exclusion theorem on H -minor free graphs.

The usage of bidimensional structures and width parameters in Algorithms goes back to the parameter of treewidth, introduced

in the Graph Minors series. Treewidth is now ubiquitous in algorithm design and expresses the degree of topological resemblance of a graph to the structure of a tree. Its algorithmic importance dates back in the early 90's to the powerful meta-algorithmic result of Courcelle asserting that all graph problems expressible in Monadic Second Order Logic can be solved in linear time on graphs of bounded treewidth. Bounded treewidth can be guaranteed by the exclusion of certain bidimensional structures. Intuitively, this exclusion is what enables the application of a series of classic algorithmic techniques (divide-and-conquer, dynamic programming, finite automata) for problems of certain descriptive complexity. This phenomenon was perhaps the first strong indication of the deep interleave between graph structure and logic in graph algorithms. However, a deeper understanding of it became more evident during the last decade and produced powerful meta-algorithmic techniques.

Apparently, graph-theoretic fundamentals emerging from the Graph Minors project developed by Robertson and Seymour, are used currently in several areas of Computer Science and Discrete Mathematics. Algorithmic fertilization of these ideas occurred mostly in the context of parameterized complexity and its foundational links to logic. The course of developing a structural algorithmic graph theory revealed strong connections between Graph Theory, Algorithms, Logic, and Computational Complexity and joined a rapidly developing community of researchers from Theoretical Computer Science and Discrete Mathematics.

Dagstuhl seminar 13121 brought together some of the most active researchers on this growing field.

4.18 Future Internet

Organizers: Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz
Seminar No. 13131

Date: March 24–27, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.3.75

© Creative Commons BY 3.0 Unported license

© Jon Crowcroft, Markus Fidler, Klara Nahrstedt, and Ralf Steinmetz



Participants: Zdravko Bozakov, Florin Ciucu, Jon Crowcroft, Ruben Cuevas-Rumin, Hermann de Meer, David Dietrich, Markus Fidler, Philip Brighten Godfrey, Christian Gross, David Hausheer, Markus Hofmann, Matthias Hollick, Tobias Hoßfeld, Brad Karp, Martin Karsten, Wolfgang Kellerer, Karl Klug, Paul J. Kühn, Jörg Liebeherr, Laurent Mathy, Martin Mauve, Michael Menth, Max Mühlhäuser, Paul Müller, Klara Nahrstedt, Panagiotis Papadimitriou, Rastin Pries, Ivica Rimac, Silvia Santini, Nadi Sarrar, Jonathan M. Smith, Ralf Steinmetz, Dominik Stingl, Phuoc Tran-Gia, Oliver P. Waldhorst, Klaus Wehrle, Michael Zink, Martina Zitterbart

The recent vision of the “Future Internet” attracts significant networking research and causes controversial debates on the actions to be taken. Clean-slate initiatives envision a fresh start that put fundamental principles of networking into question. Avoiding any constraints of the current Internet implementation, the ambition of the clean-slate approach is to understand and design the ‘right’ network architecture. Evolutionary approaches, on the other hand, seek incremental improvements, assuming that the Internet can –as in the past– be fixed to accommodate the changing needs of users and applications.

Numerous initiatives on the Future Internet, like FIND, GENI funded by the NSF, FIRE, 4WARD by the EU, and G-LAB by the BMBF, reflect the importance of the topic. Characteristic for numerous Future Internet initiatives is an experimental approach using testbed facilities such as the GENI or the G-Lab platform.

Challenges that are of central importance for the Future Internet fall into the following categories:

Network design: computer networks and the Internet obey certain architectural guidelines that reflect experience gained in the art of network design, such as layered reference models or the Internet end-to-end argument. While these principles are backed up by the success of the Internet, it has to be noted that the network exhibits major architectural restrictions, e.g., regarding mobility, security, and quality of service. Computer networking as a relatively young field of research can benefit significantly from architectural reconsiderations that are initiated by clean-slate initiatives. While today, network theory is largely descriptive, this Dagstuhl seminar investigated the potentialities of a prescriptive network theory, which could justify a methodical rule/equation-based approach for the design of future networks.

Virtualization: the virtualization paradigm revolutionized the use of computers and data centers, where the flexibility and mobility of virtual machines offers tremendous potential,

posing, however, significant new challenges for networking. On the other hand, the virtualization paradigm has already many applications in networking, e.g., in virtual private networks or overlay networks. Currently, virtualization finds its way into network components, e.g., routers, itself, where the decoupling of logical entities from the physical substrate enables major innovations, e.g., concurrent (possibly post-IP) networks, infrastructure as a service, redundant shadow configurations, in network management, and in energy efficiency. Furthermore, the provisioning of service-oriented virtual networks across multiple infrastructure providers creates the need for separation between the network operations and the physical infrastructure. This is expected to change the way that virtual networks are managed, debugged, and operated. The Dagstuhl seminar contrasted different approaches to network virtualization and investigated their applications.

Experimental research: the Internet standardization process relies on running code and real world verification. An essential prerequisite for the transfer of research results is building of large scale testbed networks. These are frequently implemented as virtual, Software Defined Networks that run concurrently to a production network using the same hardware. The Dagstuhl seminar revisited the experimental approach and gathered lessons learned and best practices.

During the seminar, we discussed and (partly) answered the following questions:

Is a prescriptive network theory feasible?

Today, network research is largely descriptive, e.g., there exist methods and tools to model communication networks and protocols, to analyze their performance, or to verify their correctness. The design of new networks, however, lacks a prescriptive network theory that provides necessary rules and equations that specify how a network for a given purpose

has to be built. Instead, network design relies heavily on previous experience and best practices frequently resulting in incremental works. In contrast, the clean-slate Future Internet approach seeks to build a new Internet architecture from scratch. In this case the design space is entirely open requiring decisions regarding functional and non-functional aspects, e.g.,

- Where to implement reliable/unreliable and connection-less/connection-oriented?
- Where (end systems or network) and in which layer to keep state information?
- Where and how to achieve security, quality of service, and dependability?
- How to split locators and identifiers?

Given the examples above, we discussed:

- How can a prescriptive approach to network theory be formulated?
- What are the perspectives and the fundamental limits of the candidate approaches?
- What are the prospects of the approach if successful?

Which insights can the experimental, testbed-based approach reveal?

Many approaches to the Future Internet are experimentally driven and centered around a testbed that ideally if successful becomes the first running prototype of the Future Internet. Clearly, testbeds are indispensable to implement running code as a proof-of-concept, whereas their use for understanding networking and for establishing new principles and paradigms can be debated. In the seminar we elaborated on this question to provide answers to:

- Which insights can be expected?
- Which exemplary fundamental insights did emerge from testbeds?
- For which use cases are testbeds meaningful, e.g., to engineer details, to approach concepts weakly understood, to understand the impact of users, etc.?
- How should a testbed platform look like, which properties must be provided?
- How can testbeds be benchmarked to achieve comparability and validity?

What are the challenges for wide-area service-oriented virtual networks?

The virtualization paradigm gained a lot of attention in networking as it provides numerous useful applications and promises to solve a number of important issues, such as the gradual deployment of new networking solutions in parallel to a running production network. Considering existing networking technologies, it becomes apparent that virtual networks and virtual network components are already being used in a multitude of different ways and in different layers, e.g., Virtual LANs (VLANs), Virtual Private Networks (VPNs), the Virtual Router Redundancy Protocol (VRRP), or in form of overlay networks to name a few. Furthermore, the abundance of resources offered by commodity hardware can turn it into a powerful and highly programmable platform for packet processing and forwarding. The virtualization of such programmable network elements can provide network slices which are highly customized for particular network services and applications. The topics that were discussed at the seminar include:

- Resource discovery and provisioning of virtual networks across multiple domains, given that infrastructure providers will not be willing to expose their topology, resource information and peering relationships to third-parties;
- Virtualization of network components (e.g., resource allocation, isolation issues);
- Scaling of virtual resources to meet variable service demand;
- Use cases of network virtualization.

The seminar comprised a one minute madness session for introduction and for statements on the Future Internet, a breakout session for group work on the topic of prescriptive network theory, as well as podium discussions on experimentally driven research and on the use cases of SDN.

We would like to thank all presenters, scribes, and participants for their contributions and lively discussions. Particular thanks go to the team of Schloss Dagstuhl for their excellent organization and support. We also would like to thank Anil Madhavapeddy for his feedback and comments on SDN.

4.19 Formal Verification of Distributed Algorithms

Organizers: Bernadette Charron-Bost, Stephan Merz, Andrey Rybalchenko, and Josef Widder
Seminar No. 13141

Date: April 1–5, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.4.1

© Creative Commons BY 3.0 Unported license

© Bernadette Charron-Bost, Stephan Merz, Andrey Rybalchenko, and Josef Widder



Participants: Béatrice Bérard, Péter Bokor, Borzoo Bonakdarpour, Pierre Castéran, Bernadette Charron-Bost, Marie Duflot, Cormac Flanagan, Matthias Függer, Alexey Gotsman, Serge Haddad, Gerwin Klein, Igor Konnov, Fabrice Kordon, Akash Lal, Victor Luchangco, Stephan Merz, Uwe Nestmann, Thomas Nowak, Eric Ruppert, John Rushby, Andrey Rybalchenko, André Schiper, Klaus Schneider, Philippe Schnoebelen, Wilfried Steiner, Murali Talupur, Serdar Tasiran, Helmut Veith, Christoph Weidenbach, Jennifer L. Welch, Josef Widder, Karsten Wolf

While today's society depends heavily on the correct functioning of distributed computing systems, the current approach to designing and implementing them is still error prone. This is because there is a methodological gap between the theory of distributed computing and the practice of designing and verifying the correctness of reliable distributed systems. We believe that there are two major reasons for this gap: On the one hand, distributed computing models are traditionally represented mainly in natural language, and algorithms are described in pseudo code. The classical approach to distributed algorithms is thus informal, and it is not always clear under which circumstances a given distributed algorithm actually is correct. On the other hand, distributed algorithms are designed to overcome non-determinism due to issues that are not within the control of the distributed algorithm, including the system's timing behavior or faults of some components. Such issues lead to a huge executions space which is the major obstacle in applying verification tools to distributed algorithms.

The rationale behind our Dagstuhl seminar was that closing the methodological gap requires collaboration from researchers from distributed algorithms and formal verification. In order to spur the interaction between the communities, the program contained the following overview talks on the related subjects:

- *Distributed algorithms* – Eric Ruppert (York University)
- *Semi-automated proofs* – John Rushby (SRI)
- *Parameterized model checking* – Muralidhar Talupur (Intel)

In addition to the tutorials, we organized several open discussion rounds. The seminar participants identified modeling issues as a central question, which confirmed one of our motivation for the seminar, namely, the lack of a universal model for

distributed algorithms. Hence, one of the discussion rounds was exclusively devoted to this topic. Unlike sequential programs, whose semantics is well understood and closely follows the program text, the executions of distributed algorithms are to a large extent determined by the environment, including issues such as the distribution of processes, timing behavior, inter-process communication, and component faults. Models of distributed algorithms and systems embody different assumptions about how the environment behaves. These hypotheses are often left implicit but are of crucial importance for assessing the correctness of distributed algorithms. The discussions during the seminar raised the awareness of these issue among the researchers, and showed that research in this area is a necessary first step towards a structured approach to formal verification of distributed algorithms. In addition to modeling, we discussed issues such as benchmarks, implementation of distributed algorithms, or application areas of distributed algorithms.

To round-off the technical program, we had several short presentations by participants who presented their past and current work in the intersection of formal methods and distributed algorithms, and a joint session with the other seminar going on concurrently at Dagstuhl on Correct and Efficient Accelerator Programming. The topics of the talks spanned large parts of the concerned areas, for instance, there were talks on model checking techniques such as partial order reductions or abstractions, and their applications to distributed algorithms; several talks focused on proof assistants, and how they can be used to verify distributed algorithms; some talks considered concurrent systems, and some focused on transactional memory. The atmosphere during these sessions was very constructive, and the short talks were always followed by elaborate and insightful discussions.

4.20 Correct and Efficient Accelerator Programming

Organizers: Albert Cohen, Alastair F. Donaldson, Marieke Huisman, and Joost-Pieter Katoen
Seminar No. 13142

Date: April 1–4, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.4.17

© Creative Commons BY 3.0 Unported license

© Albert Cohen, Alastair F. Donaldson, Marieke Huisman, and Joost-Pieter Katoen

Participants: Jade Alglave, Adam Betts, Albert Cohen, Christian Dehnert, Dino Distefano, Alastair F. Donaldson, Jeremy Dubreil, Benoit Dupont de Dinechin, Ganesh L. Gopalakrishnan, Sebastian Hack, Lee Howes, Marieke Huisman, Christina Jansen, Joost-Pieter Katoen, Jeroen Ketema, Alexander Knapp, Georgia Kouveli, Alexey Kravets, Anton Lokhmotov, Roland Meyer, Cedric Nugteren, Zvonimir Rakamaric, Oliver Reiche, Philipp Rümmer, Ana Lucia Varbanescu, Sven Verdoolaege, Jules Villard, Heike Wehrheim, Anton Wijs, Marina Zaharieva-Stojanovski, Dong Ping Zhang



In recent years, massively parallel accelerator processors, primarily GPUs, have become widely available to end-users. Accelerators offer tremendous compute power at a low cost, and tasks such as media processing, simulation and eye-tracking can be accelerated to beat CPU performance by orders of magnitude. Performance is gained in energy efficiency and execution speed, allowing intensive media processing software to run in low-power consumer devices. Accelerators present a serious challenge for software developers. A system may contain one or more of the plethora of accelerators on the market, with many more products anticipated in the immediate future. Applications must exhibit portable correctness, operating correctly on any configuration of accelerators, and portable performance, exploiting processing power and energy efficiency offered by a wide range of devices.

The seminar focussed on the following areas:

- Novel and attractive methods for constructing system-independent accelerator programs;
- Advanced code generation techniques to produce highly optimised system-specific code from system-independent programs;
- Scalable static techniques for analysing system-independent and system-specific accelerator programs both qualitatively and quantitatively.

The seminar featured five tutorials providing an overview of the landscape of accelerator programming:

- Architecture – Anton Lokhmotov, ARM
- Programming models – Lee Howes, AMD
- Compilation techniques – Sebastian Hack, Saarland University
- Verification – Ganesh Gopalakrishnan, University of Utah
- Memory models – Jade Alglave, University College London

In addition, there were short presentations from 12 participants describing recent results or work-in-progress in these areas, and two discussion sessions:

- Domain specific languages for accelerators;
- Verification techniques for GPU-accelerated software.

Due to the “correctness” aspect of this seminar, there was significant overlap of interest with a full week seminar on *Formal Verification of Distributed Algorithms* running in parallel. To take advantage of this overlap a joint session was organised, featuring a talk on verification of GPU kernels by Alastair Donaldson, Imperial College London (on behalf of the *Correct and Efficient Accelerator Programming* seminar) and a talk on GPU-accelerated runtime verification by Borzoo Bonakdarpour, University of Waterloo, on behalf of the *Formal Verification of Distributed Algorithms* seminar.

4.21 Drawing Graphs and Maps with Curves

Organizers: Sara Fabrikant, Stephen G. Kobourov, Martin Nöllenburg, and Monique Teillaud
Seminar No. 13151

Date: April 7–12, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.4.34

© Creative Commons BY 3.0 Unported license

© Stephen G. Kobourov, Martin Nöllenburg, and Monique Teillaud



Participants: Patrizio Angelini, Michael Bekos, David Eppstein, Fabrizio Frati, Eric Fusy, Martin Gronemann, Jan-Henrik Haurert, Herman J. Haverkort, Michael Hemmer, Danny Holten, Michael Kaufmann, Stephen G. Kobourov, Sylvain Lazard, Maarten Löffler, Tamara Mchedlidze, Wouter Meulemans, Lev Nachmanson, Benjamin Niedermann, Arlind Nocaj, Martin Nöllenburg, Sergey Pupyrev, Helen C. Purchase, Andreas Reimer, Maxwell J. Roberts, Günter Rote, André Schulz, Aidan Slingsby, Bettina Speckmann, Monique Teillaud, Torsten Ueckerdt, Kevin Verbeek, Alexander Wolff, Jo Wood, Kai Xu

Graphs and networks, maps and schematic map representations are frequently used in many fields of science, humanities and the arts. The need for effective visualization and aesthetically pleasing design is attested by the numerous conferences and symposia on related topics, and a history that is several centuries old. From Mercator's maps dating to the 1500's, to interactive services such as Google Earth, geography and cartography have generated and solved many theoretical and practical problems in displaying spatial data effectively and efficiently. From Euler's visualization of the bridges of Königsberg in the 1700's, to Facebook's social networks, graph drawing has also proven a fertile area for theoretical and practical work. More recent is the notion of highly schematized maps and graphs, with the classic examples of statistical value-by-area cartograms by Raisz and Henry Beck's London Tube map, both dating back to the 1930's.

A key challenge in graph and cartographic visualization is designing cognitively useful spatial mappings of the underlying data that allow people to intuitively understand the displayed information. Such work draws on the intellectual history of several traditions, including information visualization, human-computer interaction, psychology, cognitive science, graphic design, cartography, and art. The synthesis of relevant ideas from these fields with new techniques can lead to new and better visualizations to help us keep pace with the torrents of data confronting us.

Although a great deal is known, both in theory and in practice, about drawing graphs and maps with straight-line segments, there are few corresponding results about circular-arc drawings in particular, and curve drawings in general. The use of circular arcs in place of straight-line segments opens a new chapter in drawing graphs and maps from both theoretical and practical points of view. Specifically, we are interested in the interplay between practical requirements of drawing with curves, arising in cartography and GIS, and theoretical results in computational

geometry and graph drawing. Such work is motivated by perception research which indicates that representing paths with smooth geodesic trajectories aids in comprehension, as well as by the aesthetic appeal of smooth curves; see Fig. 4.3 and Fig. 4.4.

■ Aims of the Seminar

The main goal of this seminar was to bring together researchers with interests in drawing graphs and maps with curves coming from information visualization, psychology, cognitive science, human-computer interaction, graph drawing, computational geometry, cartography, and GIS. It follows in a tradition of several previous similarly structured Dagstuhl seminars on graph drawing and map visualization. From April 7th to April 12th a group of 34 junior and senior researchers from eight different countries gathered in Dagstuhl. Being a *small* seminar with a target participation of 30 persons, the seminar was fully booked, which shows that this seemingly narrow topic still raises a lot of interest in the different communities. We all came together to discuss open research questions and engage in new collaborations around visualizations that replace straight lines with circular arcs and curves. This topic opens a great deal of theoretical and practical possibilities and with this in mind, the specific aims of the Dagstuhl seminar were:

- To learn about the state of the art of the use of curves in the different research areas. We invited a small number of survey lectures to define a common ground for interdisciplinary work.
- To organize an exhibition of art and visual designs on the common theme of curves contributed by participants and artists, and use this to stimulate discussion.
- To identify specific theoretical and practical open problems that need to be solved in order to make it possible to draw graphs and maps with circular arcs and curves.




Fig. 4.3
Lombardi graph drawings [1]. Brinkman graph, Dyck graph, and F_{40} (dodecahedron double cover).

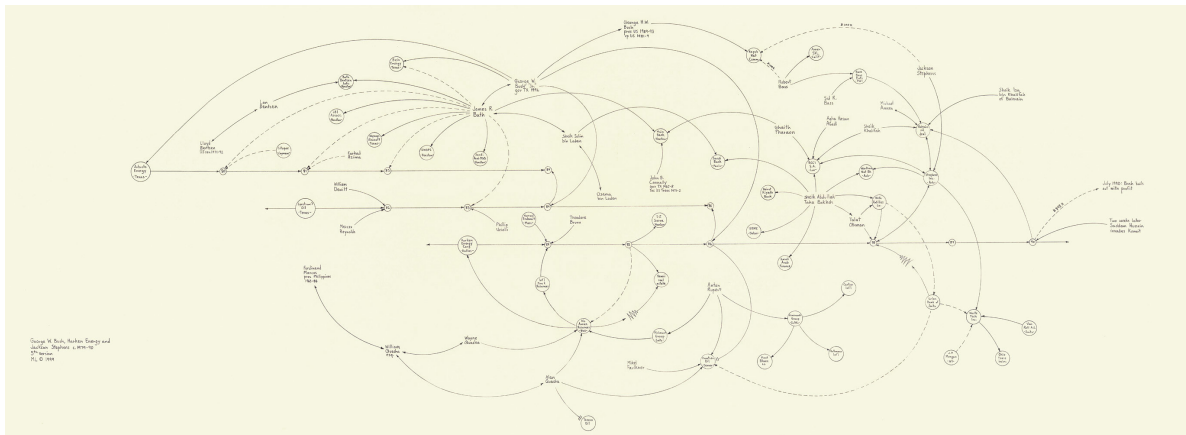


Fig. 4.4
One of Mark Lombardi's pieces. *George W. Bush, Harken Energy, and Jackson Stevens ca. 1979–90, 1999.* Graphite on paper, 20 × 44 inches. Courtesy Pierogi Gallery and Donald Lombardi. Photo credit: John Berens.

- To form smaller working groups around some of the identified problems and to initiate a collaborative research process for finding answers and solutions to these problems.
- To report about the progress made in the working groups in a plenary session for getting feedback and further input from members of the other groups.
- To continue the joint research efforts beyond the seminar week and eventually publish those results.

■ Achievements of the Seminar

The achievements in the seminar were numerous and varied. The subsequent chapters of this report summarize the more important ones.

1. On Monday and Tuesday, we enjoyed seven survey lectures. David Eppstein opened with a broad overview of the use of curves in visualization of graphs and networks. Günter Rote talked about algorithms for approximating polygonal curves by simpler curves and sequences of biarcs. Sylvain Lazard illustrated connections with algebra and geometry when dealing with curves. Jo Wood surveyed the use of curves in cartography and information visualization. Helen Purchase discussed perception theories and empirical studies on the use of curves in visualization, and Maxwell Roberts discussed the question whether curvilinear metro maps have cognitive benefits over traditional straight-line schematic maps. Finally, Monique Teillaud and Michael Hemmer overviewed the

history of the open source project CGAL, the Computational Geometry Algorithms Library, and then discussed specific CGAL packages that are relevant for drawing circular arcs and smooth algebraic curves. Beyond the survey and review talks, we also heard a presentation by Wouter Meulemans about the use of curved schematization of geometric shapes, where the results were obtained via a user study of the participants in the seminar.

2. We also had two short impromptu presentations and software demos. In particular, Günter Rote presented an ipelet to transform polygons into splines in the drawing editor ipe. Jan-Henrik Hauernt reported about work in progress and showed a demo on morphing polygonal lines so that edge lengths and angles behave as consistently as possible over time.
3. A number of relevant open problems were formulated early in the seminar and six working groups formed around some of the problems. The groups then worked by themselves, formalizing and solving their specific theoretical and practical challenges. Below is a list of the working group topics.
 - a. **Smooth Orthogonal Drawings:** What is the complexity of recognizing whether a given 4-planar graph admits a smooth orthogonal drawing of edge complexity 1?
 - b. **Confluent Drawing:** What is the complexity of determining whether a given graph has a so-called *strict* confluent drawing?
 - c. **Automated Evaluation of Metro Map Usability:** What

are good, objective, quantifiable criteria by which curvilinear metro maps can be evaluated? Can such criteria be used so that linear maps can likewise be compared both with each other and also with curvilinear maps?

- d. **Universal Point Sets for Planar Graph Drawings with Circular Arcs:** What can be said about universal point sets for drawing planar graphs if curves are used instead of straight-line segments?
 - e. **Labeling Curves with Curved Labels:** How can points on a smooth curve be labeled automatically using curved labels?
 - f. **Graphs with Circular Arc Contact Representation:** Which graphs can be represented by contacts of circular arcs?
4. We had an excellent exhibition entitled “Bending Reality: Where Arc and Science Meet”. This exhibition is the third one in a series of exhibitions that accompany Dagstuhl seminars where aesthetics and art are naturally part of the scientific topics. It was on display from April 8 to April 21, 2013. Moreover, for the first time in Dagstuhl history, this exhibition is made permanently available as a virtual exhibition that can be accessed at <http://www.dagstuhl.de/ueber-dagstuhl/kunst/13151>.

The last three days of the seminar were dedicated to working group efforts. Several of the groups kept their focus on the original problems as stated in the open problem session, while other groups modified and expanded the problems. On the last day of the seminar we heard progress reports from all groups. The results of two of the groups have recently been accepted to international conferences, and we are expecting further research publications to result directly from the seminar.

Arguably the best, and most-appreciated, feature of the seminar was the opportunity to engage in discussion and interactions with experts in various fields with shared passion about curves. The aforementioned exhibition “Bending Reality” helped make the topics of the seminar more visible and raised new questions. In summary, we regard the seminar as a great success. From the positive feedback that we got it is our impression that the participants enjoyed the unique scientific atmosphere at Schloss Dagstuhl and profited from the scientific program. We are grateful for having had the opportunity to organize this seminar and thank the scientific, administrative, and technical staff at Schloss Dagstuhl.

■ References

1 Christian A. Duncan, David Eppstein, Michael T. Goodrich, Stephen G. Kobourov, and Martin Nöllenburg.

Lombardi drawings of graphs. *J. Graph Algorithms and Applications*, 16(1):85–108, 2012.

4.22 Interface of Computation, Game Theory, and Economics

Organizers: Sergiu Hart, Éva Tardos, and Bernhard von Stengel
Seminar No. 13161

Date: April 14–19, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.4.666
 © Creative Commons BY 3.0 Unported license
 © Sergiu Hart, Éva Tardos, and Bernhard von Stengel

Participants: Yakov Babichenko, Simina Branzei, Markus Brill, Jing Chen, Giorgos Christodoulou, Constantinos Daskalakis, Gabrielle Demange, Nikhil R. Devanur, Shahar Dobzinski, Paul Dütting, Edith Elkind, Michal Feldman, Felix Fischer, Paul W. Goldberg, Yannai A. Gonczarowski, Sergiu Hart, Jason D. Hartline, Penelope Hernandez Rojas, Martin Hoefer, Nicole Immorlica, Ramesh Johari, Thomas Kesselheim, Max Klimm, Elias Koutsoupias, Kevin Leyton-Brown, Katrina Ligett, Brendan Lucier, Jeffrey MacKie-Mason, Vangelis Markakis, Vahab Mirrokni, Hervé Moulin, Rudolf Müller, Sigal Oren, Mallesh Pai, Dimitrii V. Pasechnik, Aaron Roth, Tim Roughgarden, William H. Sandholm, Rahul Savani, Guido Schäfer, Michael Schapira, Ilya R. Segal, Vasilis Syrgkanis, Éva Tardos, Berthold Vöcking, Rakesh V. Vohra, Bernhard von Stengel, Jens Witkowski



The aim of this seminar was to study research issues at the interface of computing, game theory and economics. It facilitated discussions among people working in different disciplines. The majority of participants were academics from computer science departments, and the others (about one third) from other disciplines such as economics or corporate research departments of Google or Microsoft. All have strong cross-disciplinary interests.

Economic transactions on the internet are of ever-increasing importance. In order to execute and support them algorithmically, it is important to understand the agents' incentives on one hand and computational constraints on the other hand. This is studied in approaches to mechanism design and auctions, which formed a large part of the topics of this workshop.

Theoretical and practical issues of *mechanism design* were topics of the following presentations: epistemic implementations with belief levels (Jing Chen), translating agent-provided inputs to optimization (Constantinos Daskalakis), reward schemes (Shahar Dobzinski), the difficulties of allocating more than one good (Sergiu Hart), advertisement exchanges (Vahab Mirrokni), mechanisms for the private supply of a public good (Rudolf Müller), truthfulness versus privacy (Aaron Roth), composing mechanisms (Vasilis Syrgkanis), and allocating indivisible objects (Rakesh V. Vohra).

Aspects of *auctions* concerned “expressiveness” about preferences (Paul Dütting), the approximate optimality of marginal revenue maximization (Jason D. Hartline), improving the design of online advertising auctions (Kevin Leyton-Brown), commitment (Katrina Ligett), inefficiency of multi-unit auctions (Vangelis Markakis), symmetric auctions (Mallesh Pai), interdependent values (Tim Roughgarden), and spectrum auctions (Ilya Segal).

Understanding the interconnectedness of complex economic systems requires models and theories for the underlying *network* structures and their dynamics. Networks were studied with respect to social segregation (Nicole Immorlica), practical market

applications (Ramesh Johari), online creation (Thomas Kesselheim), competition (Brendan Lucier), and social contagion (Sigal Oren).

Social models, with bridges to mechanism design, were studied in presentations on division protocols (Simina Branzei), randomized social choice (Markus Brill), ranking methods (Gabrielle Demange), power changes in voting games (Edith Elkind), and incentives beyond selfishness (Guido Schäfer).

Achieving and computing an equilibrium in *dynamic models of large interactions* such as games and market models was studied for large aggregative games (Yakov Babichenko), new price updating in markets (Nikhil R. Devanur), payoff queries for games (Paul W. Goldberg), limit processes for evolutionary games (Bill Sandholm), and tournament competitions (Bernhard von Stengel).

The topics were chosen by the presenters, not by the organizers. The rather strong emphasis on mechanism design and auctions (which may have caused one single critical feedback comment on “too much groupthink”) reflects the strong current interest in this area, in line with its economic importance, for example as the source of the riches of Google and other internet search engines.

4.23 Pointer Analysis

Organizers: Ondřej Lhoták, Yannis Smaragdakis, and Manu Sridharan
Seminar No. 13162

Date: April 14–19, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.4.91

© Creative Commons BY 3.0 Unported license

© Ondřej Lhoták, Yannis Smaragdakis, Manu Sridharan



The Dagstuhl seminar on *Pointer Analysis* brought together experts in pointer analysis and researchers building demanding clients of pointer analysis, with the goal of disseminating recent results and identifying important future directions. The seminar was a great success, with high-quality talks, plenty of interesting discussions, and illuminating breakout sessions.

■ Research Context

Pointer analysis is one of the most fundamental static program analyses, on which virtually all others are built. It consists of computing an abstraction of which heap objects a program variable or expression can refer to. Due to its importance, a large body of work exists on pointer analysis, and many researchers continue to study and develop new variants. Pointer analyses can vary along many axes, such as desired precision, handling of particular language features, and implementation data structures and optimizations. Given the subtle implications of these design choices, and the importance of low-level details often excluded from conference-length papers, it can be difficult even for pointer analysis experts to understand the relationship between different analysis variants. For a non-expert aiming to use pointer analysis in a higher-level client (for verification, optimization, refactoring, etc.), choosing the right analysis variant can be truly daunting.

Pointer analysis is a mature area with a wealth of research results, at a temptingly close distance from wide practical applicability, but not there yet. The breakout application of precise analysis algorithms has seemed to be around the corner for the past decade. Although research ideas are implemented and even deployed in limited settings, several caveats always remain. These include assumptions about client analyses (i.e., the pointer analysis algorithm is valid only under assumptions of how the information will be used), assumptions about the analyzed program (e.g., that some language features are absent or that

Participants: Jose Nelson Amaral, Gogul Balakrishnan, Eric Bodden, Bor-Yuh Evan Chang, Isil Dillig, Thomas Dillig, Julian Dolby, Samuel Z. Guyer, Christian Hammer, Laurie J. Hendren, Uday Khedker, Ondrej Lhotak, Benjamin Livshits, Welf Löwe, Mark Marron, Matt Might, Ana Milanova, Anders Moeller, Mayur Naik, Hakjoo Oh, Erhard Plödereder, Xavier Rival, Yannis Smaragdakis, Gregor Snelting, Manu Sridharan, Bjarne Steensgaard, Dimitris Vardoulakis

their presence does not affect the analysis outcome), assumptions about modularity (e.g., that the code to be analyzed constitutes the whole program), etc. The right engineering packaging of pointer analysis algorithms as well as a convenient characterization of their domain of applicability are still elusive.

In this light, the seminar aimed to emphasize the relationship of pointer analysis algorithms with client analyses, as well as practical deployment issues. The seminar brought together researchers working on pointer analysis for various programming languages with researchers working on key analysis clients. Our main goals were (1) to deepen understanding of the relationships between existing pointer analysis techniques, and (2) to gain a better understanding of what pointer analysis improvements are required by clients, thereby setting an exciting agenda for the area going forward.

■ Seminar Format

Our seminar employed a somewhat unusual format for participant talks, intended to encourage a deeper discussion of each participant’s work. Each participant was allotted a 40-minute slot to present their work, consisting of 20 minutes of presentation and 20 minutes of discussion. The presentation and discussion times in each slot were enforced using a chess clock: when a question arose during a talk, the clock was “flipped” to discussion time, and after the discussion, it was flipped back to speaker time. (The times were not very strictly enforced; in some cases, the audience would “donate” time to the speaker to complete his/her presentation.) This format had two key benefits:

- It enabled discussion to freely occur during the talk, removing the worry that the speaker would have no time left to complete his/her presentation.
- It encouraged the audience to ask more questions, in order to “use up” the allotted audience time.

Overall, the format was very successful in encouraging good discussion, and most participants enjoyed it.

In addition to talks, we held four 90-minute breakout sessions. The session topics were proposed by participants before and during the seminar and voted on by participants. The sessions were scheduled two at a time, and participants could choose which session to attend. The discussions held in these sessions were quite illuminating. Finally, the last half-day of the seminar was spent on additional discussion of the breakout session topics, and on an initial effort to collectively improve the Wikipedia article on pointer analysis.¹⁰

■ Seminar Results

Recent advancements in pointer analysis have come from several different directions:

- Formulations (CFL, Datalog)—highly-complex analyses have been specified in terms of concise specifications, by utilizing declarative notations.
- Greater precision—interesting analyses that maintain finer-grained abstractions while maintaining scalability have been invented.
- Optimizations—data structures such as BDDs have been used to make complex analyses feasible.
- Demand-driven, refinement—the analysis problem has been specialized effectively when pointer information only needs to be computed for select program sites.
- Partial programs—analyses have been formulated to work without fully analyzing all libraries, or even all application code.

Such advances were discussed in detail during many participant talks in the seminar, and in the breakout sessions.

Recent work in pointer analysis has been driven by new clients for the analysis and by new programming languages. Along with ongoing use of pointer analysis in traditional optimizing compilers, recent years have seen many other clients emerge that require effective pointer analysis, e.g., in the areas of program verification and bug finding, refactoring, and security. These clients were well-represented by seminar attendees, who gave many interesting talks on novel uses of pointer analysis (particularly in the security domain). The rich exchanges between researchers building novel clients and those with pointer analysis expertise were one of the most valuable aspects of the seminar. Additionally, one breakout session covered the difficulties in designing an effective general pointer-analysis API that is suitable for a wide variety of clients.

Mainstream programming has been transitioning to increasingly heap-intensive languages: from C-like languages to object-oriented languages like Java and C#, and more recently to script-

ing languages like JavaScript and Ruby. As languages become more heap-intensive, the need for effective pointer analysis is greater, motivating continuing work in this area. The seminar talks covered a wide and diverse set of languages, each with its own considerations. A few talks covered pointer analysis for higher-level languages such as JavaScript and MATLAB. Such languages are becoming increasingly popular, and they are very heap-intensive compared to C-like languages, motivating the need for better pointer analysis. A couple of talks presented techniques for control-flow analysis of functional languages like Scheme. While the pointer analysis and control-flow analysis communities often use similar techniques, the relationships between the techniques is often obscured by differing terminology and presentation styles. The presentations on control-flow analysis and the corresponding discussions were helpful in bridging this gap.

The seminar included a good deal of discussion on practical issues with pointer analysis, including evaluation methodologies and issues arising in real-world deployments. A key theme that arose from these discussions was the need for pointer analysis to be at least partially unsound to be useful in practice, and how this need for unsoundness has not been explained properly in the literature. Analyses that made soundness compromises for practicality were deemed “soundy,” a tongue-in-cheek term that caught on quickly among participants. Recently, some seminar participants presented a well-received PLDI Fun and Interesting Topics (FIT) talk on the notion of “soundiness,” and several participants have agreed to collectively co-author a publishable document on the topic.

■ Conclusions

Overall, the *Pointer Analysis* Dagstuhl seminar was a great success. The seminar brought together 27 researchers from both academia and industry (including Google, IBM, Microsoft, NEC), with a good mix of junior and senior researchers. There were many interesting talks, with deep discussion facilitated by the chess clock time maintenance. The seminar facilitated interaction between pointer analysis experts and researchers building novel clients (a key goal for the seminar from the beginning), and also between researchers working on analyses for a variety of languages. Breakout sessions enabled further discussion of certain particularly interesting topics. In particular, there were invaluable discussions of many practical issues that often get short shrift in conference papers. These discussions sparked the notion of “soundiness,” which may have broader impact via a future publication.

¹⁰ See http://en.wikipedia.org/wiki/Pointer_analysis.

4.24 Customizing Service Platforms

Organizers: Luciano Baresi, Andreas Rummler, and Klaus Schmid
Seminar No. 13171

Date: April 21–26, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.4.114
 © Creative Commons BY 3.0 Unported license
 © Luciano Baresi, Andreas Rummler, and Klaus Schmid



Participants: Marco Aiello, Luciano Baresi, Karina Barreto Villela, Deepak Dhungana, Peter Dolog, Schahram Dustdar, Holger Eichelberger, Gregor Engels, Sam Guinea, Waldemar Hummer, Christian Inzinger, Patricia Lago, Grace A. Lewis, Georg Leyh, Tiziana Margaria, Nenad Medvidovic, Nanjangud C. Narendra, Leonardo Passos, Cesare Pautasso, Manuel Resinas Arias de Reyna, Florian Rosenberg, Antonio Ruiz Cortés, Andreas Rummler, Klaus Schmid, Jacek Serafinski, Damian Andrew Tamburri, Frank van der Linden, Wenjun Wu, Uwe Zdun

■ Background

Service-orientation has become a major trend in computer science over the last decade. More recently cloud computing is leading into the same direction: a virtualization of resources and service offerings. Especially cloud computing is getting very significant attention by companies. While the initial idea in service orientation was to have the relevant services standardized and distributed across the internet, we also see that an increasing amount of customization must be done to really meet customer needs. As in traditional system development, one size fits all is not enough.

This seminar focused on the notion of service platforms, a concept including, but not limited to, cloud computing. A service platform is a combination of technical infrastructure along with domain-specific or business-specific services built according to the service-oriented development paradigm. Especially the latter in practice often requires significant customization in order to be practically useful. Providing such customizations on a massive scale cost-effectively is an extremely demanding task. This is a lesson that has been learned hard by a number of companies in traditional software engineering. As a consequence the concept of product line engineering was conceived.

The focus of this seminar was to explore the range of different approaches towards customized service offerings in current – and future – service-based environments. In particular, it was a goal to address the potential for a combination of service-orientation with product line engineering ideas. In this regard, this seminar was the first of its kind.

■ Diversity of Topics

The expected diversity of inputs that was desired for the seminar was well achieved. This is shown by the diversity of individual presentations. Also the working groups that were

established had participants from multiple communities. These working groups discussed the following topics:

Quality Assurance and Validation in the Context of Customization:

Here, a broad range of different problems and techniques could be identified, related both to problems of varying of the object of the quality assurance as well as to the variation of the expectations (qualities).

Mobility Devices and Customization: This working group focused particularly on the difficulties that arise from a mobile context with a lot of variation over time and limited resources.

Architecting for Platform Customization: Architectures are fundamental to any software system, so this group addressed what architectural techniques are important to create customizable platforms.

Energy-Aware Customization: Here, the focus was on the issue of energy-awareness and, in particular, energy-efficiency, which is particularly relevant to mobile platforms. By adequate customization, this can be improved for a platform.

Customizing Service Platforms for Cloud Computing:

Modern cloud computing environments pose new challenges and provide new opportunities for customizing service platforms. It turned out that the cloud context provides a number of very special problems and technologies for addressing them.

Customizing Service Platforms for Agile Networked Organizations:

The organizational context of service platform needs to be taken into account as well as a platform needs to fit to the relevant business context. Hence customization needs to be done on both levels in a synchronized manner.

Binding time aspects of service platform customization:

This working group focused on when (i.e., in which lifecycle phase) the customization is done, as this has significant impact on the details of the technologies that can be used.

■ Reflections on the Format

A main goal of the seminar was to have a significant portion of the time for discussion. In order to achieve this, we decided to not require presentations from everyone associated with a long introduction round. Rather, we decided to ask everyone for a poster to present her- or himself and describe the personal interest

and relation to the topic. Overall this novel approach was well received by the participants. The poster walls were set up in the coffee break area outside the room. (Thanks to everyone at Dagstuhl for their support.) This allowed for a casual browsing of the posters in every coffee break during the seminar. Each poster also had a picture of the participant, this also helped to get to know each other.

4.25 VaToMAS – Verification and Testing of Multi-Agent Systems

Organizers: Alessio R. Lomuscio, Sophie Pinchinat, and Holger Schlingloff

Seminar No. 13181

Date: April 28 to May 3, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.5.151

© Creative Commons BY 3.0 Unported license

© Alessio R. Lomuscio, Sophie Pinchinat and Holger Schlingloff



Participants: Thomas Ágotnes, Carlos Areces, Guillaume Aucher, Alexandru Baltag, Ezio Bartocci, Ioana Boureanu, Nils Bulling, Louise A. Dennis, Michael Fisher, Tim French, Valentin Goranko, Stefan Gruner, Dimitar Guelev, Yuri Gurevich, Andreas Herzig, Wojtek Jamroga, François Laroussinie, Alessio R. Lomuscio, Nicolas Markey, Bastien Maubert, Stephan Merz, Aniello Murano, Wojciech Penczek, Sylvain Peyronnet, Jerzy Pilecki, Sophie Pinchinat, Franco Raimondi, Jean-François Raskin, Markus Roggenbach, Ina Schaefer, Holger Schlingloff, Gerardo Schneider, Henning Schnoor, François Schwarzentruher, Dmitry Shkatov, Ron van der Meyden, Hans Van Ditmarsch, Ramanathan Venkatesh, Karsten Wolf

Multi-agent systems (MAS) are distributed computing systems in which the individual components, or agents, interact with each other by means of communication, negotiation, cooperation etc., in order to meet private and common goals. The agent model finds applications in a variety of key applications of high-impact to society including web-services, autonomous vehicles, and e-government. But if MAS are to deliver on their promise to drive future applications, they need to be reliable.

MAS are typically specified and reasoned about by a variety of modal formalisms, including a variety of different logics. There are presently several, compartmented communities tackling questions pertaining to the correctness of MAS: researchers in model checking, model based testing, and controller synthesis. There presently is very little personal interaction among the scientists from different communities. The aim of this seminar was to bring these communities together, get exposure to each others' solutions to similar aims, and ultimately enhance their future interaction.

The topics concentrated on the intersection of the fields:

- Model checking of temporal-epistemic logic, alternating logics, and BDI logics
- Model based test generation for embedded systems
- Controller synthesis for self-organizing systems

In model checking, usually a model of the system and a property to be verified are given. In model based test generation, the goal is to construct a test suite from a model which establishes confidence in a certain property. In synthesis, a property and a model of computation are given, from which a strategy (a system model) is to be built. Both the test generation and the controller synthesis problem are closely related to model checking – in order to check the satisfiability of certain alternating time temporal logic (ATL) formulas in a model, one needs to construct a strategy for the participating agents.

The purpose of the seminar was to establish a common understanding of the problems in the different technologies of these application areas. It was expected that increased interaction between these three fields would stimulate new results and techniques of both theoretical relevance and practical usefulness.

Besides survey talks (60 minutes) on common technologies, attendees gave short contributions (30 minutes) and lightning presentations (15 minutes) on current research results and discussion rounds on open problems and research agendas. Additional technical sessions, including software demos, were organised spontaneously by the attendees for two of the evenings.

Attendees also contributed to the seminar by taking part in the lively discussions organised on topics of importance in the area. These were held in some of the afternoons but also during informal occasions outside the usual seminar hours such as after dinner. This helped bridge some of the gaps between the subdisciplines and rectify some misconception about each other's work.

Specifically, research topics of the seminar included:

- Logics and specification formalisms for MAS
- Verification and model checking for interactive systems
- Model-based testing for MAS
- Explicit, symbolic, and SAT-based algorithms for module checking
- Test case generation and synthesis
- Synthesis of winning strategies for games

The goals of the seminar were

- to obtain a common understanding of base technologies and intersections between these topics
- to collect a state-of-the-art picture of recent research results in the fields
- to confront methods from model checking and test generation for MAS

- to clarify terminology, research agendas and open problems
- to define a set of benchmark problems for verification and testing of MAS
- to bring together different communities formerly not interacting

The research topics were also discussed in relation with embedded systems applications such as:

- Verification of cyber-physical systems
- Validation of autonomous robots

It was felt that the seminar helped the participants to reach a common and shared understanding on the roles of logic, verification and testing as well as their interplay in the context of multi-agent systems

4.26 Meta-Modeling Model-Based Engineering Tools

Organizers: Tony Clark, Robert B. France, Martin Gogolla, and Bran V. Selic
Seminar No. 13182

Date: April 28 to May 3, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.4.188

© Creative Commons BY 3.0 Unported license

© Tony Clark, Robert B. France, Martin Gogolla, Bran V. Selic



Participants: Colin Atkinson, Tony Clark, Benoit Combemale, Larry Constantine, Lukas Diekmann, Catherine Dubois, Michalis Famelis, Robert B. France, Martin Glinz, Martin Gogolla, Lars Hamann, Øystein Haugen, Gabor Karsai, Steven Kelly, Thomas Kühne, Vinay Kulkarni, Stephen J. Mellor, Pieter J. Mosterman, Pierre-Alain Muller, Leonel Domingos Telo Nóbrega, Ileana Ober, Marian Petre, Dorina C. Petriu, Louis Rose, Bernhard Rumpe, Martina Seidl, Bran V. Selic, Perdita Stevens, Laurence Tratt, André van der Hoek, Markus Völter, Jon Whittle, Dustin Wüest

The 33 participants at the Meta-Modeling Model-Based Engineering Tools (M³BET) Dagstuhl Seminar were brought together to explore how model-based engineering (MBE) techniques can be used to improve the quality of software modeling tools. The participants included expert researchers, practitioners and tool developers in the software/system modeling and the human computer interaction communities. The discussions aimed to answer the following question: Can MBE techniques be used to produce more effective MBE tools, and, if so, how should it be done?

The vision underlying the seminar is one in which technologists create tool models that specify desired tool features, and tool modeling frameworks that are used to analyze, compose, transform, simulate and otherwise manipulate the tool models. In the vision, tool developers will use tool models and frameworks to produce useful, usable and cost-effective software modeling tools.

■ Seminar Organization

On the first day the seminar objective and outcomes were presented. The seminar objectives, as presented on that day, was to better understand the “what, why, and how” of tool models, and initiate work on (1) languages for tool modeling, (2) MBE methods and technologies for tool development, and (3) tool modeling frameworks.

The planned outcomes were (1) reports on the results of group discussions, (2) a research roadmap for achieving the tool modeling vision, (3) potential solutions for achieving the vision, and (4) initiation of new research collaborations among participants.

To help shape and initiate the discussions, the organizers proposed the following as an initial set of breakout group topics:

Tool capabilities. The intent was that discussions in this group would focus on identifying the software tool capabilities that should be captured in tool models, and on how these capabilities could be captured in tool metamodels. This covers discussions on (1) how metamodels can be used to describe tool capabilities in a manner that supports generation of high-quality tool components, (2) the utility and feasibility of defining tool metamodels, (3) potential benefits associated with and purposes served by a tool metamodel, and (4) challenges associated with developing an effective metamodel (i.e., a metamodel that is fit-for-purpose).

Tool qualities. Discussions in this group would aim to answer questions about desirable tool qualities (e.g., What issues determine tool adoption and why?). This includes key questions related to, for example, usability/human factors, scalability, interoperability, as well as non-technical but important considerations related to organization goals, culture, and processes.

Tool ecosystems. A tool framework can be thought of as a technological ecosystem that involves both tools as well as tool users. Discussions in this group would seek answers to questions such as: What are the features of a good tools framework? Are there candidate frameworks available? If so, are they sufficient or do they need to be extended?

Tool development methods. Discussions in this group would focus on answering the following questions: How can MBE be applied to the development of MBE tools? What types of languages are suitable for describing tools? How can tool quality issues be addressed by such methods?

■ Working Groups

During the discussions on group topics it was decided to base the groups on tool concerns and issues that the participants had some stake in. It turned out that the concerns and issues that arose from the discussions were mostly derived from those underlying the groups proposed by the organizers.

The concerns and issues were then clustered into two groups based on participant interests. Group A consisted of usability, utility, and broader non-technical concerns (e.g., designing tools that support how developers work and think, designing and performing usability studies, adapting tool features to user expertise and desired level of formality, marketing/business/cultural concerns). Group B consisted of the more technical concerns, for example, concerns related to tool development methods, scalability, support for separation of concerns, tool quality assessment and benchmarking.

Group B concerns were further grouped into two categories: Composition, and Methods and Quality concerns. The Composition concerns included issues related to tool, language, and model composition, and the use of multi-models with heterogeneous semantics.

Three Working Groups, each focusing on one of the above concern groups, were formed on the seminar's first day.

■ Summary and Future Work

One of the major insights gained during the seminar was that a good understanding of the utility of tool models and the identification of appropriate forms of tool models/metamodels requires one to first address more fundamental tool development and assessment concerns. On hindsight, this should not have been a surprising result; effective tool models would have to capture significant tool development and assessment experience and knowledge and thus such experience and knowledge needs to be distilled first. The seminar provided a good forum for discussing and organizing the experience and knowledge of the participants. Post-seminar collaborations that will utilize these results to develop an initial set of tool models/metamodels were initiated at the seminar.

In addition to the planned collaborations on tool models, the participants also agreed to engage in the following post-seminar activities:

- Publications: The following publications are planned
 - A special issue of the Software and System Modeling (SoSyM) journal that will include articles that focus on the concerns and issues discussed at the seminar.
 - A paper that discusses problems associated with current software modeling tools.
- Workshops: Workshops in which participants will discuss and develop tool models/meta-models will be held at conferences such as MODELS 2014 and ICSE 2014.

4.27 Tree Transducers and Formal Methods

Organizers: Sebastian Maneth and Helmut Seidl
Seminar No. 13192

Date: May 5–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.5.1

© Creative Commons BY 3.0 Unported license

© Sebastian Maneth



Participants: Henrik Björklund, Johanna Björklund, Adrien Boiret, Bruno Courcelle, Loris d'Antoni, Frank Drewes, Emmanuel Filiot, Zoltan Fülöp, Olivier Gauwin, Daniel Gildea, Kazuhiro Inaba, Florent Jacquemard, Jan Janousek, Naoki Kobayashi, Marco Kuhlmann, Pavel Labath, Aurélien Lemay, Sebastian Maneth, Wim Martens, Uwe Mönnich, Keisuke Nakano, Joachim Niehren, Damian Niwinski, Chih-Hao Luke Ong, Pierre-Alain Reynier, Kai T. Salomaa, Helmut Seidl, Frédéric Servais, Jean-Marc Talbot, Sophie Tison, Jan Van den Bussche, Margus Veanes, Heiko Vogler

The Dagstuhl seminar 13192 “Tree Transducers and Formal Methods” was a short two and a half day seminar that took place from May 5th to 8th, 2013. The aim was to bring together researchers from various research areas related to the theory and application of tree transducers. Tree transducers are a classical formalism in computer science, dating back to the early days of compilers and syntax-directed translation. Recently, interest in tree transducers has been revived due to surprising new applications in areas such as XML databases, security verification, programming languages, and linguistics. This seminar was meant to inspire the exchange of theoretical results and practical requirements related to tree transducers. These points were addressed in particular:

- Expressiveness versus Complexity: Which transducers offer the best trade-offs between expressiveness and complexity?
- Implementability under Resource Restrictions: Which transducer models can be executed by devices with bounded resources, e.g., for processing XML streams?
- New Applications: What new challenges do the different application areas of tree transducers raise? What new solutions have been found?
- Open Problems: Which are the most pressing open problems in tree transducer theory?

The seminar fully satisfied our expectations. The 33 participants from 13 countries (Australia, Belgium, Canada, Czech, France, Germany, Great Britain, Hungary, Japan, Poland, Slovakia, Sweden, and the US) had been invited by the organizer Sebastian Maneth to give particular survey talks about their recent research on applications and theory of tree transducers. There were talks focusing on very practical issues such as Margus Veanes’ talk on software verification using symbolic tree transducers (which kicked off the meeting), and also talks on highly challenging theoretical results such as the talk by Emmanuel Filiot on their recent breakthrough of proving that one-wayness of a two-way

word automaton is decidable. The other application areas, besides verification, were (1) tree processing (related to databases and search) (2) learning, and (3) linguistics.

The first talk by Veanes on symbolic transducers was followed by Jan Janousek about using pushdown automata to search for tree patterns, in linear order of trees. Symbolic transducers, from a theoretical point of view, were discussed in Heiko Vogler’s talk in the afternoon. Input driven pushdown automata, also known as nested word automata or visibly pushdown automata, were discussed with respect to descriptiveness complexity by Kai Salomaa. The second morning session of the first day was devoted to MSO translations, first about its theory with respect to word and tree translations by Bruno Courcelle, and then concerning a one-pass and linear time implementation model for MSO tree translations: the streaming tree transducer by Loris d’Antoni. The first afternoon section was about higher-order transducers, recursion schemes, and verification, given by Kazuhiro Inaba, Luke Ong, and Naoki Kobayashi. They discussed the open problem of proving context-sensitivity of the unsafe OI-hierarchy, results on model checking of higher-order recursion schemes, and practical approaches to type checking unsafe higher-order tree transducers.

The second day started with theoretical results about word and tree transducers by Emmanuel Filiot and Sebastian Maneth. The latter one was about deciding two database notions, namely determinacy and rewriting, for top-down and MSO tree transducers. Next was a sequence of talks about streaming, by Joachim Niehren, Pavel Labath, and Keisuke Nakano. They discussed practical aspects of early query answering, streaming of macro tree transducers using parallel streams, and stack attributed tree transducers, respectively. Related to streaming was the following talk by Frédéric Servais which surveyed recent results on visibly pushdown transducers. The following three talks discussed learning algorithms: first about tree series by Johanna Björklund

and Frank Drewes, and then about top–down tree transformations by Adrien Boiret. The last talk of the second day was Florent Jacquemard and Sophie Tison’s survey about tree automata with constraints.

The final day started with a talk about natural language processing using transducers, given by Daniel Gildea. It presented applications of multi bottom-up tree transducers to machine translation of natural language. It was followed by a talk by Uwe Mönnich on logical definitions of mildly context-sensitive grammar formalisms. A survey on “the tree-based approach” to natural language grammars was given by Marco Kuhlmann. Damian Niwinski’s talk connected to the session of the first day

on higher-order schemes: they are equivalent to panic automata, the invention and topic of Damian. An important practical consideration is incremental evaluation: it was discussed for XPath by Henrik Björklund and for succinct regular expressions by Wim Martens.

We thank Schloss Dagstuhl for the professional and inspiring atmosphere it provides. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers’ needs. For instance, elaborate research discussions in the evening were followed by musical intermezzi of playing piano trios by Beethoven and Mozart, or by table tennis matches and sauna sessions.

4.28 Information Visualization – Towards Multivariate Network Visualization

Organizers: Andreas Kerren, Helen C. Purchase, and Matthew O. Ward
Seminar No. 13201

Date: May 12–17, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.5.19

© Creative Commons BY 3.0 Unported license

© Andreas Kerren, Helen C. Purchase, and Matthew O. Ward



Participants: James Abello, Daniel Archambault, Katy Börner, Maura Conway, Stephan Diehl, Tim Dwyer, Peter Eades, Niklas Elmqvist, Jean-Daniel Fekete, Helen Gibson, Carsten Görg, Hans Hagen, Benjamin David Hennig, Danny Holten, Christophe Hurter, T. J. Jankun-Kelly, Daniel A. Keim, Jessie Kennedy, Andreas Kerren, Stephen G. Kobourov, Oliver Kohlbacher, Robert Kosara, Lothar Krempel, Kwan-Liu Ma, Guy Melançon, Silvia Miksch, Martin Nöllenburg, A. Johannes Pretorius, Helen C. Purchase, Jonathan C. Roberts, Falk Schreiber, John T. Stasko, Alexandru C. Telea, Jarke J. van Wijk, Tatiana von Landesberger, Matthew O. Ward, Michael Wybrow, Kai Xu, Jing Yang, Michelle X. Zhou, Björn Zimmer

■ Introduction

Information Visualization (InfoVis) is a research area that focuses on the use of visualization techniques to help people understand and analyze data. While related fields such as Scientific Visualization involve the presentation of data that has some physical or geometric correspondence, Information Visualization centers on abstract information without such correspondences, i.e., it is not possible to map this information into the physical world in most cases. Examples of such abstract data are symbolic, tabular, networked, hierarchical, or textual information sources.

The first two Dagstuhl Seminars on Information Visualization aimed to cover more general aspects of our field, such as interaction, evaluation, data wrangling, and collaboration, or focused on higher level topics, for instance, the value of InfoVis or the importance of aesthetics. Besides the Dagstuhl reports that are typically published directly after a seminar [1, 2, 4, 5], there were also follow-up publications for both seminars. The participants of Seminar #07221 wrote book chapters which have been consolidated into a Springer book [7]; the organizers of the same seminar published a workshop report in the Information Visualization journal [6]. For the second Seminar #10241, a special issue in the same journal was published [3].

The goal of this third Dagstuhl Seminar on Information Visualization was to bring together theoreticians and practitioners from Information Visualization, HCI, and Graph Drawing with a special *focus on multivariate network visualization*, i.e., on graphs where the nodes and/or edges have additional (multidimensional) attributes. The integration of multivariate data into complex networks and their visual analysis is one of the big challenges not only in visualization, but also in many application areas. Thus, in order to support discussions related to the visualization of real world data, we also invited researchers from selected application areas, especially bioinformatics, social sciences, and software engineering. The unique *Dagstuhl climate* ensured an

open and undisturbed atmosphere to discuss the state-of-the-art, new directions, and open challenges of multivariate network visualization.

■ Seminar Topics

The following themes were discussed during the seminar. The seminar allowed attendees to critically reflect on current research efforts, the state of field, and key research challenges today. Participants also were encouraged to demonstrate their system prototypes and tools relevant to the seminar topics. In consequence, topics emerged in the seminar week and were the focus of deeper discussions too.

- **Focus on biochemistry/bioinformatics:** In the life sciences, huge data sets are generated by high-throughput experimental techniques. Consequently, biologists use computational methods to support data analysis. The information in many experimental data sets can be either represented as networks or interpreted in the context of various networks. How can our current techniques help to analyze primary and secondary data in the context of such networks, and how can different network types be combined?
- **Focus on social science:** Graph drawing techniques have been used for several years for the visualization and analysis of social networks, but other social science fields (e.g., geography, politics, cartography, and economics) also make use of data visualization. How can (or do) our network visualizations support these domains?
- **Focus on software engineering:** In the application domain of software engineering, various graphs and data attached to graphs (e.g., software metrics) play a dominant role in the static and dynamic analysis of programs. Which of these problems are conceptually similar to graph-related problems in biology or social sciences and how can multivariate

network visualization support specific tasks, such as software architecture recovery?

- **Approaches and methods:** There already exist a number of technical approaches, algorithms, and methods to inter-actively visualize multivariate networks. Which ones are suitable for solving specific tasks in our applications areas? What is their potential? What are their limitations? By identifying the range of approaches that do exist, can we see the potential for new, innovative visualization ideas?
- **Challenges in visualizing multivariate networks:** Multivariate networks are large and complex and their complexity will increase in the future. Thus, not all problems can be solved in the short term. What are the current challenges?
- **Time-dependent/dynamic networks:** Many networks that are considered in practice change over time with respect to their topology and/or their attributes. How can we best visualize networks and attributes that change over time?
- **Interaction:** How can we best support the navigation, exploration and modification of multivariate networks?
- **Multiple networks at different scales:** How can we integrate, combine, compare more than one multivariate network at different scales? In this context, the term of so-called multi-modal networks is often used in literature. What does this term mean exactly? Can we visualize a range of different information types concurrently?
- **Tasks:** What range of tasks can multivariate network visualization support? Are there general tasks for all application domains?
- **Novel metaphors:** What type of visualization metaphors should we use beyond node-link diagrams? What would be the benefit in doing so?

■ Outcomes

The organizers and participants decided to write a book on multivariate network visualization to be published as LNCS issue

by Springer. The possibility of publishing this Springer book was confirmed by the Editor-in-Chief of LNCS already before the start of the seminar. Working groups have been invited to submit a book chapter building on their discussions and findings, and writing is underway. The final chapters are to be submitted by November 3, 2013, with a planned publication date of Spring 2014. A preliminary book structure was presented at the end of the seminar:

1. Introduction
 - a. Definition of multivariate networks, typical representations
2. Domain Application Data Characteristics in Context of Multivariate Networks
 - a. Biology
 - b. Social Sciences
 - c. Software Engineering
3. Tasks
4. Interaction
5. Metaphors (Visual Mappings beyond Node-Link)
6. Multiple and Multi-Domain Networks
7. Temporal Networks
8. Scalability
9. Summary/Conclusion

The Dagstuhl team performed an evaluation at the end of the seminar week. The results of this survey (scientific quality, inspiration to new ideas/projects/research/papers, insights from neighboring fields, ...) were throughout very good to excellent. Only a few single improvements were proposed by participants, for example, more junior researchers should be invited to come into contact with world-class researchers. And more domain experts should be invited to be spread out across the breakout groups. Another issue was that the time available for group work should be extended in future seminars.

■ References

- 1 Andreas Kerren, Catherine Plaisant, and John T. Stasko. 10241 Abstracts Collection: Information Visualization. In Andreas Kerren, Catherine Plaisant, and John T. Stasko, eds., *Information Visualization*, no. 10241 in Dagstuhl Seminar Proceedings, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010.
- 2 Andreas Kerren, Catherine Plaisant, and John T. Stasko. 10241 Executive Summary: Information Visualization. In Andreas Kerren, Catherine Plaisant, and John T. Stasko, eds., *Information Visualization*, no. 10241 in Dagstuhl Seminar Proceedings, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2010..
- 3 Andreas Kerren, Catherine Plaisant, and John T. Stasko. Information Visualization: State of the Field and New Research Directions. *Information Visualization*, 10(4):269–270, 2011.
- 4 Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North. 07221 Abstracts Collection: Information Visualization – Human-Centered Issues in Visual Representation, Interaction, and Evaluation. In Jean-Daniel Fekete, Andreas Kerren, Chris North, and John T. Stasko, eds., *Information Visualization – Human-Centered Issues in Visual Representation, Interaction, and Evaluation*, no. 07221 in Dagstuhl Seminar Proceedings, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2007.
- 5 Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North. 07221 Executive Summary: Information Visualization – Human-Centered Issues in Visual Representation, Interaction, and Evaluation. In Jean-Daniel Fekete, Andreas Kerren, Chris North, and John T. Stasko, eds., *Information Visualization – Human-Centered Issues in Visual Representation, Interaction, and Evaluation*, no. 07221 in Dagstuhl Seminar Proceedings, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany, 2007.
- 6 Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North. Workshop Report: Information Visualization – Human-Centered Issues in Visual Representation, Interaction, and Evaluation. *Information Visualization*, 6(3):189–196, 2007.
- 7 Andreas Kerren, John T. Stasko, Jean-Daniel Fekete, and Chris North, editors. *Information Visualization: Human-Centered Issues and Perspectives*, volume 4950 of LNCS. Springer, Berlin, Heidelberg, 2008.

4.29 Automated Reasoning on Conceptual Schemas

Organizers: Diego Calvanese, Sven Hartmann, and Ernest Teniente
Seminar No. 13211

Date: May 19–24, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.5.43
 © Creative Commons BY 3.0 Unported license
 © Diego Calvanese, Sven Hartmann, and Ernest Teniente



Participants: Alessandro Artale, Thomas Baar, Mira Balaban, Joachim Biskup, Xavier Blanc, Achim D. Brucker, Diego Calvanese, Marco A. Casanova, Carolina Dania, Sophie Dupuy-Chessa, David W. Embley, Ingo Feinerer, Enrico Franconi, Geri Georg, Parke Godfrey, Martin Gogolla, Sven Hartmann, Stephen J. Hegner, C. Maria Keet, Roman Kontchakov, Mirco Kuhlmann, Michael Leuschel, Jorge Lobo, Carsten Lutz, Stephan Mäs, Jerzy Marcinkowski, Marco Montali, Alessandro Mosca, Xavier Oriol, Elena V. Ravve, Guillem Rull, Klaus-Dieter Schewe, Ernest Teniente, Bernhard Thalheim, Dániel Varró, Qing Wang, Michael Zakharyashev

This Dagstuhl Seminar brought together 37 researchers from 16 countries across disciplines related to automated reasoning on conceptual schemas. The participants' expertise covered the three most popular languages used to specify the conceptual schema, i.e., Entity-Relationship (ER), Unified Modeling Language (UML) and Object-Role Modeling (ORM); either addressing reasoning only on the static (i.e., structural) schema alone or reasoning also on the elements of a conceptual schema that capture the dynamic (i.e., behavioral) aspects of a system.

Monday and Tuesday were devoted to short presentations from the participants of their most recent achievements in the field.

On Wednesday and Thursday morning the participants were allocated to three different groups, in parallel break out sessions, each one of them addressing a different aspect related to the topic of the workshop:

- On the practical applicability of current techniques for reasoning on the structural schema;
- Reasoning about the conceptual schema components capturing dynamic aspects;
- New challenges for automated reasoning on conceptual schemas.

The organizers asked each group to share the experiences of their participants and to try to identify the most pressing and challenging research issues or open problems for the aspect it addressed. Each group presented a summary of their results on Thursday afternoon. Thursday evening and Friday morning were devoted to a discussion about the outcomes of each group aiming at trying to come up with a roadmap for automated reasoning on conceptual schemas, something which was shown to be harder than expected.

4.30 Computational Methods Aiding Early-Stage Drug Design

Organizers: Andreas Bender, Hinrich Göhlmann, Sepp Hochreiter, and Ziv Shkedy
Seminar No. 13212

Date: May 19–14, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.5.78

© Creative Commons BY 3.0 Unported license

© Andreas Bender, Hinrich Göhlmann, Sepp Hochreiter, and Ziv Shkedy

Participants: Dhammika Amaratunga, Andreas Bender, Ulrich Bodenhofer, Chas Bountra, Javier Cabrera, Aakash Chavan Ravindranath, Hinrich Göhlmann, Jelle J. Goeman, Sepp Hochreiter, Wolfgang Huber, Murat Iskar, Adetayo Kasim, Samuel Kaski, Günter Klambauer, Leo Lahti, Justin Lamb, Johannes Mohr, Gianluca Pollastri, Ziv Shkedy, Willem Talloen, Oswaldo Trelles, Bie Verbist, Jörg Kurt Wegner



Besides discussing scientific findings enabled by computational approaches, the seminar successfully stimulated discussions between scientists from different disciplines and provided an exceptional opportunity to create mutual understanding of the various challenges and opportunities. It created understanding for technical terms and concepts and served as a catalyst to explore new ideas.

As a concrete example, it challenged the feasibility of utilizing chemical structure information for identifying correlations with biological data. Rather than attempting to define a most suitable way of translating chemical structure information into computer understandable form (e.g., via fingerprinting algorithms such as ECFP), the notion of utilizing functional readouts such as gene expression profiles was favored for prioritizing candidate drugs that demonstrate a favorable balance of desired and undesired compound effects.

4.31 Belief Change and Argumentation in Multi-Agent Scenarios

Organizers: Jürgen Dix, Sven Ove Hansson, Gabriele Kern-Isberner, and Guillermo Simari
Seminar No. 13231

Date: June 2–7, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.6.1

© Creative Commons BY 3.0 Unported license

© Jürgen Dix, Sven Ove Hansson, Gabriele Kern-Isberner, and Guillermo Simari



Participants: Edmond Awad, Pietro Baroni, Ringo Baumann, Pierre Bisquert, Alexander Bochman, Martin Caminada, Célia da Costa Pereira, Jürgen Dix, Florence Dupin de St-Cyr, André Fuhrmann, Dov M. Gabbay, Aditya K. Ghose, Massimiliano Giacomin, Sven Ove Hansson, Andreas Herzig, Anthony Hunter, Gabriele Kern-Isberner, Sebastien Konieczny, Patrick Krümpelmann, Daniel Lehmann, Beishui Liao, Pierre Marquis, Maria Vanina Martinez, Peter Novak, Nir Oren, Odile Papini, Matei Popovici, Mauricio Reis, Tjitze Rienstra, Ken Satoh, Jan Sefranek, Gerardo I. Simari, Guillermo R. Simari, Andrea Tettamanzi, Matthias Thimm, Serena Villata, Emil Weydert, Stefan Woltran, Zhiqiang Zhuang

Belief change and argumentation theory both belong to the wide field of knowledge representation, but their focal points are different. Argumentation theory provides frameworks for reasoning by setting up formal structures that allow the processing and evaluation of arguments for or against a certain option. Here, focus is put on dialectical deliberation and on finding justifications for decisions. Belief change theory has its focus on the adjustments of previously held beliefs that are needed in such processes. However, the interrelations between the two fields are still for the most part unexplored.

Both the fields of argumentation theory and belief revision are of substantial relevance for multi-agent systems which are facing heavy usage in industrial and other practical applications in diverse areas, due to their appropriateness for realizing distributed autonomous systems. Moreover, the topics of this seminar address recent research questions in the general area of decision making and are innovative in the combination of methods.

The seminar took place June 3rd–7th 2013, with 39 participants from 16 countries. The program included overview talks, individual presentations by the participants and group discussions. Overview talks ranged from 30 to 35 minutes, individual presentations were about 25 minutes long, including questions. We specifically asked participants not to present current research (their next conference paper), but rather asked to relate their research to argumentation/belief revision and how it could be used in agent theories.

Participants were encouraged to use their presentations to provide input for the discussion groups. We organized two discussion groups that each met twice (they took place in the afternoon, before and after the coffee break). Each group was headed by two organizers as discussion leaders.

The seminar concluded with the presentation of the group discussions on Friday morning and a wrap-up of the seminar.

From the discussion groups, some core topics arose which will help to focus further scientific work: Semantical issues concerning belief revision and argumentation were seen to be of major importance, and a layered view on both argumentation and belief revision, separating the underlying logic from the argumentation layer resp. revision layer helped to provide common grounds for the two communities. Both these topics proved to be very successful to stimulate scientific discourse, gave rise to interesting questions that might lead to papers and projects in the future, and look promising to allow a deeper analysis and a better understanding of the links between the two areas. Furthermore, a strong interest in having more applications and benchmarks became obvious, and a road map collecting informations on that is planned.

The organizers agreed to put together a special issue of *Annals of Mathematics and Artificial Intelligence* on *Argumentation and Belief revision* and invite papers on the use of methods and tools from belief change theory in argumentation theory, on the use of methods and tools from argumentation theory in belief change theory, on systems and frameworks that contain elements from both belief change and argumentation, and on practical applications of argumentation or belief revision in multi-agent systems or knowledge representation.

4.32 Indexes and Computation over Compressed Structured Data

Organizers: Sebastian Maneth and Gonzalo Navarro

Seminar No. 13232

Date: June 2–7, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.6.22

© Creative Commons BY 3.0 Unported license

© Sebastian Maneth and Gonzalo Navarro

Participants: Edmond Awad, Pietro Baroni, Ringo Baumann, Pierre Bisquert, Alexander Bochman, Martin Caminada, Célia da Costa Pereira, Jürgen Dix, Florence Dupin de St-Cyr, André Fuhrmann, Dov M. Gabbay, Aditya K. Ghose, Massimiliano Giacomini, Sven Ove Hansson, Andreas Herzig, Anthony Hunter, Gabriele Kern-Isberner, Sebastien Konieczny, Patrick Krümpelmann, Daniel Lehmann, Beishui Liao, Pierre Marquis, Maria Vanina Martinez, Peter Novak, Nir Oren, Odile Papini, Matei Popovici, Mauricio Reis, Tjitze Rienstra, Ken Satoh, Jan Sefranek, Gerardo I. Simari, Guillermo R. Simari, Andrea Tettamanzi, Matthias Thimm, Serena Villata, Emil Weydert, Stefan Woltran, Zhiqiang Zhuang



The Dagstuhl Seminar “Indexes and Computation over Compressed Structured Data” took place from June 2nd to 7th, 2013. The aim was to bring together researchers from various research directions of compression and indexing of structured data. Compression, and the ability to compute directly over compressed structures, is a topic that is gaining importance as digitally stored data volumes are increasing at unprecedented speeds. Of particular interest is the combination of compression schemes with indexes that give fast access to particular operations. The seminar was meant to inspire the exchange of theoretical results and practical requirements related to compression and indexing. These points were addressed in particular

- Tractability versus Intractability for Algorithmic Problems on Compressed Data
- Compression Algorithms for Strings, Trees, and Graphs
- Indexes for Compressed Data
- Algorithms for Compressed Data
- Better Search Results: Ranking and TF/IDF
- Applications of Structure Compression to other Areas

The seminar fully satisfied our expectations. The 34 participants from 11 countries (Canada, Chile, Denmark, Finland, Germany, Great Britain, Italy, Israel, Japan, Spain, and US) had been invited by the organizers to give survey talks about their recent research related to the topic of the seminar. The talks covered topics related to compression (e.g. grammar-based string compression) databases (e.g., XML, and top- k query answering), data structures (e.g. wavelet tries), string matching, and ranged to broad application areas such as biology. Most talks were followed by lively discussions. Smaller groups formed naturally which continued these discussions later.

We thank Schloss Dagstuhl for the professional and inspiring atmosphere. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers’ needs. For instance, elaborate research discussions in the evening were followed by local wine tasting or by heated sauna sessions.

4.33 Virtual Realities

Organizers: Guido Brunnett, Sabine Coquillart, Robert van Liere, and Gregory Welch
Seminar No. 13241

Date: June 9–14, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.6.37

© Creative Commons BY 3.0 Unported license

© Guido Brunnett, Sabine Coquillart, Robert van Liere and Gregory Welch



Participants: Carlos Andujar, Steffi Beckhaus, Roland Blach, Wolfgang Broll, Pere Brunet, Guido Brunnett, Sabine Coquillart, Carolina Cruz-Neira, Ralf Dörner, Steven K. Feiner, Uwe Freiherr von Lukas, Bernd Fröhlich, Henry Fuchs, Martin Göbel, Raphael Grasset, Jens Herder, Tobias Höllner, Charles E. Hughes, Masahiko Inami, Victoria Interrante, Bernhard Jung, Panagiotis D. Kakkis, Marcelo Kallmann, Yoshifumi Kitamura, Kiyoshi Kiyokawa, Gudrun Klinker, Ernst Kruijff, Torsten Kuhlen, Marc Erich Latoschik, Anatole Lecuyer, Robert W. Lindeman, Paul Milgram, Mark Mine, Betty Mohler, Tabitha C. Peck, Jerome Perret, John Quarles, Christian Sandor, Dieter Schmalstieg, Andreas Simon, Oliver Staadt, Anthony Steed, Jeanine Stefanucci, Frank Steinicke, Susumu Tachi, Robert van Liere, Gregory F. Welch, Gabriel Zachmann

Virtual Reality (VR) is a multidisciplinary area of research aimed at interactive human computer mediated simulations of artificial environments. An important aspect of VR-based systems is the stimulation of the human senses – usually sight, sound, and touch – such that a user feels a sense of presence in the virtual environment. Sometimes it is important to combine real and virtual objects in the same real or virtual environment. This approach is often referred to as Augmented Reality (AR), when virtual objects are integrated into a real environment. Research in VR and AR encompasses a wide range of fundamental topics, including: 3D interaction, presence, telepresence and tele-existence, VR modelling, multi-model systems, and human factors. Typical VR applications include simulation, training, scientific visualization, and entertainment, whereas typical AR applications include computer-aided manufacturing or maintenance, and computer-aided surgery or medicine.

The main goal of the seminar was to bring together leading international experts and promising young researchers to discuss current VR and AR challenges and future directions.

The organization built on the experiences from the previous seminar “Virtual Realities 2008”. The format of the seminar included sessions with standard presentations as well as parallel breakout sessions devoted to “hot-topics” in VR and AR research. It was the desire of the participants of the seminar that sufficient time for plenary discussion and working groups was scheduled. Before the seminar, the organizers solicited topics for the working groups. During the first days of the seminar these working groups were formed and a schedule was created. Plenary sessions were also scheduled to allow the working groups to report and discuss their findings.

Eight plenary sessions of presentations were scheduled throughout the week. Each session usually consisted of three 15 minute presentations followed by a 45 minute moderated discussion. Abstracts of the presentations are collected in the

next chapter. The Monday afternoon plenary sessions were devoted to the topics of Telepresence and Human Embodiment. Tuesday morning the topics Applications and Health/Wellbeing were presented. Wednesday morning was devoted to a session on Virtual Environments. The Thursday morning sessions were on Commercial/Business aspects of VR and Authoring/Content. The last session was devoted to Augmented Reality.

Seven working groups were created and parallel breakout sessions held throughout the week. Each working group reported their findings in plenary sessions. The following lists the titles of the working groups:

- Real Time Interactive Systems – Architecture Issues
- VR Current State and Challenges
- 3D User Interfaces
- Avatars in Virtual Reality
- Scientific Visualization and VR
- Characterising Interactions in Virtual (and/or Real) Environments
- Unconventional Mixed Environments

4.34 Parallel Data Analysis

Organizers: Artur Andrzejak, Joachim Giesen, Raghu Ramakrishnan, and Ion Stoica
Seminar No. 13251

Date: June 16–21, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.6.66

© Creative Commons BY 3.0 Unported license

© Artur Andrzejak, Joachim Giesen, Raghu Ramakrishnan, and Ion Stoica

Participants: Artur Andrzejak, Ron Bekkerman, Joos-Hendrik Böse, Sebastian Breß, Patrick Briest, Jürgen Broß, Lutz Büch, Michael J. Cafarella, Surajit Chaudhuri, Tyson Condie, Giuseppe Di Fatta, Rodrigo Fonseca, Johannes Fürnkranz, Joao Gama, Joachim Giesen, Philipp Große, Max Heimel, Yves J. Hilpisch, Anthony D. Joseph, George Karypis, Shonali Krishnaswamy, Soeren Laue, Frank McSherry, Jens K. Müller, Klaus Mueller, Srinivasan Parthasarathy, Tom Peterka, Raghu Ramakrishnan, Ion Stoica, Domenico Talia, Alexandre Termier, Markus Weimer, Hans-Martin Will, Matei Zaharia, Osmar Zaiane



■ Motivation and goals

Parallel data analysis accelerates the investigation of data sets of all sizes, and is indispensable when processing huge volumes of data. The current ubiquity of parallel hardware such as multi-core processors, modern GPUs, and computing clusters has created an excellent environment for this approach. However, exploiting these computing resources effectively requires significant efforts due to the lack of mature frameworks, software, and even algorithms designed for data analysis in such computing environments.

As a result, parallel data analysis is often being used only as the last resort, i.e., when the data size becomes too big for sequential data analysis, and it is hardly ever used for analyzing small and medium-sized data sets though it could be also beneficial for there, i.e., by cutting compute time down from hours to minutes or even making the data analysis process interactive. The barrier of adoption is even higher for specialists from other areas such as sciences, business, and commerce. These users often have to make do with slower, yet much easier to use sequential programming environments and tools, regardless of the data size.

The seminar participants have tried to address these challenges by focusing on the following goals:

- Providing user-friendly parallel programming paradigms and cross-platform frameworks or libraries for easy implementation and experimentation.
- Designing efficient and scalable parallel algorithms for machine learning and statistical analysis in connection with an analysis of use cases.

■ The program

The seminar program consisted of individual presentations on new results and ongoing work, a plenary session, as well as work in two working groups. The primary role of the focus

groups was to foster the collaboration of the participants, allowing cross-disciplinary knowledge sharing and insights. Work in one group is still ongoing and targets as a result a publication in a magazine.

The topics of the plenary session and the working groups were the following ones:

- Panel “From Big Data to Big Money”
- Working group “A”: Algorithms and applications
- Working group “P”: Programming paradigms, frameworks and software.

4.35 Interoperation in Complex Information Ecosystems

Organizers: Andreas Harth, Craig A. Knoblock, Kai-Uwe Sattler, and Rudi Studer
Seminar No. 13252

Date: June 16–19, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.6.82

© Creative Commons BY 3.0 Unported license

© Andreas Harth, Craig A. Knoblock, Kai-Uwe Sattler, and Rudi Studer



Participants: Felix Bießmann, Christian Bizer, Stefan Decker, Stefan Dessoach, John Domingue, Valeria Fionda, Yolanda Gil, Claudio Gutierrez, Armin Haller, Andreas Harth, Melanie Herschel, Aidan Hogan, Katja Hose, Martin Junghans, Kjetil Kjernsmo, Craig A. Knoblock, Ulf Leser, Andrea Maurino, Sheila McIlraith, Bernhard Mitschang, Felix Naumann, Daniela Nicklas, Giuseppe Pirrò, Axel Polleres, Kai-Uwe Sattler, Rene Schubotz, Steffen Staab, Thomas Steiner, Rudi Studer, Giovanni Tummarello, Raju Vatsavai, Ruben Verborgh

Individuals, enterprises and policy makers increasingly rely on data to stay informed and make decisions. The amount of available digital data grows at a tremendous pace. At the same time, the number of systems providing and processing data increases, leading to complex information ecosystems with large amounts of data, a multitude of stakeholders, and a plethora of data sources and systems. Thus, there is an increasing need for integration of information and interoperation between systems.

Due to the ubiquitous need for integration and interoperation, many research communities have tackled the problem. Recent developments have established a pay-as-you-go integration model, where integration is seen as a process starting out with enabling only basic query functionality over data and iteratively spending targeted integration effort as the need for more complex queries arises. Such an ad-hoc model is in contrast to previous integration models which required the construction of a mediated schema and the integration of schema and data before any queries – even simple ones – could be answered. The move towards less rigid integration systems can be traced back to many communities: the database community established Dataspaces as a new abstraction for information integration; the Semantic Web community provided ontologies and logic-based modelling in a web context; finally, the Web community established the Hypermedia principle which enables decentralized discovery and ad-hoc unidirectional interlinking in very large information systems.

Current systems for data integration focus on query-related aspects. However, to enable real interoperation, updates and invocation of functionality are required. Mobile applications, for example, require both access to information and functionality. We want to broaden the scope of research on data integration towards a vision of interoperation between systems (i.e., systems that not only exchange and integrate their data but also link functionality)

and investigate how an iterative model can be established for the interoperation of systems.

The seminar has multiple objectives:

- to bring together researchers from these diverse communities to identify common themes and to exploit synergies,
- to develop the theoretical foundations and an understanding of architectures and methods,
- to develop a research agenda and road-map towards a vision of web-scale integration and interoperation.



Fig. 4.5
Andrea Neumann – Laveline. Part of the Dagstuhl art collection and donated by: Kurt Geihs, Anca Muscholl, Ute Vollmar, Roland Vollmar, Igor Walukiewicz, Reinhard Wilhelm, and participants in Dagstuhl Seminars 08512, 07451, 08441, 09201, 09301, 09361 and 09501.

4.36 Theory of Evolutionary Algorithms

Organizers: Benjamin Doerr, Nikolaus Hansen, Jonathan L. Shapiro, and L. Darrell Whitley
Seminar No. 13271

Date: June 30 to July 5, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.7.1

© Creative Commons BY 3.0 Unported license

© Benjamin Doerr, Nikolaus Hansen, Jonathan L. Shapiro, and L. Darrell Whitley



Participants: Youhei Akimoto, Khulood Alyahya, Dirk V. Arnold, Anne Auger, Karl Bringmann, Dimo Brockhoff, Francisco Chicano, Kenneth A. De Jong, Benjamin Doerr, Carola Doerr, Anton V. Eremeev, Tobias Friedrich, Marcus Gallagher, Tobias Glasmachers, Nikolaus Hansen, Michael Hellwig, Christian Igel, Thomas Jansen, Daniel Johannsen, Joshua D. Knowles, Timo Kötzing, Marvin Künnemann, Jörg Lässig, William B. Langdon, Per Kristian Lehre, Johannes Lengler, Andrei Lissovoi, Luigi Malago, Rachael Morgan, Samadhi Nethmini Nallaperuma, Frank Neumann, Pietro S. Oliveto, Adam Prugel-Bennett, Jonathan E. Rowe, Günter Rudolph, Tom Schaul, Manuel Schmitt, Jonathan L. Shapiro, Peter F. Stadler, Sebastian U. Stich, Dirk Sudholt, Andrew M. Sutton, Olivier Teytaud, Lothar Thiele, L. Darrell Whitley, Christine Zarges

Evolutionary algorithms (EAs) are stochastic optimization methods that are based on principles derived from natural evolution. Mutation, recombination, and selection are iterated with the goal of driving a population of candidate solutions toward better and better regions of the search space.

In recent years, new methods have been developed at a rapid pace. Some of the advancements for continuous optimization methods have been enabled by focusing on how evolutionary algorithms can be compared and contrasted to more traditional gradient based methods. Arguably, evolutionary algorithms are one of the best methods now available for derivative-free optimization (DFO) on higher dimensional problems.

Another area of rapid recent advancement is in the area of run-time analysis for evolutionary algorithms applied to discrete optimization problems. Here, some techniques could be successfully borrowed from traditional algorithm analysis, but many new techniques were necessary to understand the more complicated stochastic processes arising from nature-inspired algorithms.

EA theory has gained much momentum over the last few years and has made numerous valuable contributions to the field of evolutionary computation. Much of this momentum is due to the Dagstuhl seminars on “Theory of Evolutionary Algorithms”, which has become the leading meeting for EA theorists in the world.

■ Specific Topics

This year, the following topics had the particular attention of organizers, speakers both of overview and specialized talks, and participants of the breakout sessions (also called “working parties” or “working groups” in other Dagstuhl seminars).

Advanced Runtime Analysis Methods. One difficulty common to the analysis of most randomized search heuristics

is that, while in principle these are nothing more than randomized algorithms, their particular nature disallows the use of many methods used in the classical analysis of the randomized algorithms community. The particular difficulties include dealing with populations (instead of a single search point as in other local optimizers) or recombination (instead of mutation only, which creates a search point close to the parent). Both the fitness level method and various variants of the drift analysis method were greatly improved in the last three years to cope with these difficulties. Also, the fixed budget view on runtime analysis was recognized as an alternative way of analyzing the performance of randomized search heuristics, and may better reflect performance indicators used by practitioners.

Complexity Theory for Randomized Search Heuristics. Complexity theory is one of the corner stones of classical computer science. Informally speaking, the *black-box complexity* of a problem is the number of fitness evaluations needed to find its solution. Unfortunately, it turns out that some notoriously hard problems like the clique problem in graphs have a ridiculously small black-box complexity. In their 2010 GECCO award winning paper, Lehre and Witt presented a promising way out of this dilemma. They introduced a restricted version of black-box complexity that on the one hand still covers most known evolutionary approaches, but on the other hand forbids the counter-intuitive tricks that led to the undesired results in the first approach. Following up on this work, several variants of black-box complexity have been suggested. During the seminar, in particular during the breakout session on this topic, these were intensively discussed, new variations have been proposed, both from the theory perspective and from practitioners, and a new approach was presented explaining how to gain new and better evolutionary algorithms from black-box complexity results.

Theory of Natural Evolutionary Algorithms.

Recently, the idea of conducting a natural gradient descent in the space of sampling probability distributions has been introduced in evolution strategies. The idea offers a very principled design technique for search algorithms that sample from a parameterized distribution. Comparable to classical deterministic optimization, an iterated gradient descent is performed on the distribution parameters. The remarkable difference is that the curvature information on this space is known a priori. A natural descent that is based on the inner product from the Fisher information matrix uses this curvature and is comparable to a Newton method. This new and promising idea is lesser-known and largely unexploited for evolutionary computation. This is a completely new topic for this seminar series, but it is related to previous work on Covariance Matrix Adaptation.

Theory for Multi-Objective Optimization. One of the most explosive areas of growth both within evolutionary algorithms and in derivative-free optimization is multi-objective optimization. This is because good evolutionary algorithms now exist that can cover complex Pareto fronts for 2 to 12 objectives. This gives practitioners a much more informative view of the trade-offs that are possible when facing a multi-objective decision, and can also reveal trade-offs that otherwise would never be seen: for example if we are wishing to minimize cost and maximize quality, there can be “knees” at specific locations on

the Pareto front where one might dramatically improve quality while incurring only a slight increase in cost. This is why multi-objective optimization methods that “map” the Pareto front are exciting. Yet, there is not a great deal of work on the theory of multi-objective optimization. Evolutionary algorithms are the method of choice for derivative-free multi-objective optimization and there is a great need to bring together theoreticians who are interested in evolutionary algorithms and those practitioners who are developing multi-objective optimization methods. This was another new topic for this seminar series.

Landscape Analysis. Landscape Analysis is an old idea: one should be able to compute features of a search space that can be used to guide search. One of the problems is that the kinds of metrics that one might wish to know about usually take exponential time to compute exactly. However, recent work has shown that some NP-hard problems (TSP, Graph Coloring, MAXSAT) can be decomposed to the point that Fourier methods can be used to exactly compute statistic moments of the search space (and subspaces of the search space) in polynomial time; these computation normally require exponential time for arbitrary problems. How can this information be used to guide the search, and to potentially replace heuristics with more exact information? New results in this area open new opportunities for exploration at this seminar series.

4.37 Computer Science in High Performance Sport – Applications and Implications for Professional Coaching

Organizers: Koen A.P.M. Lemmink, Stuart Morgan, Jaime Sampaio, and Dietmar Saupe
Seminar No. 13272

Date: June 30 to July 3, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.7.29

© Creative Commons BY 3.0 Unported license

© António Lopes



Participants: Youhei Akimoto, Khulood Alyahya, Dirk V. Arnold, Anne Auger, Karl Bringmann, Dimo Brockhoff, Francisco Chicano, Kenneth A. De Jong, Benjamin Doerr, Carola Doerr, Anton V. Ereemeev, Tobias Friedrich, Marcus Gallagher, Tobias Glasmachers, Nikolaus Hansen, Michael Hellwig, Christian Igel, Thomas Jansen, Daniel Johannsen, Joshua D. Knowles, Timo Kötzing, Marvin Künnemann, Jörg Lässig, William B. Langdon, Per Kristian Lehre, Johannes Lengler, Andrei Lissovoi, Luigi Malago, Rachael Morgan, Samadhi Nethmini Nallaperuma, Frank Neumann, Pietro S. Oliveto, Adam Prugel-Bennett, Jonathan E. Rowe, Günter Rudolph, Tom Schaul, Manuel Schmitt, Jonathan L. Shapiro, Peter F. Stadler, Sebastian U. Stich, Dirk Sudholt, Andrew M. Sutton, Olivier Teytaud, Lothar Thiele, L. Darrell Whitley, Christine Zarges

From June, 30th to July, 3rd, 2013 a seminar on “Computer Science in High Performance Sport – Applications and Implications for Professional Coaching” was held at Schloss Dagstuhl – Leibniz Center for Informatics. After 2006, 2008, and 2011 this seminar was the fourth on computer science in sport that was held in Dagstuhl.

Following the tradition, this seminar brought together experts from computer science together with experts from sports science to explore the options of interdisciplinary work.

This year emphasis was put on the interface between computer science and the high performance sport, in particular on coaching. The seminar focused on barriers that prevent coaches from embracing sport and computer science, and, how data can be presented in a more meaningful way so that coach’s expertise is enabled by science.

During the seminar, several participants presented their current research lines, ongoing work and open problems were discussed, focusing on three sub-themes: (1) coach-specific computer applications to address issues of communication and real-time application, (2) the pipeline from data acquisition to processing to analysis to visualization, and (3) modelling and simulation.

Twenty-seven invited participants, among which there were sports and computer scientists and coaches, gave a total of 25 talks and had enriching discussions about sport science. Problems, solutions, and benefits between computer science and sport science into high-performance coaching were discussed, and considered current developments in data acquisition, positional tracking, filtering, signal processing, game modelling, match analysis, performance analysis and optimization, computer-supported training, computer visualization and communication, 3D motion reconstruction, and serious games.

Once again, the Dagstuhl seminar concept provided benefits for the experts from different fields and countries that otherwise would hardly have met for an opportunity to exchange their ideas and inspire visions for the future of computer science and sport science in professional sport and coaching in an informal way. Several ideas were presented to try reduce gap between sport science and high performance coaching and new projects were discussed among the participants. Discussion led to current and future trends and challenges that require implementation on high performance sports coaching, such as: mobile computing, multimedia, data visualization, performance reconstruction and real time feedback.

4.38 Duality in Computer Science

Organizers: Mai Gehrke, Jean-Eric Pin, Victor Selivanov, and Dieter Spreen
Seminar No. 13311

Date: July 28 to August 2, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.7.54

© Creative Commons BY 3.0 Unported license

© Mai Gehrke, Jean-Eric Pin, Victor Selivanov, and Dieter Spreen

Participants: Jorge Almeida, Andrej Bauer, Veronica Becher, Ulrich Berger, Vasco Brattka, Karin Cvetko-Vah, Matthew de Brecht, Hannes Diener, Willem L. Fouché, Mai Gehrke, Serge Grigorieff, George Hansoul, Reinhold Heckmann, Peter Hertling, Achim Jung, Klaus Keimel, Andreas Krebs, Ganna Kudryavtseva, Hans-Peter Albert Künzi, Rutger Kuyper, Jimmie D. Lawson, Yoshihiro Maruyama, Paul-Andre Mellies, Luca Motto Ros, Alessandra Palmigiano, Arno Pauly, Jean-Eric Pin, Davide Rinaldi, Giuseppe Rosolini, Antonino Salibra, Matthias Schröder, Peter M. Schuster, Helmut Schwichtenberg, Victor Selivanov, Dieter Spreen, Paul Taylor, Sebastiaan A. Terwijn, Sam Van Gool, Lorijn van Rooijen, Klaus Weihrauch, Pascal Weil, Marc Zeitoun



This seminar concentrated on applications of duality in computation, semantics, and formal languages.

Duality and computation. Consider the area of exact real number computation. Real numbers are abstract infinite objects. Computing machines, on the other hand, can only transform finite objects. However, each real number is uniquely determined by the collection of rational open intervals that contain it, or a certain sub-collection thereof. Rational intervals can be finitely described as a pair of rationals. So, in order to compute with real numbers one has to compute with certain properties, i.e., one no longer works in the space of the reals but in the algebra generated by these properties. In doing so, the open intervals are considered as first-class objects and the concept of point is taken as a derived one. This is exactly the approach of pointless topology which tries to develop analytical concepts in a pointfree way, hereby using constructive logic.

Duality and semantics. In logic, dualities have been used for relating syntactic and semantic approaches. Stone's original result is in fact of this type as it shows that clopen subsets of Stone spaces provide complete semantics for classical propositional logic. This base case has been generalized in various directions. There is a general scheme underlying this work: given a logic, construct its Lindenbaum algebra which in these cases is a Boolean algebra with unary operators. Jonsson-Tarski duality relates such algebras to binary relational structures which in the modal case are just Kripke frames. In this setting, a wide spectrum of duality tools are available, e.g. for building finite models, for obtaining interpolation results, for deciding logical equivalence and other issues. For infinitary logics, Stone-type dualities have also played an important role starting with Scott's groundbreaking first model of the lambda-calculus

which is a Stone space. Subsequently Abramsky, Zhang and Vickers developed a propositional program logic, the logic of finite observations. More recently work of Jung, Moshier, and others has evolved this link between infinitary and finitary logics in the setting of logics for computation much further.

Duality and formal languages. The connection between profinite words and Stone spaces was already discovered by Almeida, but Pippenger was the first to formulate it in terms of Stone duality. Gehrke, Pin and Grigorieff lately systematized and extensively developed this discovery which led to new research efforts in formal language theory. A final goal is a general theory of recognition.

The seminar brought together researchers from mathematics, logic and theoretical computer science that share an interest in the fields of computing with infinite data, semantics and formal languages, and/or the application of duality results. The researchers came from 12, mostly European, countries, but also from Argentina, Japan, Russia, South Africa, and the United States.

Some of the specific questions that were investigated in talks and discussions:

- Explore the use of the link between finitary and infinitary Stone dualities in other settings than semantics;
- Explore the link between complexity theory and semantics provided by the connection via duality theory;
- Identify the relationship between game semantics and dual spaces;
- Explore the link between the profinite semi-groups used in formal language theory and logics given by state-based transition systems or categorical models)

4.39 “My Life, Shared” – Trust and Privacy in the Age of Ubiquitous Experience Sharing

Organizers: Alessandro Acquisti, Ioannis Krontiris, Marc Langheinrich, and Martina Angela Sasse

Seminar No. 13312

Date: July 28 to August 2, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.7.74

© Creative Commons BY 3.0 Unported license

© Alessandro Acquisti, Ioannis Krontiris, Marc Langheinrich, and Martina Angela Sasse



Participants: Alessandro Acquisti, Mads Schaarup Andersen, Zinaida Benenson, Bettina Berendt, Claudio Bettini, Rainer Böhme, Ian Brown, Claude Castelluccia, Delphine Christin, Alexander De Luca, Tassos Dimitriou, Frank Dürr, Deborah Estrin, Simone Fischer-Hübner, Michael Friedewald, Raghu K. Ganti, Jens Grossklags, Seda F. Gürses, Thomas Heimann, Ioannis Krontiris, Marc Langheinrich, Renè Mayrhofer, Joachim Meyer, Anthony Morton, David Phillips, Jo Pierson, Sören Preibusch, Kai Rannenber, Norman Sadeh, Martina Angela Sasse, Marcello Paolo Scipioni, Katie Shilton, Sarah Spiekermann

Advancements in smart phones and sensing technology have bolstered the creation and exchange of user generated content, resulting in new information flows and data-sharing applications. Through such applications, personal mobile devices are used to uncover and share previously private elements of people’s own everyday experiences. Examples include using smartphones or wearable sensors to collect and share context information (e.g., activities, social context, sports performance, dietary or health concerns). These flows of personal information have two distinct characteristics: they happen seamlessly (in real time, without necessarily the conscious participation of the user), and they are shared with a user’s family, social circles, or even publicly.

This new paradigm repositions individuals as producers, consumers, and remixers of a vast set of data with potential many economic and societal benefits. However, as sharing practices become more fluid than in desktop-based online environments, control over personal information flows becomes harder to maintain.

The goal of Dagstuhl Seminar 13312 “*My Life, Shared*” – *Trust and Privacy in the Age of Ubiquitous Experience Sharing* was to advance a research agenda in trust and privacy that addresses not only the evolution of the pervasive technologies underlying these trends (e.g., smartphones, wearable sensors), but also the surrounding societal and economic context, and to identify the resulting qualitative changes to the privacy landscape.

With that in mind, the seminar created an interdisciplinary discussion forum and a set of organised presentations around four broad areas: 1) tools and protocols, 2) usability and control tools, 3) behavioural decisions, and 4) social implications. Each area saw a selected set of participants present their work and views in the context of a short presentation, followed by an in-depth discussion session. From these discussions the organizers collected the main challenges and opportunities, and grouped them around four major themes: “Personal Data Services”, “Social Justice”, “Tool

Clinics”, and “Consequence-based Privacy Decision-making”. Each theme was subsequently discussed during one and a half days in four individual working groups, which presented their findings at the end of the seminar.

The full report not only contains the abstracts of the initial presentations but also the findings of the four thematic working groups. Below we summarize the main findings from these working groups.

Theme 1: Personal Data Service (PDS). A “Personal Data Service (PDS)” represents a trusted container for aggregating, storing, processing and exporting personal data. In principle, all data regarding the user (either user-generated or obtained from other sources, e.g. service providers) should be accessible to this container, including data about the user collected and published by others. Users are in control of all data stored in the PDS, which includes the option to share or sell parts of this data. In addition to storing data, the PDS can execute code to process this data locally.

By considering both a household- and a health-related scenario, the working group identified some of its properties and functionalities and sketched a possible system architecture that would include such a container. In a detailed discussion of benefits and risks, the working group concluded that there were still several issues to be investigated and real challenges that needed to be addressed before a PDS framework could be implemented and deployed, such as:

- Creating *incentives* to initial data providers to engage and open up the personal data APIs that are needed to fuel the PDS and associated applications.
- Creating *utility* from stored data: data fusion, sense making, and visualization that will lead to meaningful and actionable and sustainable engagement of the end user with their data.
- Addressing *privacy*: even though the PDS can increase transparency, awareness and engagement of users with their

data, it is neither obvious nor guaranteed that PDS will resolve user privacy problems and several of them remain open.

Theme 2: Social Justice. Privacy issues in participatory sensing are symptoms of broader concerns about the impact of sensing on social justice. Framing a social justice research agenda for participatory sensing requires the operationalization of concepts like fairness, human flourishing, structural change, and balances of power for system design, use, and regulation. The working group discussed how one might begin to operationalize these concepts for the design of data collection features, processing, sharing, and user interfaces. The group developed an analysis tool – a social justice impact assessment – to help system designers consider the social justice implications of their work during the design phase. The participants identified and presented several open questions that could spark future research, such as:

- If one assumes that participatory sensing will lead to *greater transparency*, will such transparency equally impact individuals, powerful people, and institutions?
- Do the powerful always end up *subverting transparency* schemes? Or can sensing change that tendency, for example by making facts visible to consumers and citizens, enabling organized responses (unionization)?
- What are the forums for *encouraging collective action* in participatory sensing? Can one encourage system designers to consider social justice during design by framing design as a collective action problem? Can participatory sensing open new avenues for consumers and citizens to organize collective action?

Theme 3: Tool Clinics, Privacy researchers and practitioners are working largely in isolation, concentrating on people's use of different user interfaces for privacy control, largely ignoring existing cross-disciplinary collaboration techniques. A "tool clinic" could encourage a collaborative (re)consideration of a technological solution, research technique or other artefact, in order to critically assess its design, development and deployment from multiple perspectives. A tool clinic can be used to provide a setting for those who are developing the solutions to rethink the framing and presentation of their solutions. The objective is to reflect from different perspectives on practices around the development, encoding, use, domestication, decoding and sustainability of a tool to gain quasi-ecological validation. The working group recommended to develop a tool clinic as a new event format for a scientific conference, ideally at a renowned computer-science

conference. This would combine the tool-centric nature of a demo session, the protected space of work-in-progress afforded by a workshop, and the mentoring spirit of a doctoral workshop. The format of a tool clinic session could typically consist of three steps:

1. Identifying particular affordances of the technological solution, research technique or other artefact and possible (unintended) consequences for people and society;
2. Gathering perspectives and practices of different experts, disciplines and/or stakeholders (e.g. users, policy makers, industry, etc.) linked with the development, deployment and sustainable evolution of a particular tool, solution, technique or artefact;
3. Informing and advising on technological design of the tool or solution, in order to avoid negative consequence and to further positive outcome.

Theme 4: Consequence-based Privacy Decisions.

Recent research shows that people not only want to control their privacy but are actually trying to do so. An appropriate privacy-respectful user interface should thus show users the consequences of making different privacy choices, rather than framing the choices only in technical terms regarding system parameters, which users often do not understand and do not care about. Providing tools to increase user comprehension of potential consequences is one of the next big challenges to be addressed in the field of privacy respectful user interfaces. In addition to helping users make better choices in terms of privacy protection, this will also allow them to make better informed decisions and hence, implement the notion of informed consent. The attempt to develop user interaction in this direction requires research on a number of issues that have so far received relatively little attention and concern, such as:

- *Expression of potential consequences:* The consequences should be expressed in a way that is comprehensible by different user categories from novices to experts.
- *Decision support:* Users could be further helped in their privacy decisions by external information sources. Studies to determine the responses to different kinds of information sources, different formats, and information from different groups of users will be necessary.
- *Minimal effort:* Introducing additional tools to help users make informed decisions may add significant overhead to the interaction. While this overhead may be the price to pay for better privacy protection, it should be limited to the minimum.

4.40 Reinforcement Learning

Organizers: Peter Auer, Marcus Hutter, and Laurent Orseau
Seminar No. 13321

Date: August 4–9, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.8.1
 © Creative Commons BY 3.0 Unported license
 © Peter Auer, Marcus Hutter, and Laurent Orseau



Participants: Peter Auer, Manuel Blum, Robert Busa-Fekete, Yann Chevaleyre, Marc Deisenroth, Thomas G. Dietterich, Christos Dimitrakakis, Lutz Frommberger, Jens Garstka, Mohammad Ghavamzadeh, Marcus Hutter, Rico Jonschkowski, Petar Kormushev, Tor Lattimore, Alessandro Lazaric, Timothy Mann, Jan Hendrik Metzen, Gergely Neu, Gerhard Neumann, Ann Nowé, Laurent Orseau, Ronald Ortner, Joëlle Pineau, Doina Precup, Mark B. Ring, Manuela Ruiz-Montiel, Scott Sanner, Nils T. Siebel, David Silver, Orhan Sönmez, Peter Sunehag, Richard S. Sutton, Csaba Szepesvári, William Uther, Martijn van Otterlo, Joel Veness, Jeremy L. Wyatt

Reinforcement Learning (RL) is becoming a very active field of machine learning, and this Dagstuhl Seminar aimed at helping researchers have a broad view of the current state of this field, exchange cross-topic ideas and present and discuss new trends in RL. It gathered 38 researchers together. Each day was more or less dedicated to one or a few topics, including in particular: The exploration/exploitation dilemma, function approximation and policy search, universal RL, partially observable Markov decision processes (POMDP), inverse RL and multi-objective RL. This year, by contrast to previous EWRL events, several small tutorials and overviews were presented. It appeared that researchers are nowadays interested in bringing RL to more general and more realistic settings, in particular by alleviating the Markovian assumption, for example so as to be applicable to robots and to a broader class of industrial applications. This trend is consistent with the observed growth of interest in policy search and universal RL. It may also explain why the traditional treatment of the exploration/exploitation dilemma received less attention than expected.



Fig. 4.6
Ina Rolshoven – Oelbild. Part of the Dagstuhl art collection and donated by: Roland Vollmar, Ute Vollmar, Reinhard Wilhelm, and participants in Dagstuhl Seminar 99041.

4.41 The Critical Internet Infrastructure

Organizers: Georg Carle, Jochen Schiller, Steve Uhlig, Walter Willinger, Thomas C. Schmidt, and Matthias Wählisch

Seminar No. 13322

Date: August 4–9, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.8.27

© Creative Commons BY 3.0 Unported license

© Georg Carle, Jochen Schiller, Steve Uhlig, Walter Willinger, Thomas C. Schmidt, and Matthias Wählisch



Participants: Bernhard Ager, Lothar Braun, Randy Bush, Georg Carle, Nikolaos Chatzis, Kenjiro Cho, David Choffnes, Alberto Dainotti, Roland Dobbins, Anja Feldmann, Timothy G. Griffin, Thomas Häberlen, Ethan Katz-Bassett, Stefan Katzenbeisser, Rossella Mattioli, Matthew Roughan, Jochen Schiller, Johann Schlamp, Thomas C. Schmidt, Yuval Shavitt, Georgios Smaragdakis, Rade Stanojevic, Steve Uhlig, Matthias Wählisch, Walter Willinger, Bill Woodcock

The Internet was designed to offer open data transfer services on a planetary scale. However, its success has turned it into a mission-critical infrastructure of vital importance for most countries, businesses, and industries. The aim of this seminar is to bring together the research and network operator communities to discuss and analyze the Internet as a critical infrastructure. We will address the vulnerability of the Internet from a number of different angles (e.g., physical infrastructure, control plane, data plane, services, etc.), with an emphasis on the core transport infrastructure as well as the content delivery side. The seminar will contribute to a better understanding of the Internet as a system of interdependent elements and pursue extensions of current research perspectives to consider novel (and maybe unusual) approaches to studying the local or region-specific substrates as parts of the Internet's global ecosystem.

Rethinking Perspectives on the Internet Backbone.

Analyzing the mutual impact between ASes, the vulnerability and efficiency of the backbone requires the identification of ASes and their role in mutual transit. In particular, stakeholders do not want to ground their internal data exchange on weak third parties. In Internet terms, AS interconnections between key players of a country should be part of a transparently visible Internet ecosystem. However, the Internet is a globally distributed network without boundaries, which makes the identification of locally relevant subparts hard. This seminar aims at being a platform to leverage new and uncommon research perspectives that go beyond the Internet backbone as a globally distributed system.

Methodologies to Analyze the Internet Structure.

To analyze the Internet as a critical infrastructure, a clear picture is required about the kind and granularity of data needed to obtain relevant results and draw valid conclusions, even if the available dataset is restricted. Sampling and inference are

common methods to assess the impact of the limited view on the real Internet. Current approaches to model the Internet backbone need to be revisited to reflect the Internet as critical infrastructure. The mapping of logical Internet nodes (ASes) to concrete entities (companies, points of presence etc.) as well as its annotation with meta data (e.g., administrative contact points) have been identified as important to cover the Internet structure from a non-technical perspective.

Paradigms Overlaying IP Connectivity. Delivering content to the end users is one of the main objectives of the Internet. In the early Internet, end users accessed content directly from a primary source. With the advent of CDNs this has changed. A single CDN operates as replication and distribution network for many content publishers, which brings data closer and more efficiently to end users. In fact, a very large portion of the current Internet content is maintained by only a limited number of CDNs, creating a limited competition in this area. Until now, this oligarchy has not been thoroughly studied, especially in the context of the Internet as a critical infrastructure.

Original Goals of the Seminar. The research questions to be pursued and answered include:

- How can we define and extract a locally relevant view of the globally distributed Internet?
- Which metrics are appropriate to measure the importance of Internet stakeholders and their mutual relationships?
- Which countermeasures and improvements are feasible to protect the Internet as critical infrastructure without narrowing its flexibility and openness?
- To what extent can we analyze the Internet structure in short time frames?
- What is the role of specific ASes for reliably interconnecting the Internet infrastructure of a country?
- How can we reveal weak transits and unintentionally strong

dependencies between ASes and specific regions of the world?

- How can we predict Internet scale consequences of large scale problems (what-if-questions)?

The complexity of the Internet makes it equally complex to give complete answers to these questions. This seminar helped us to start *touching* the questions. During our discussions it was clear that it is not only important to continue the work on these challenges but that it is also worth to follow up with a more specific focus on measurement aspects.

4.42 Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time

Organizers: Thore Husfeldt, Ramamohan Paturi, Gregory B. Sorkin, and Ryan Williams
Seminar No. 13331

Date: August 11–16, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.8.333

© Creative Commons BY 3.0 Unported license

© Thore Husfeldt, Ramamohan Paturi, Gregory Sorkin, and Ryan Williams



Participants: Per Austrin, Hans L. Bodlaender, Marek Cygan, Holger Dell, Andrew Drucker, Michael Elberfeld, Fedor V. Fomin, Serge Gaspers, Pinar Heggernes, Timon Hertli, Thore Husfeldt, Kazuo Iwama, Petteri Kaski, Eun Jung Kim, Joachim Kneis, Mikko Koivisto, Łukasz Kowalik, Dieter Kratsch, Stefan Kratsch, Alexander S. Kulikov, Daniel Lokshtanov, Shachar Lovett, Dániel Marx, Matthias Mnich, Jesper Nederlof, Rolf Niedermeier, Ramamohan Paturi, Peter Rossmanith, Rahul Santhanam, Saket Saurabh, Dominik Scheder, Stefan Schneider, Kazuhisa Seto, Gregory B. Sorkin, Ewald Speckenmeyer, David Steurer, Stefan Szeider, Navid Talebanfar, Suguru Tamaki, Virginia Vassilevska Williams, Magnus Wahlström, Ryan Williams

Computational complexity has demonstrated that thousands of important computational problems, spanning the sciences, are intimately linked: either they all have polynomial time algorithms, or none does. Nearly all researchers believe that $P \neq NP$, and that these problems do not all have low time complexity. However, they must be solved, one way or another, which means relaxing the requirements for “solving” a problem. One natural requirement to relax is the running time. Problems are often solved in practice by algorithms with worst-case exponential time complexity. It is of interest to find the *fastest* algorithm for a given problem, be it polynomial, exponential, or something in between.

This relaxation has revealed a finer-grained notion of problem complexity than NP-completeness. By definition all NP-complete problems are equivalent as far as the existence of polynomial time algorithms is concerned. However, the exact time complexities of these problems can be very different, just as their best approximation ratios can vary.

Algorithms for satisfiability represent well the progress in the field and the questions that arise. The theory of NP-completeness says that the Circuit Sat problem and 3-Sat are polynomial time equivalent. However, from the exact, exponential time perspective, the two problems look radically different.

For 3-Sat (and k -Sat in general), algorithms faster than the exhaustive search of 2^n assignments have been known for many years and are continually improved. The analysis of the randomized PPSZ algorithm for 3-Sat has recently been improved to $O(1.308^n)$, so currently the best known algorithm for this problem is also very simple. The best known deterministic algorithm runs in time $O(1.331^n)$, and is obtained by derandomizing earlier local search algorithms. A very natural DPLL-type algorithm for Formula Sat in fact has good performance on linear size formulas. All of these results represent major conceptual contributions.

No such progress has been made for general Circuit Sat. In fact, such results would have major implications in circuit complexity: even a $1.99^n \text{poly}(m)$ time algorithm for satisfiability of circuits with n inputs and m gates would imply exponential size *lower bounds* for solving problems with circuits. Between 3-Sat and Circuit Sat, there are also intermediate problems such as CNF-Sat that have resisted all efforts to produce an $O(1.99^n)$ time algorithm.

The basic algorithmic techniques to avoid exhaustive search are now consolidated in the field’s first textbook, (Fomin and Kratsch, *Exact Exponential Algorithms*, Springer 2010) though they are still being extended and refined. For example, there is now a general framework for making various exponential time dynamic programming algorithms, such as standard algorithms for Knapsack and Subset Sum, run in polynomial space. The fast zeta transform, which plays a central role in the implementation of inclusion-exclusion algorithms, continues to be actively researched. And “measure-and-conquer” methods for analyzing branching/backtracking algorithms continue to be enhanced.

However, many other powerful techniques have been explored only recently. One idea is to find combinatorial structures (such as matchings) by looking for corresponding algebraic objects (such as polynomials). The idea dates to Edmonds if not Tutte, but was introduced by Koutis for exponential time path and packing problems, leading to an $2^k \text{poly}(n)$ algorithm to find a k -path in a graph and a breakthrough $O(1.67^n)$ time algorithm for finding a Hamiltonian path, improving the 50-year-old previously best algorithm.

Other open problems in the field have been attacked by intricate, dedicated analyses; for example, there is now an algorithm for scheduling partially ordered jobs in $O((2 - \epsilon)^n)$ time.

Parameterized complexity is a closely related field that also investigates exponential time computation. Fundamentally, the field is interested in the dichotomy between algorithms that admit running times of the form $f(k)\text{poly}(n)$ (called fixed-parameter tractability) and those that do not, leading to qualitative hardness notions like $W[1]$ -hardness. This field continues to make great progress, with the parameterized tractability of many fundamental problems just being discovered. Just recently the first fixed-parameter algorithms for finding topological minors and the multi-cut problem have been found.

However, many recent results in that area are interested in determining (typically exponential) growth rate of the function f , instead of just establishing its existence. For example, a recent, very successful focus of parameterized complexity is the existence of problem *kernels* of polynomial size, or their absence under assumptions from classical computational complexity. In another direction, very strong lower bounds for algorithms parameterized by treewidth can now be shown under hypotheses about the exponential time complexity of Sat.

A *quantitative theory* of computational complexity of hard problems would address questions like why it is that 3-Sat can be solved in 1.4^n but CNF-Sat apparently cannot be solve. Ideally, we could hope for a characterization of the exact complexity of NP-complete problems, perhaps under some plausible assumptions. There is a growing body of work on the exact complexity of NP-complete problems which draws heavily from parameterized complexity theory. The *Exponential Time Hypothesis* (ETH), which posits that 3-Sat cannot be solved in $2^{o(n)}$ time, has given a strong explanatory framework for why some classes of problems admit improved algorithms while others are resistant. The results surrounding ETH show that if 3-Sat could be solved in subexponential time, then many other NP problems would also have subexponential algorithms.

Another compelling conjecture is the *Strong Exponential Time Hypothesis* (SETH) that CNF Satisfiability cannot be solved in 1.999^n time on formulas with n variables and cn clauses (for sufficiently large c). SETH has implications for k -Sat, other graph problems, and parameterized computation. There is less consensus about the truth of SETH; nevertheless, studying its implications will help better understand what makes CNF so difficult. A counting version of the hypothesis, #ETH, has recently been introduced to study the exponential time complexity of counting problems, such as the permanent and the Tutte polynomial.

Connections to other fields are being discovered as well, such as the importance of exponential time algorithms to the study of lower bounds in circuit complexity, as mentioned above.

For another example, a celebrated recent result in the complexity of approximation algorithms exhibits an $\exp(O(n^\epsilon))$ time approximation algorithm for Khot's Unique Games problem. This suggests that approximating unique games is a significantly easier task than solving NP-hard problems such as 3-Sat. The key to the algorithm is a new type of graph decomposition based on spectral methods. This decomposition method may well lead to more developments in exponential algorithms.

Furthermore, there are surprising connections between SETH and various other well-studied questions from other areas such as communication complexity and the 3-Sum hypothesis used in computational geometry and data structures. The instance compressibility notion introduced in the study of kernelisation turns out to be connected to the construction of hash functions.

These results show that increased attention to exponential time algorithms leads to progress of the highest caliber in well-established areas of the theory of computation.

4.43 Verifiably Secure Process-Aware Information Systems

Organizers: Rafael Accorsi, Jason Crampton, Michael Huth, and Stefanie Rinderle-Ma
Seminar No. 13341

Date: August 18–23, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.8.341

© Creative Commons BY 3.0 Unported license

© Rafael Accorsi, Jason Crampton, Michael Huth, and Stefanie Rinderle-Ma



Participants: Rafael Accorsi, Lujo Bauer, Anne Baumgraß, Nataliia Bielova, Achim D. Brucker, David Cohen, Jason Crampton, Christopher Dearlove, Guido Governatori, Christian Günther, Gregory Z. Gutin, Michael Huth, Fuyuki Ishikawa, Limin Jia, Mark Jones, Günter Karjoth, Felix Klaedtke, Agnes Koschmider, Jim Huan Pu Kuo, Andreas Lehmann, Niels Lohmann, Raimundas Matulevicius, Marco Montali, Charles Morisset, Alessandro Mosca, Michael Norrish, Andreas Oberweis, Alexander Paar, Stefanie Rinderle-Ma, Anne Rozinat, Thomas Stocker, Mark Strembeck, Meike Ullrich, Matthias Weidlich, Edgar Weippl, Nicola Zannone, Uwe Zdun, Maria Zhdanova

Business processes play a major role in many commercial software systems and are of considerable interest to the research communities in Software Engineering, and Information and System Security. A process-aware information system provides support for the specification, execution, monitoring and auditing of intra- as well as cross-organizational business processes.

Designing and enacting secure business processes is as tricky as “Programming Satan’s Computer”, as Ross Anderson and Roger Needham observed in a paper with that title. Recent fraud disasters show how subtle secure process engineering and control can be. The Swiss bank UBS suffered from a rogue trader scandal in 2011, which led to a loss of a then-estimated US\$2 billion, was possible because the risk of trades could be disguised by using “forward-settling” Exchange-traded Funds (ETF) cash positions. Specifically, processes that implemented ETF transactions in Europe do not issue confirmations until after settlement has taken place. The exploitation of this process allows a party in a transaction to receive payment for a trade before the transaction has been confirmed. While the cash proceeds in this scheme cannot be simply retrieved, the seller may still show the cash on their books and possibly use it in further transactions. Eventually, the mechanics of this attack allowed for a carousel of transactions, thereby creating an ever growing snowball. Similar analyses, usually based upon insider threats, can also be made for fraud cases such as the well-documented Société Générale case, but also for the WorldCom and Parmalat cases.

Addressing these problems requires, on the one hand, strong security and compliance guarantees. On the other hand, these guarantees must be substantiated by formal methods ensuring a verifiably secure business process enactment. It should be noted that these concerns are not confined to the financial service sector or to insider threats. For example, the planned unification of European data protection law into a sole Data Protection Regulation law is likely to change the statutory duties of the

private sector. Under this plan, companies will be legally required to report any breaches of this regulation and may be liable to penalties in the range of 2–5% of their global annual turnover. European industries seem to be ill-prepared to ensure that their information systems and processes will comply with the security requirements of that upcoming regulation, and the threat of substantial fines means that there is an urgent need to create more resilient systems and processes, which calls for more research within the thematic scope of this seminar.

At the interface of security requirements, business needs, and compliance methodologies we can ask many practically relevant research questions, and their answers are bound to have significant impact in academia and industry alike. Relatively little work has been done, however, on adapting or creating new formal methods with which one can check that processes are compliant with rules, preserve demanded privacy constraints, and enforce desired security policies at the same time.

One main purpose of the seminar was to present the state of the art in research within the three communities of Security, Verification, and Process-Aware Information Systems to all three communities in an accessible manner and with a view of identifying important research topics at the intersections of these communities. In addition, that exercise was also meant to explore what strategic activities could help in promoting research at the junction of these communities. This agenda was pursued through a mix of keynotes, technical presentations, break-out groups under the WorldCafe method, sessions with free-style discussions, and tool demonstrations.

We now highlight some of the key questions and findings that emerged during that week of work – we refer to the online archive of presentation slides, papers, and abstracts for more detailed discussions and findings. Three action items that seemed of particular importance to the participants were:

1. The need for a classification of security properties that are relevant for process-aware information systems, and an understanding of what formal methods might be able to analyse such properties.
2. The need for a set of concrete examples of business processes that are annotated with security considerations or constraints. These might be examples from the real world that have been sufficiently sanitized and anonymized.
3. The need for a review/survey article on the state of the art in formal methods, written for non-experts and ideally for an audience that deals with security, privacy, audit or business processes.

It was also asked what makes formal methods and tools “practical” in this problem space; there was concern about the scalability of these methods, but also about the considerable effort it would take to transfer foundational tools to real application domains – were such somewhat routine but important transfer may not be supported by standing funding models. Concerns were also voiced about the current research in security and privacy, which tends to ignore recent innovations in process composition, such as choreographies.

Another point of considerable interest made concerned the organization of research in this problem space. At the moment, researchers work on problems within their areas and when they begin to collaborate with people from another area this is then more of a bottom-up process where techniques and tools across areas are combined to see what problems one could now solve. It was remarked that it may often be more effective to take a top-down approach in which key problems of the inter-area

domain are first formulated and then researchers from the areas get together and try to come up with solutions that draw from their own tool boxes but that may also invent new tools for the problem at hand.

There was also a lot of discussion about what is so distinctive about *process-aware* information systems, and whether these differences to conventional information systems offer perhaps also opportunities. For example, it was discussed whether there is value in validating such systems at a high level of abstraction without considering how such processes get implemented in IT infrastructures and abstraction layers. The participants had mixed views on such merits but it was felt that validation at that level would be easier to realize and that the identification of weaknesses or vulnerabilities at that layer would no doubt be of value.

Another problem mentioned was the need to support legacy systems, and that this need would not go away. Faced with this, it appears that formal validation techniques will have to be able to reason about composed systems in which some parts only have a somewhat well defined interface, but whose internal behavior cannot be guaranteed or predicted to a good degree.

Finally, it was also noted that some of the research problems that suggest themselves to the specialists may not be issues in the field. For example, we may want trusted system composition across organizations but there may not be the need to formally validate such trust since contractual or other legal mechanisms may be in place that incentivize parties to honor that trust, and that give parties a means of seeking damages in case that trust has been violated. On the other hand, such legal mechanisms may not be adequate in the upcoming Internet of Things where 2-party, end-to-end composition will be the exception and not the norm.

4.44 ICT Strategies for Bridging Biology and Precision Medicine

Organizers: Andreas Dress, Titus Kühne, and Laxmi Parida
Seminar No. 13342

Date: August 18–23, 2013 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.3.8.342

© Creative Commons BY 3.0 Unported license

© Jonas S. Almeida, Andreas Dress, Titus Kühne, and Laxmi Parida



Participants: Jonas S. Almeida, Bernhard Balkenhol, Mark Braunstein, Robert Burk, Stefan Decker, Helena F. Deus, Andreas Dress, Jochen Dreß, David Gilbert, Anja Hennemuth, Scott Kahn, Ina Koch, Titus Kühne, Hans Lehrach, Pietro Lio', Markus Löffler, Wolfgang Maaß, Klaus Maisinger, Eric Neumann, Laxmi Parida, Alex Pothén, Eric Prud'hommeaux, Joel Saltz, Walter Schubert, Andrea Splendiani, Marc Van Regenmortel, Susana Vinga, Peter Walden, Zhenbing Zeng

Water, water, everywhere, nor any drop to drink. So goes Coleridge's *Rime of the Ancient Mariner*. Until recently, the same went for data: everywhere, but not of much use so far, neither for deriving new medical insights nor for improving medical care.

However, three key developments currently help to overcome this problem: the rapid adoption of electronic medical records [1], the dramatic advances in molecular biology [2], and, just as dramatic, the growing pervasiveness of social computing environments combined with a new attitude towards participatory health management [3–5]. The result is an exciting medley of initiatives devoted to supporting healthcare related information flow ranging from patient-facing resources such as *PatientsLikeMe* [6] to initiatives such as *MD-Paedegree* [7] (EU's FP7) that provides a physician-centric sort of '*PatientsLikeMine*' analogue addressing treatment choices in paediatrics.

Managing the *creative deconstruction* [8] involved in advancing towards systems medicine requires fundamentally changing the use of ICT in both, healthcare and biomedical research. It requires in particular to take account of the new paradigm of *web-centric computing* which is a basic prerequisite for all these initiatives.

Reflecting these concerns, a Dagstuhl Perspectives Workshop on *ICT Strategies for Bridging Biology and Medicine* was held to discuss a wide range of fundamental and foundational issues. These ranged from architectural considerations to data-access policies including *Open/Linked Data*, *the Semantic Web*, *Pervasive Hardware Ecosystems*, *Medical Clouds*, *Patient-Partic-*

ipation Frameworks, *'Healthcare 4.0'*, *Analytical Tools*, and *Medical Education*. Clearly, the required changes can only be achieved by initiatives of a broader scale and scope than what can be accommodated within the existing academic organisations. They need to always involve *all* stakeholders in the healthcare environment. In response to these challenges, the discussions led to the following theses and postulates:

- (i) An *open-data policy* for healthcare-related information systems is a fundamental and urgent imperative.
- (ii) Following the *usiness-IT alignment* paradigm [9], healthcare should – on all levels – be supported by secure IT-platforms enabling clinical workflow engines that map healthcare-related processes while integrating pertinent data-analysis, visualisation, and engineering tools.
- (iii) Such platforms should also take full advantage of advances provided by *cloud services*, *pervasive computing ecosystems*, and the *semantic web*.
- (iv) The *participatory potential* of the Web should be exploited to advance new forms of partnership in the healthcare environment.
- (v) The acquisition of *ICT literacy* must become a required part of biomedical education.
- (vi) Specifically in Germany, the Bundesnetzagentur should be encouraged to setting up a Working Group *Medizinische Netze* to explore options for a *Medical Cloud* within the German healthcare environment.

References

- 1 Tracy D. Gunter and Nicolas P. Terry *The Emergence of National Electronic Health Record Architectures in the United States and Australia: Models, Costs, and Questions*. J Med Internet Res 7:1 (2005).
- 2 Susan Desmond-Hellmann et al. *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease*. National Research Council (US), Committee on A Framework for

- Developing a New Taxonomy of Disease. The National Academies Press, Washington D.C., USA (2011).
- 3 Wikipedia. http://en.wikipedia.org/wiki/Social_computing
 - 4 Barbara A. Israel et al. *Community-based participatory research: policy recommendations for promoting a partnership approach in health research*. *Education for Health*, 14:2 (2001):182–197.
 - 5 Melanie Swan. *Emerging Patient-Driven Health Care Models: An Examination of Health Social Networks, Consumer Personalised Medicine and Quantified Self-Tracking*. *Int. J. Environ. Res. Public Health* 6 (2009):492–525.
 - 6 Ben Heywood et al. <http://www.patientslikeme.com/>.
 - 7 Bruno Dallapiccola et al. <http://www.md-paedegree.eu/>. The European Commission.
 - 8 Eric J. Topol. *The Creative Destruction of Medicine: How the Digital Revolution Will Create Better Health Care*. Basic Books, New York, NY, USA (2012).
 - 9 Wim van Grembergen and Steven De Haes. *Enterprise Governance of IT: Achieving Strategic Alignment and Value*. Springer, New York Heidelberg Dordrecht London (2009).

4.45 Coding Theory

Organizers: Hans-Andrea Loeliger, Emina Soljanin, and Judy Walker
Seminar No. 13351

Date: August 25–30, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.8.136

© Creative Commons BY 3.0 Unported license

© Hans-Andrea Loeliger, Emina Soljanin, and Judy Walker



Participants: Daniel Augot, Angela Barbero, Alexander Barg, Eimear Byrne, Pascale Charpin, Gerard Cohen, Stark C. Draper, Iwan M. Duursma, Salim El Rouayheb, Marcelo Firer, Heide Gluesing-Luerssen, Elisa Gorla, Marcus Greferath, Hamed S. Hassani, Michael Heindlmaier, Tor Helleseth, Werner Henkel, Tracey Ho, Tom Høholdt, Jørn Justesen, Axel Kohnert, Margreta Kuijper, P. Vijay Kumar, Michael Lentmaier, Hans-Andrea Loeliger, Felice Manganiello, Muriel Medard, Sihem Mesnager, Olgica Milenkovic, Katherine Morrison, Joachim Rosenthal, Vladimir Sidorenko, Vitaly Skachek, Roxana Smarandache, Patrick Solé, Emina Soljanin, Alex Sprintson, Vladimir D. Tonchev, Anna-Lena Trautmann, Bane Vasic, Pascal Vontobel, Judy L. Walker, Wolfgang Willems, Oyvind Ytrehus, Jiun-Hung Yu

While coding theory has evolved into an essential ingredient of contemporary information technology, it remains a fascinating area of research where many fundamental ideas of information theory and mathematics meet. Indeed, the diversity and profundity of recent new ideas in, and new applications of, coding theory is impressive. The following themes were of primary interest at the seminar:

Codes on graphs include turbo codes, low-density parity check codes, and a variety of similar codes. Due to the recent new idea of “spatial coupling”, such codes can now be designed to achieve the Shannon capacity of most communication channels with practical encoders and decoders. Such codes are a perfect nurturing ground for cross-fertilization of ideas between computer science, electrical engineering, and mathematics. The mathematical tools in this area include ideas from graph theory, probability, algebra, discrete mathematics, and statistical physics.

Algebraic coding theory continues to be of supreme theoretical and practical interest. Prime examples of this area are Reed-Solomon codes, codes from algebraic geometry, and codes obtained from algebraically constructed graphs. Recent advances in the field include, in particular, list-decoding algorithms for various classes of algebraic codes. Emerging relationships between this area and codes on graphs appear to be promising for future research.

Polar codes (discovered by Arikan in 2008) are a breakthrough of utmost significance. Such codes are provably capacity-achieving on very many channels with very low-complexity (and very practical) encoders and decoders. These codes rely on a new large-system limit that combines information theory and coding theory more smoothly than any prior coding technique. The investigation of such codes, including their combination with other coding techniques (such as codes on graphs and algebraic codes), is an exciting new area of research.

Network coding aims at improving data transmission (throughput, reliability, latency, etc.) in networks. This area is still quite young, but it has begun to influence the design of methods and protocols of content delivery in the internet. There is a diverse set of network coding problem formulations, and network coding can be (and has been) studied within a number of different theoretical frameworks, such as algebraic, combinatorial, information theoretic, and linear programming frameworks.

Codes for cloud applications are about distributed storage of large amounts of data. Diverse requirements on reliability, access latency, updatability, and repairability pose entirely new challenges for coding theory.

In addition, there were also two talks on topics in coding theory inspired by biology.

The seminar brought together 45 high-caliber researchers with backgrounds and interests in these different areas. The seminar was held in the usual Dagstuhl style, with a rather light program of formal presentations and much room for informal interaction. It was interesting and stimulating to hear of developments outside one’s own speciality, and (to the best of our knowledge) all attendants greatly enjoyed the seminar.

4.46 Interaction with Information for Visual Reasoning

Organizers: David S. Ebert, Brian D. Fisher, Petra Isenberg, and Shixia Liu
Seminar No. 13352

Date: August 25–30, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.8.352
 © Creative Commons BY 3.0 Unported license
 © David S. Ebert, Brian D. Fisher, and Petra Isenberg

Participants: Simon Atfield, Anastasia Bezerianos, Sheelagh Carpendale, Peter C.-H. Cheng, Fanny Chevalier, Christopher Collins, Mary Czerwinski, David S. Ebert, Thomas Ertl, Brian D. Fisher, Steve Franconeri, Kelly Gaither, Wayne D. Gray, Hans Hagen, Petra Isenberg, Tobias Isenberg, Daniel Keefe, David Kirsh, Jörn Kohlhammer, Heidi Lam, Bongshin Lee, Chris North, Catherine Plaisant, Margit Pohl, Huamin Qu, Kamran Sedig, Jinwook Seo, Christian Tominski, Xiaoru Yuan, Michelle X. Zhou



Scientific and information visualization researchers routinely build and evaluate interactive visualization systems to aid human reasoning. However, this work is often disconnected from the methodological and theoretical tools developed by the cognitive and social sciences to address the complexities of human thought processes. Those tools and methods can help us to understand human perception and understanding of data visualization, but typically do not address how rich interaction with computational processes could be engineered to support better decision-making. Yet, an increasing number of researchers are turning to the question of how to best engineer interaction techniques for visualization and how to best study and understand their influence on cognition, insight formation, and also efficiency and effectiveness of work. The goal of this seminar was to bring together researchers in cognitive science and psychology with researchers in the field of visualization to discuss the value that interaction can bring to visualization, how best to study it, and how research on interaction in cognitive science can be best integrated into visualization tools and systems to the benefit of domain experts or casual users of these tools.

4.47 Crowdsourcing: From Theory to Practice and Long-Term Perspectives

Organizers: Claudio Bartolini, Tobias Hoßfeld, Phuoc Tran-Gia, and Maja Vukovic
Seminar No. 13361

Date: September 1–4, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.4.9.1

© Creative Commons BY 3.0 Unported license

© Tobias Hoßfeld, Phuoc Tran-Gia, and Maja Vukovic



Participants: Abraham Bernstein, Kathrin Borchert, Alessandro Bozzon, Cristina Cabanillas, Joseph Davis, Gianluca Demartini, Klaus Diepold, Matthias Hirth, Tobias Hoßfeld, Andreas Hotho, Deniz Iren, Christian Keimel, Shinichi Konomi, Vassilis Kostakos, Markus Krause, Martha A. Larson, Babak Naderi, Nhatvi Nguyen, Munindar P. Singh, Phuoc Tran-Gia, Maja Vukovic, Florian Zeiger

Over the past several years crowdsourcing has emerged as a new research theme, but also as a new service platform and Internet for harnessing the skills of the large, network-connected crowd on-line. Whilst the research community has not just yet recognized crowdsourcing as an entirely new discipline, many research challenges remain open and need to be addressed to ensure its successful applications in academia, industry and public sectors. Crowdsourcing research intersects many existing domains and brings to the surface new challenges, such as crowdsourcing as a novel methodology for user-centered research; development of new services and applications based on human sensing, computation and problem solving; engineering of improved crowdsourcing platforms including quality control mechanisms; incentive design and gamification of work; usage of crowdsourcing for professional business; theoretical frameworks for evaluation. Crowdsourcing, as a new means of engaging human capital online is increasingly having an impact on the Internet and its technical infrastructure, on society, and the future of work.

With crowdsourcing gaining momentum and becoming mainstream, the objective of this Dagstuhl seminar was to lead coordination of research efforts in the different communities, especially in US currently leading the crowdsourcing market and in Europe. The seminar engaged experts from the different research fields (e.g. sociology to image processing) as well as experts from industry with a practical background on the deployment, operation or usage of crowdsourcing platforms. From industry, real-world problem statements, requirements and challenges, position statements, innovative use cases, and practical experiences are tackled and discussed. The collection and analysis of practical experiences of the different crowdsourcing stakeholders were key outcomes of the Dagstuhl Seminar. The seminar was structured so that the participants use existing use cases, as a driver in the discussion to envision future perspectives

of this domain. To move forward, we identified the need for a common terminology, classification and taxonomy of crowdsourcing systems, as well as evaluation frameworks; and have already proposed a blueprint of the same. The impact of crowdsourcing from different perspectives has been discussed, by participants' viewpoints stemming from societal, business, economic, legal and infrastructure perspectives.

From platform provider side, Nhatvi Nguyen showed the actual challenges in operating a crowdsourcing platform. As industry use case, the example of enterprise crowdsourcing was presented by Maja Vukovic, where the rapid generation of a snapshot of the state of IT systems and operation is conducted by means of crowdsourcing. This allows for massive cost savings within the company by uncovering knowledge critical to IT services delivery. Crowdsensing is another industry use case presented in the seminar by Florian Zeiger. Environmental sensing in the area of safety and security was discussed from industry point of view along with the challenges and open questions, e.g. user privacy, data quality and integrity, efficient and reliable data collection, as well as architectural decisions and flexible support of various business models. A concrete application for crowdsensing is radiation sensing as shown by Shinichi Konomi.

Beyond this, there were also discussions on multimedia related use cases. Crowdsourcing can be efficiently used for describing and interpreting multimedia on the Internet and allows to better address other aspects of multimedia with meaning for human beings. Martha Larson provided examples of these aspects like the emotional impact of multimedia content, and judgments concerning which multimedia is best suited for a given purpose. Klaus Diepold applied crowdsourcing to move subjective video quality tests from the lab into the crowd. The resulting ratings are used to train mathematical model for predicting subjective quality of video sequences. Multivariate data analysis tools are

recommended to incorporate contextual information to further validate the mathematical model. Vassilis Kostakos showed that the data quality of appropriate subjective tests may be increased by using public displays and touch screens in cities compared to online surveys. While gamification pops up as buzzword aiming among others at increased data quality, Markus Krause mentioned that the player should be put first i.e. the desires of player are paramount. In particular, task and game ideas need to be able to be linked, while fun has to be the main motivator for the game.

General approaches to improve crowdsourcing and the resulting data quality were a topic of interest by several participants. Gianluca Demartini proposes to model workers in the crowd as basis for quality assurance mechanisms. Alessandro Bozzon demanded for better conceptual abstractions for crowd tasks and processes design and (automatic) generation; better understanding of crowds properties such as (soft and hard) skills, reliability, availability, capacity, precision; and better tools for measuring and driving worker engagement. Cristina Cabanillas considered the human resource management aspects starting from workflows to crowdsourcing. Abraham Bernstein discussed human computers as part of computational processes, however, with their own strengths and issues. The three traits on human computation, that are motivational diversity, cognitive diversity, and error diversity, are embraced as strengths instead of weaknesses. While the main focus of the seminar was on technical challenges, the potential impact and long-term perspectives were discussed from an interdisciplinary point of view too, given the social and human aspects of crowdsourcing. Those issues were also raised by Phuoc Tran-Gia and Joseph G. Davis.

Overall there were 22 participants from 9 countries and 16 institutions. The seminar was held over 2.5 days, and included presentations by researcher and specific hands-on discussion sessions to identify challenges, evaluate viewpoints and develop a research agenda for crowdsourcing. The different aspects of crowdsourcing were discussed in more detail in four different working groups formed during the seminar:

- (W1) long-term perspectives & impact on economics in five years,
- (W2) theory: taxonomy and dimensions of crowdsourcing,
- (W3) industry use cases,
- (W4) crowdsourcing mechanisms and design.

Please note that a related seminar on “Cloud-based Software Crowdsourcing” (Dagstuhl Seminar 13362), organized by Michael N. Huhns, Wei Li, Martin Schader and Wei-Tek Tsai, took place in parallel to this seminar. We held a joint social event and a session on discussing research challenges and planned publications. In this late night session, on one hand ethical issues in the area of crowdsourcing were raised in a stimulus talk by Martha Larson (TU Delft). On the other hand, Munindar P. Singh (North Carolina State University) intended to provoke with his talk on the critique of current research in the area of social computing and crowdsourcing.

A comprehensive list of open problems and challenges in the area of crowdsourcing as observed and stated by the participants is another key outcome of the seminar.

4.48 Cloud-based Software Crowdsourcing

Organizers: Michael N. Huhns, Wei Li, and Wei-Tek Tsai
Seminar No. 13362

Date: September 1–4, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.34

© Creative Commons BY 3.0 Unported license

© Michael N. Huhns, Wei-Tek Tsai, and Wenjun Wu



Participants: Shaukat Ali, Xiaoying Bai, Xavier Blanc, Kyle Chard, Schahram Dustdar, Michael N. Huhns, Robert Kern, Donghui Lin, Greg Little, Xinjun Mao, Michael Maximilien, Dave Murray-Rust, Khrystyna Nordheimer, Dirk Riehle, Ognjen Scekcic, Lionel Seinturier, Hong-Linh Truong, Wei-Tek Tsai, Huaimin Wang, Wenjun Wu, Gang Yin, Tao Yue

Crowdsourcing software development or software crowdsourcing is an emerging software engineering approach. Software development has been outsourced for a long time, but the use of Internet with a cloud to outsource software development to the crowd is new. Most if not all software development tasks can be crowdsourced including requirements, design, coding, testing, evolution, and documentation. Software crowdsourcing practices blur the distinction between end users and developers, and allow the co-creation principle, i.e., a regular end-user becomes a co-designer, co-developer, and co-maintainer. This is a paradigm shift from conventional industrial software development to a crowdsourcing-based peer-production software development. This seminar focused on the notion of cloud-based software crowdsourcing, with the following goals:

1. to establish a theoretical framework for applying software crowdsourcing, and identify the important design patterns and highly interactive and iterative processes in a cloud-based infrastructure.
2. to propose and design a reference architecture for software crowdsourcing
3. to develop and finalize the research roadmap for software crowdsourcing for the next five years

The grand research challenge in cloud-based software crowdsourcing is how to embrace elements from the two aspects: cloud infrastructure and software crowdsourcing. Metaphorically, it can be regarded as synergy between two clouds – machine cloud and human cloud, towards the ultimate goal of developing high-quality and low cost software products. This seminar intended to bring together scientists from both fields to tackle the major research problems in this emerging research area.

More than twenty researchers, who work on different domains such as crowdsourcing, human-computer interaction, cloud computing, service oriented computing, software engineering and

business management attended the seminar. In addition to regular 5-minute talks from every participant in the seminar, the organizer arranged a keynote speech delivered by Prof Schahram Dustdar, which summarizes large-scale collective problems solving research enabling software crowdsourcing. The topics covered by their presentations can be roughly categorized into three groups: software crowdsourcing process and models, crowdsourcing cloud infrastructure and human crowd management. To promote in-depth discussion among these topics, we also divided people into five discussion groups including:

Crowd Source Software Engineering Design-Group:

This group identified the three main areas in the design of software crowdsourcing: processes, models, and techniques. It highlighted the importance of standardized generic models of software crowdsourcing study, and explored multiple crowdsourcing techniques, especially virtual team formation and quality assessment.

Worker-centric design for software crowdsourcing:

This group focused on the crowd management in software crowdsourcing and aimed to answering the question about how to make a sustainable software crowdsourcing industry. Discussion in the group covered the major issues such as careers and reputation development of workers, trust among workers and “employers” (task solicitors) on crowdsourcing markets, virtual team selection and team building.

Cloud-based Software Crowdsourcing Architecture:

This group discussed the possible common architectures of crowd-sourcing applications and explored two complementary architectural approaches.

Experimentation Design for Software Crowdsourcing:

The central topic of this group is about how to design a valid and reproducible experiment for software crowdsourcing research. The group had extensive discussion on software

crowdsourcing experiment approaches and the major crowdsourcing infrastructures.

Infrastructure and Platform:

This group reviewed the motivations to construct the crowdsourcing platform, analyzed architecture design issues, and proposed a educational platform for software crowdsourcing.

During the session of our seminar, Dagstuhl also set up a parallel seminar named “Crowdsourcing: From Theory to Practice and Long-Term Perspectives”, which mostly focused on general crowdsourcing research and service platforms. Software crowdsourcing can be regarded as one of the most complex crowdsourcing activities that often need intense dedication from workers with high-level skills of software engineering. Thus,

there are some interesting overlapping areas such as worker incentive and quality assurance, between our seminar and the parallel seminar. To foster collaboration among the two groups, we hold a joint discussion session for introducing and sharing findings from each group, followed by an evening session with two presentations from the general crowdsourcing group.

We believe this seminar is a good start for software crowdsourcing research. Finding and consensus generated from the seminar have been formalized in the wiki page of software crowdsourcing (http://en.wikipedia.org/wiki/Crowdsourcing_software_development) to give a clear definition and initial reference architecture of cloud-based crowdsourcing software development. More efforts will be put into the growth of the research community and production of joint research publications.

4.49 Quantum Cryptanalysis

Organizers: Serge Fehr, Michele Mosca, Martin Rötteler, and Rainer Steinwandt

Seminar No. 13371

Date: September 8–13, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.59

© Creative Commons BY 3.0 Unported license

© Serge Fehr, Michele Mosca, Martin Rötteler, Rainer Steinwandt



Participants: Aleksandrs Belovs, Daniel J. Bernstein, Johannes A. Buchmann, Andrew Childs, Frédéric Dupont-Dupuis, Serge Fehr, Katalin Friedl, Markus Grassl, Nadia Heninger, Peter Høyer, Gabor Ivanyos, Stacey Jeffery, Stephen P. Jordan, Thijs Laarhoven, Bradley Lackey, Tanja Lange, Yi-Kai Liu, Alexander May, Kirill Morozov, Michele Mosca, Maris Ozols, Youming Qiao, Martin Rötteler, Louis Salvail, Miklos Santha, Christian Schaffner, John M. Schanck, Nicolas Sendrier, Daniel Smith, Rolando Somma, Fang Song, Rainer Steinwandt, Krysta Svore, Wim van Dam, Joop van de Pol, Maarten van den Nest, Frank Wilhelm-Mauch

■ Motivation and Background

This (second) quantum cryptanalysis seminar aimed at improving our understanding of quantum attacks against modern cryptographic schemes, a task that is closely related to the question of plausible quantum computational hardness assumptions. By bringing together researchers who work in the field of quantum computing with those who work in the field of classical cryptography, the seminar aimed at identifying practical approaches to achieve cryptographic security in the presence of quantum computers. A lesson learned from an earlier edition of this seminar (Dagstuhl Seminar 11381) was that statements about the security of cryptographic schemes in the presence of a quantum attacker require the study and characterization of quantum security parameters. Those parameters measure the amount of resources that have to be spent in order to “break” a system. In this spirit, the following three topics turned out to be particularly relevant for the seminar:

- *Quantum attacks on currently deployed schemes:* Derive quantitative estimates for the resources (like no. of qubits and quantum gates) that are needed to carry out quantum attacks with cryptographically relevant parameter choices.
- *New quantum algorithms to attack potential new hardness assumptions:* For instance, can quantum algorithms be used to improve on classical solutions for computational problems in lattices or for the decoding of error-correcting codes?
- *Quantum computational assumptions:* Which problems are currently considered as intractable, even for a quantum computer, and possibly might have the potential to be of cryptographic interest? Examples are certain hidden shift and hidden subgroup problems.

One indicator for the importance of these topics for the seminar was that most talks addressed (at least) one of them. The invited group of researchers as well as the organizing team was chosen to offer a balance of expertise from the different relevant

■ Seminar Organization

The seminar involved 37 participants from around the globe, ranging from young researchers to colleagues with many years of interdisciplinary research experience. For young researchers the interdisciplinary set-up of the seminar offered an excellent opportunity to make new connections beyond the familiar research communities. Based on the experience from the predecessor (Dagstuhl Seminar 11381), we decided for a schedule which has enough flexibility to add presentations that grow out of discussions during the week, and indeed these additional slots could be brought to good use. We made an effort to keep the number of presentations limited to have ample time for open discussions between presentations. Having two research communities present at the meeting, it also seemed realistic to assume that not all participants are familiar with the latest developments in the complementing discipline. Placing survey presentations on critical topics early in the schedule was well received by the participants.

To ensure an adequate connection with the technological state-of-the-art of implementing quantum computers, one of the survey presentations was specifically devoted to this subject, and the seminar included discussions on implementation aspects of quantum computing. Keeping with the Dagstuhl tradition and the tradition of the predecessor, for Wednesday afternoon we did not schedule any presentations, allowing seminar participants to enjoy a hike in the woods, a visit to Trier, or to use the time for longer technical discussions.

■ Achievements and Next Steps

As in the first edition of this seminar, there were many fruitful discussions across discipline boundaries. At the time of writing the report for the seminar, two seminar participants had already published a preprint with a generalization of a previously known

quantum attack to a more general class of algebraic structures. We expect further publications to come forward in the coming months. While we are still far from a thorough understanding of the cryptanalytic potential of quantum computing, synergetic collaborations of seminar participants have helped greatly to advance the state-of-the-art in quantum cryptanalysis.

The seminar also successfully facilitated the exchange among colleagues from academia, government, and industry. We believe that in regard to a standardization of post-quantum cryptographic solutions, this type of exchange across community boundaries is valuable and deserves to be intensified further in future meetings.

4.50 Integration of Tools for Rigorous Software Construction and Analysis

Organizers: Uwe Glässer, Stefan Hallerstede, Michael Leuschel, and Elvinia Riccobene
Seminar No. 13372

Date: September 8–13, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.74

© Creative Commons BY 3.0 Unported license

© Uwe Glässer, Stefan Hallerstede, Michael Leuschel, and Elvinia Riccobene



Participants: Shaukat Ali, Xiaoying Bai, Xavier Blanc, Kyle Chard, Schahram Dustdar, Michael N. Huhns, Robert Kern, Donghui Lin, Greg Little, Xinjun Mao, Michael Maximilien, Dave Murray-Rust, Khrystyna Nordheimer, Dirk Riehle, Ognjen Scekcic, Lionel Seinturier, Hong-Linh Truong, Wei-Tek Tsai, Huaimin Wang, Wenjun Wu, Gang Yin, Tao Yue

■ Motivation

Dagstuhl Seminar 06191 had been a success in establishing the “ABZ” joint conference for the different state-based modelling communities (e.g., ASM, B, VDM, Z, TLA⁺) with venues in London (2008), Orford, CA (2010), Pisa (2012) and Toulouse (2014). It was a first step toward bringing these communities closer together. However, the conference, although being a place where the researchers meet, does not produce in itself a significant number of collaborations across the communities. The organisers of this seminar consider such collaborations vital in order to achieve a larger impact academically and industrially.

■ Aims of the seminar

The seminar aims to

1. Inspire exchange and joint use of formal modelling tool technologies
2. Establish long-term cross-community collaboration
3. Work towards a common vision on formal modelling

Points 2 and 3 are particularly important for future tool developments, a common methodical foundation, and more economic use of the necessary and available resources.

■ Preparation

At first the organisers intended to give the participants of the seminar case studies in advance that the participants could work on prior to the seminar to showcase their methods and tools. However, a downside of this common organisational practice is that most attendees arrive at the seminar with well-prepared, polished formal models and presentations. This would have resulted in conference-style presentations, not leaving much room for cross-community group work on problems with mixed-method

approaches. Thus, no substantial gain above and beyond what the “ABZ” conferences already accomplish would have been achieved. Hence, the organisers decided to take the risk not to ask for advance preparation but have all the work done collaboratively during the seminar. This was thought to create a more open atmosphere and leave room for discussion. The organisers chose candidates for case studies to be carried out during the seminar and asked the participants to explore solutions with diverse methods in small, often mixed, groups formed dynamically based on interests. A tentative schedule for the week was published prior to the seminar. It was adapted by the organisers every night, taking into account the actual progress by the work groups and feedback received in plenum discussions held every day in the late afternoon or evening.

■ Execution

On Sunday evening the organisers held a three hour meeting to prepare day 1 of the seminar. It was decided that, in general, evenings should be left for the participants to socialise. The case study for day 1 needed to be well-chosen to engage the participants in the seminar. It was required that

- a single problem should be treated to minimise presentation overhead necessary for explaining the model
- the problem to be solved should not be trivial but solvable within 3 hours
- the problem should not leave too much room for interpretation so that the models, methods and tools used are more readily comparable
- the problem to be solved should come with a sketch of a solution so that focus would be on the modelling activity itself and not on finding the smartest solution.

The decision was made to use the problem of “Derivation of a termination detection algorithm for distributed computations” (EWD840) by E. W. Dijkstra.

On day 1, all but one group succeeded in producing a formal model. (That group produced their model on day 2 in the after-lunch session.) At the end of day 1, the organisers felt that not enough discussion across community boundaries was taking place. This would have to be addressed in the following days. The planning for day 2, payed specific attention to this aspect. In the evening of day 1, a two-hour planning meeting among the organisers was held. It was decided that a good way of getting the different communities involved in discussions would be to reshuffle the groups of day 1 somewhat. To carry out a comparison between the methods and tools, each group would have some members that produced the original model and some “envoys” of a group that had modelled the problem in a different notation. (This turned out to work well. By lunchtime on day 2, live discussions across the community boundaries had effectively started.) On suggestion of the participants, an originally planned plenum discussion on tool integration was carried out in three groups dealing with methodology, abstract syntax and low-level integration. (The actual number of groups was decided together with all participants in the beginning of the corresponding session. Even though it may have appeared frustrating at times for some participants, the organisers thought involving all participants in some of the decision making would also improve everyone’s commitment.) The organisers also started incorporating talks. This was also considered useful for breaking the routine of the seminar.

In the evening of day 2, a one hour meeting among the organisers was held to plan day 3. It was thought that the participants could be involved closer by forming new groups that should each address a problem using two different approaches and tools. The comparison would then be possible while modelling. The modelling problem chosen was the FM’99 ATM modelling challenge. Two more talks followed on day 3 and some planning

for integration meetings that should be held in smaller groups on day 4. The latter were considered to be fruitful by many with a lot of common interests being announced. In the afternoon of day 3, a shorter hike provided a welcome break, as the weather did not invite for larger excursions.

On day 3 in the evening, a 30 minute meeting of the organisers was held. The plan for day 4 was mostly to tie up the open threads from the preceding days. A short wrap up of the case studies followed by presentations giving a comparison between two methods. In the evening, a first discussion of post-seminar work was held, discussing the Dagstuhl report and a joint book on “comprehensive modelling and modelling tools”. In the evening of day 4, a 30 minute meeting by the organisers was held. It was decided that the morning of day 5 should be spent discussing the joint book. Uwe Glässer presented an alternative case for use in the book to start a discussion about the writing approach that should be taken. (On day 5, it was decided to keep the ATM study but improve its description.)

■ Outcomes and Outlook

The main outcomes of the seminar are (a) various new collaborations across community boundaries to achieve a possible integrated use of different methods and tools, and (b) concrete plans for a book on “comprehensive modelling”, a step towards a common vision of the research field. An agreement has been reached with Springer Verlag to publish the post proceedings of the seminar in the State-Of-The-Art series of Lecture Notes in Computer Science. This should improve visibility of the effort started at the seminar.

The organisers seek to get funding, e.g., by way of a network of excellence, to continue the integration work and keep up the momentum achieved during the seminar. The aim will be to develop a common vision and a more coordinated research agenda where, in particular, resources for tool development could be used more efficiently in future.

4.51 Algorithms and Scheduling Techniques for Exascale Systems

Organizers: Henri Casanova, Yves Robert, and Uwe Schwiegelshohn

Seminar No. 13381

Date: September 15–20, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.106

© Creative Commons BY 3.0 Unported license

© Henri Casanova, Yves Robert, and Uwe Schwiegelshohn



Participants: Ismail Assayad, Guillaume Aupy, Olivier Beaumont, Anne Benoit, George Bosilca, Aurélien Bouteiller, Heinrich Braun, Henri Casanova, Anthony Danalis, Carsten Franke, Bruno Gaujal, Dominik Göddeke, Christian Grimme, Abdou Guermouche, Amina Guermouche, Thomas Hérault, Julien Herrmann, Sascha Hunold, Emmanuel Jeannot, Thilo Kielmann, Thomas Lambert, Alexey Lastovetsky, Bradley Lowery, Anthony A. Maciejewski, Loris Marchal, Alix Munier, Wolfgang E. Nagel, Jean-Marc Nicod, Thomas Rauber, Paul Renaud-Goud, Yves Robert, Gudula Rünger, Rizos Sakellariou, Erik Saule, Uwe Schwiegelshohn, Howard Jay Siegel, Oliver Sinnen, Veronika Sonigo, Stefan Turek, Bora Ucar, Frédéric Vivien

Hardware manufacturers are currently deploying machines with sustained petascale performance while already looking forward to produce Exascale machines. Exascale systems are likely to contain 10^5 to 10^6 processors, each processor itself being equipped with more than 100 cores, and possibly 10^3 to 10^4 GPU cores. These systems already reach such a degree of sophistication and complexity that the conventional approach of hardware goes first and applications follow is likely to fail. Furthermore, application performance is no longer solely defined by time-to-solution but also by power consumption and resilience to fault. Many conferences and workshops are dedicated to the architecture and systems issues pertaining to Exascale computing. Instead, in this seminar we have discussed algorithmic issues (application parallelization, application scheduling, resource management, etc.) that must be addressed to make Exascale computing a tenable proposition. As seen in many of the presentations during the seminar, core elements or principles of existing applications must be modified so that they can form the building blocks of new Exascale applications while new methods specifically targeting Exascale systems must be developed for new application areas.

The presentations during the seminar covered a wide range of topics. Some of these topics were directly targeted to various aspects of “the Exascale problem”. Some topics were targeted to components of the problem, e.g., efficient execution of application kernels on a heterogeneous many-core node. Finally, yet other topics were in broader, and often more theoretical, parallel and distributed computing contexts with less immediate but possibly large impact on the future of Exascale computing. Overall, the topics presented and discussed during the workshop can be roughly categorized as follows, noting that at least half the presentations spanned more than one of these topics:

- **Fault-tolerance.** Fault-tolerance is a major concern at large scale and several presentations focused on the limitations of current checkpoint-restart fault-tolerance techniques, provid-

ing analytical studies to quantify expected performance of these solutions and comparing them to proposed new solutions. These new solutions included, for instance, the use of algorithm-specific checkpointing combined with system-level checkpointing, or the use of imperfect fault predictors.

- **Multi-criteria optimization.** A large number of presentations presented multi-criteria optimization problems, including one traditional performance metric (throughput, makespan) and one (2-criteria) or two (3-criteria) metrics relating to power consumption and/or reliability. Several works studied the use of techniques such as DVFS to trade-off performance for a lower power consumption. These multi-criteria problems were formalized, and various theoretical and practical results were obtained in attempts to solve these problems. Two main approaches were followed: (i) optimizing one metric w.r.t. constraints on the other metric(s); or (ii) obtaining Pareto-optimal solutions or determining the entire Pareto front.
- **Multiple cores.** A handful of presentations focused on the above optimization problems not on large-scale platforms but on many-core nodes with shared memory, i.e., the intended individual components of future Exascale platforms. These nodes consist of possibly heterogeneous cores, accelerators (GPUs, etc.) connected via busses and on-chip networks.
- **Novel scheduling results.** A large number of presentations included novel findings regarding the complexity of scheduling problems. These scheduling problems are of general interest for various models of parallel computation, as motivated by the above topics. Results consisted of p-time optimal algorithms, new NP-completeness results, approximation algorithms, and efficient heuristics.
- **Exascale scientific computing.** Several presentations focused on particular scientific applications (e.g., PDE solvers) or scientific kernels (e.g., matrix multiplication),

and discussed how age-old algorithms should be adapted to exploit Exascale platforms with heterogeneous components, hierarchical networks, and the need to have both efficient and rare communication primitive invocations. One presentation presented recent experience with scalable performance monitoring and performance debugging, capabilities that will be crucial in the practice of Exascale computing.

- **Programming models for Exascale.** A handful of presentations spoke to the need for novel programming models at large scale. These presentations spanned the spectrum from very (e.g., actual implementations of programming models usable today, proposals to enhance current programming standards) to theoretical (e.g., a new theoretical approach to assess the efficiency of techniques such as work stealing and least-loaded-machine-first scheduling when the number of compute nodes tends to infinity).
- **Resource and application management.** A handful of presentations discussed Exascale computing in the context of cloud computing. In other words, these presenters made a case for applying/evolving some of the concepts currently applied in cloud deployments to future Exascale platforms (e.g., service level agreements, virtualization, resource economy). A number of open problems were identified when trying to make these two “worlds” collide.

Although the presentations at the seminar were very diverse in scope, ranging from practice to theory, an interesting observation is that many works do establish strong links between practice

(e.g., particular applications, programming models) and theory (e.g., abstract scheduling problems and results). For instance, it was found that the age-old numerical linear algebra topic, far from being well-understood, in fact gives rise to a range of interrelated and interesting practical and theoretical problems that must be solved conjointly to achieve efficiency at large scale. Such observations make it plain that forums that blends practice and theory, as is the case with this seminar, are very much needed.

The seminar brought together 41 researchers from Austria, France, Germany, Ireland, Morocco, Netherlands, New Zealand, Switzerland, U.K, and U.S.A. with interests and expertise in different aspect of parallel and distributed computing. Among participants there was a good mix of senior researchers, junior researchers, postdoctoral researchers, and Ph.D. students. Altogether there were 36 presentations over the 5 days of the seminar, organized in morning and late-afternoon sessions, plus an open problem session. The program was as usual a compromise between allowing sufficient time for participants to present their work, while also providing unstructured periods that were used by participants to pursue ongoing collaborations as well as to foster new ones. The feedback provided by the participants show that the goals of the seminar, namely to circulate new ideas and create new collaborations, were met to a large extent.

The organizers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive event.

4.52 Collaboration and learning through live coding

Organizers: Alan Blackwell, Alex McLean, James Noble, and Julian Rohrhuber
Seminar No. 13382

Date: September 15–20, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.130

© Creative Commons BY 3.0 Unported license

© Robert Biddle, Alex McLean, Alan Blackwell, and Julian Rohrhuber



Participants: Sam Aaron, Robert Biddle, Alan Blackwell, Andrew R. Brown, Luke Church, Geoff Cox, Alberto de Campo, Thomas Green, Dave Griffiths, Mark Guzdial, Ellen Harlizius-Klück, Shelly Knotts, Adrian Kuhn, Thor Magnusson, Alex McLean, David Ogborn, Jochen Arne Otto, Roly Perera, Julian Rohrhuber, Juan Gabriel Alzate Romero, Uwe Seifert, Andrew Sorensen, Jan Kees van Kampen, Renate Wieser

The goal of this seminar was to understand and develop the emerging practice, characteristics and opportunities in live coding, with an emphasis on three perspectives: the humanities, computing education, and software engineering. The opening days of the seminar were broadly structured to provide thematic introductions followed by facilitated discussions on each of these three perspectives. These were interspersed with live coding performances and experiments, in order to ensure that theoretical concerns remained grounded within this discipline that fundamentally blurs the separation of concerns between theory and practice.

The second half of the seminar was problem-oriented, resulting in concrete progress on specific technical topics, together with development of a research roadmap, publications and policy strategy to realise the significant benefits that live coding promises in a number of fields. Finally, in the spirit of both practice as a form of theory and theory as a form of practice, the seminar included some exciting musical experiences – an Algorave club night in London, with performances by delegates who were traveling from other countries on their way to the seminar; an inter-continental collaborative performance hosted jointly with the IEEE VL/HCC conference in San Jose; a conceptual proposal for an interactive sound installation in the Schloss Dagstuhl garden; and live-coded jam sessions in venues ranging from the woods of the old castle, to evening cabaret entertainment in the beautiful Dagstuhl music room.

Our main findings in relation to the three contrasting research perspectives were as follows:

1. Live coding illuminates the ways in which programming can be an artistic practice, software-as-art, going beyond a mere supporting role, and illustrating that software is itself a cultural artefact, not simply an infrastructure commodity. We

see many opportunities for nuanced, cross-disciplinary contributions to the digital humanities, for example in a revitalisation of the historical connection between computation and weaving, insights into the role of practice and experiment, and an enrichment of the notion of computation itself. Indeed, as computing becomes embedded in culture, the live, everyday authorship of computation becomes a socio-political question of freedom of speech and empowerment.

2. Live coding can play an important role in computing education, because it allows programming to be demonstrated and learned in a simple but authentic context. At the same time, it can support an *affective* teaching strategy where learners are not only motivated by the production of sound, visuals and other phenomena, but are also clear on the distinctly human activity which produces them. Thereby, however, it maintains a sense of discovery of something unanticipated and not prefigured. Of particular importance for learning is the potential for deeper engagement with the non trivial nature of computing, rather than an occupation with the operation of end-user application software.
3. Live coding offers new insights with regard to software engineering processes. The history of software engineering process can be seen as a move from heavyweight lock-step approaches to more agile approaches with fast cycles of development and feedback. At their heart, the new approaches rely on collaboration, as developers, designers, and customers work together to steer the process toward mutual success. Live coding demonstrates this kind of approach in a compelling way, with simple tools, a short time frame, but still allowing improvisational collaboration between performers and various audiences.

Perhaps more significant than any of these individual considerations is an ambitious holistic vision: that live coding can entirely change the way we think about programming. Indeed, the common experience articulated at the workshop is that live coding exemplifies both the power and the excitement of programming –

in a small space, in a short time, available and accessible to anyone. Live coding exposes the *soul* of programming.

Our next steps are a series of collaborative workshops and programs to articulate and demonstrate this collection vision of a broad and expanding role for live coding.

4.53 Algorithm Engineering

Organizers: Andrew V. Goldberg, Giuseppe F. Italiano, David S. Johnson, and Dorothea Wagner

Seminar No. 13391

Date: September 22–27, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.169

© Creative Commons BY 3.0 Unported license

© Andrew V. Goldberg, Giuseppe F. Italiano, David S. Johnson, and Dorothea Wagner



Participants: Hannah Bast, Jon Bentley, Robert E. Bixby, Gerth Stølting Brodal, Kevin Buchin, Markus Chimani, Frank Dehne, Daniel Delling, Julian Dibbelt, Rudolf Fleischer, Andrew V. Goldberg, Holger H. Hoos, Falk Hüffner, Giuseppe F. Italiano, David S. Johnson, Andrea Kappes, Jyrki Katajainen, Jürgen Lerner, Kurt Mehlhorn, Ulrich Carsten Meyer, Friedhelm Meyer auf der Heide, Henning Meyerhenke, Shin-ichi Minato, Rolf H. Möhring, Matthias Müller-Hannemann, Petra Mutzel, Patrick K. Nicholson, Yoshio Okamoto, Marina Papatrifaftilou, Alejandro Salinger, Peter Sanders, Federico Santaroni, Sabine Storandt, Dorothea Wagner, Roger Wattenhofer, Renato Werneck, Christos Zaroliagis, Liang Zhao, Katharina A. Zweig

■ Topics of the Seminar

The seminar covered all methodological aspects of algorithm engineering. Examples are the scientific method in algorithmics, the use of modern computer architecture in algorithmics, and certifying algorithms. These aspects were also addressed in dedicated discussion sessions.

Science of Algorithmics. One aspect of algorithm engineering is the *scientific method*, where research on algorithms is interpreted as in other disciplines such as physics and life sciences: the observation of a phenomenon that is not yet understood is investigated via falsifiable hypotheses as explanations of the phenomena, and experimental evaluations to test these hypotheses. That way not only empirical evidence on the behavior of algorithms is attained but also new theoretical insights are sought. Experimental algorithmics is already a core component of algorithm engineering from its very beginning. However, the design of reasonable experiments, the use of meaningful test instances, and reproducibility of experiments are still issues to be discussed in order to derive a common understanding and agree on a best practice.

Manycore and GPU Algorithms. Exploiting the full potential of a *modern computer* poses many interesting new challenges for algorithm engineering: ever increasing parallelism, deep memory hierarchies, and heterogeneous architectures. Algorithms should be tailored to utilize multiple cores, but also access memory efficiently, taking into account issues such as data locality. Nowadays the use of GPUs, which are increasingly common in modern servers, is an important issue for efficient algorithm implementation. This is in particular interesting for frequently used and “classical” algorithms.

Certifying Algorithms. An effective way to ensure correct results of algorithm implementations are *certifying algo-*

rithms. The idea is to check each returned result for correctness using a simple checker. It then suffices to test or perhaps verify the checker. Making checking fast implies interesting algorithmic questions when checking is aided by certificates of correctness computed by the main algorithm.

■ Focus Topic: Web Search and Large Graphs.

Experiences from previous Dagstuhl seminars showed that the interaction between different scientific communities stimulates methodological discussions. This exchange is in particular important for neighboring scientific communities who typically meet at separate conferences. For this seminar, we focus on *web search, large graphs and social networks* in order to also address the scientific WWW and Social Media community. In these fields, methods from algorithm engineering are applied. However, these scientists typically don’t publish at the algorithm engineering conferences mentioned above but meet and publish at conferences like the “International World Wide Web Conference”, the “ACM Conference on Hypertext and Hypermedia” or the “International AAAI Conference on Weblogs and Social Media”.

Search engines work with a large amount of data, making high-performance algorithms and data structures very important. Relevant problems include fast indexing, text and query processing, and relevance computation. The latter involves a large web graph. Web-enabled applications give rise to other large graphs, such as social networks, like “friend graphs” or e-mail graphs induced by message origin-destination pairs. Algorithms on such graphs are of great interest. For example, identifying interest-based sub-communities (e.g., classical music fans) enables better service experience or contextual advertisement.

■ Aims

The aim of this seminar was to bring together researchers with different backgrounds, e.g., from algorithm and datastructures,

computational geometry, combinatorial optimization, parallel algorithms and algorithm engineering in order to strengthen and foster collaborations and to identify key research directions for the future. In particular, the seminar was intended to foster the exchange between algorithm engineering and scientists from the web search community. While the dominant goal of the seminar was the exchange of current research developments and discussion of topical subjects, it also contributed to bring algorithm engineering forward as a still evolving and expanding field in computer science. The seminar program included four dedicated discussion sessions on methodological questions, as well as research related issues like future DIMACS Implementation Challenges.

■ Conclusion

The organizers decided to schedule talks and discussions not grouped according to topics but provide a vivid mix of different research questions and results. According to the composition of the seminar participants, not all topics were covered equally well.

For example, certifying algorithms were not addressed in detail. On Monday, Renato Werneck gave a short report on the “11th DIMACS Implementation Challenge on Steiner Tree Problems”¹¹ taking place in 2013/14. The program of Monday afternoon was concluded by a panel discussion on “Empirical and Theoretical Approaches to Algorithm Design: Synergies and Opportunities”. The second panel discussion on “Benchmarks and reproducibility of experiments” took place on Tuesday. The third panel discussion on Thursday focused on “Promoting and advancing the field” and on Friday a discussion about “Teaching Algorithm Engineering” concluded the program.

The seminar hosted 39 participants. Besides presentations and panel discussions the program offered room for bilateral discussions and working groups. Schloss Dagstuhl and its staff provided a very convenient and stimulating environment. The seminar participants appreciated the cordial atmosphere which improved mutual understanding and inspiration. The organizers of this seminar wish to thank all those who helped make the workshop a fruitful research experience.

¹¹ <http://dimacs11.cs.princeton.edu/home.html>

4.54 Inter-Vehicular Communication – Quo Vadis

Organizers: Onur Altintas, Falko Dressler, Hannes Hartenstein, and Ozan K. Tonguz

Seminar No. 13392

Date: September 22–25, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.190

© Creative Commons BY 3.0 Unported license

© Onur Altintas, Falko Dressler, Hannes Hartenstein, and Ozan K. Tonguz



Participants: Natalya An, Claudio Casetti, Wai Chen, Falko Dressler, David Eckhoff, Andreas Festag, Raphaël Frank, Mario Gerla, Javier Manuel Gozalvez Sempere, Marco Gruteser, Jérôme Härri, Hannes Hartenstein, Geert Heijnen, Liviu Iftode, Stefan Jörer, Frank Kargl, Renato Lo Cigno, Giovanni Pau, Jonathan Petit, Björn Scheuermann, Florian Schimandl, Michele Segata, Christoph Sommer, Tessa Tielert, Ozan K. Tonguz, Elisabeth Uhlemann, Peter Vortisch

■ Motivation

The management and control of network connections among vehicles and between vehicles and an existing network infrastructure is currently one of the most challenging research fields in the networking domain. Using the terms Vehicular Ad-hoc Networks (VANETs), Inter-Vehicle Communication (IVC), Car-2-X (C2X), or Vehicle-2-X (V2X), many applications – as interesting as challenging – have been envisioned and (at least) partially realized. In this context, a very active research fields has developed.

There is a long list of desirable applications that can be grouped into four categories:

- eSafety applications that try to make driving safer, e.g., road hazard warning;
- traffic efficiency applications aiming at more efficient and thus greener traffic, e.g., detection of traffic jams;
- manufacturer oriented applications, e.g., automatic software updates; and
- comfort and entertainment applications, e.g., automatic map updates or video streaming.

While there are some similarities with fields like mobile ad-hoc networks or wireless sensor networks, the specific characteristics of vehicular networks require different communication paradigms, different approaches to security and privacy, or different wireless communication systems. For example, the nodes usually do not have severe power and form factor constraints, and they might be always on. On the other hand, due to high relative speeds, wireless connections may not be stable for a longer time period and the network density is expected to vary from sparse to very dense networks. Another challenging issue is the efficient use of available infrastructure, such as road side units or even cellular networks. Furthermore, IVC has strong links to other research domains, e.g., geo-informatics as it requires very precise localization and precise maps or highly scalable simulations that

are a requirement for analyzing traffic systems with hundreds or thousands of vehicles.

In the past, many specific solutions for IVC have been identified and now, industry and other stake-holders are already calling for standardization. Still, we believe that many important research questions have only been partially answered and the approaches discussed in the standardization bodies are based only on a minimum consensus of simplest solutions. Security and privacy, scalability, use of advanced communication patterns like aggregation, transmit power control, and optimal medium access are just a few of such issues.

In 2010, a first Dagstuhl Seminar (10402) was organized on the topic of inter-vehicular communication [1,2]. The motivation was to bring together experts in this field to investigate the state of the art and to highlight where sufficient solutions already existed. The main outcome of this very inspiring seminar was that there are indeed areas within this research where scientific findings are being consolidated and adapted by industry. This was the consensus of quite intriguing discussions among participants from both industry and academia. Yet, even more aspects have been identified where substantial research is still needed. These challenges have been summarized in the Dagstuhl report [1] and in an IEEE Communications Magazine article [2].

■ Objectives

It was the goal of this new seminar to again bring together leading researchers both from academia and industry to discuss if and where the previously identified challenges have been adequately addressed, and to highlight where sufficient solutions exist today, where better alternatives need to be found, and also to give directions where to look for such alternatives. Furthermore, the goal of this workshop was to go on step beyond and identify where IVC can contribute to the basic foundations of computer science

or where previously unconsidered foundations can contribute to IVC.

The 2010 Dagstuhl seminar promoted a “top-down” approach to inter-vehicle communications instead of the classical “bottom-up” approach. With the top-down approach, the effects of applications are first analyzed under the assumption that the communication system will be able to support the application. Thus, an “upper bound” can be presented on the benefits of IVC. In our discussions, we summarized all the scientific work that followed this approach after the previous Dagstuhl seminar and contrasted it with new insights based on field operational tests, safety application design and massively distributed operations.

In particular, we shifted the focus from basic networking principles to applicability in real world scenarios. In the last few years, first field operational tests have been conducted in the US (the Michigan field trial) as well as in Europe (SIM-TD in Germany, DRIVE C2X in Europe). Lessons learned from those tests applied to currently used models and concepts will bring new insights into the forthcoming research challenges. Among others, questions to be studied include the following still unanswered research challenges:

- Data analysis of current field operational tests: are they validating or invalidating current models?
- Safety applications: show stopper or driving force? What are the limitations in terms of latency and reliability of available communication principles for enabling critical safety support;
- From highly distributed to massively distributed operation: can vehicular networking based on DSRC/WAVE also support all the pedestrians and bicyclists?

We organized the 2013 seminar again as a discussion forum. Three invited keynote presentations were organized to stimulate discussions among the participants. In order to steer the discussions, we prepared four working groups that helps focusing on selected open research challenges. In addition, we also supported ad-hoc presentations on topics of the working groups. The following working groups have been formed and led to very interesting observations:

- Foundations – In this group, it was discussed, which fundamental insights gained in the vehicular networking research domain can be transferred to other domains of computer science. The other way around has been discussed as well, i.e., which areas of computer science might help fostering work in the vehicular networking and which may help overcoming open challenges.
- Field Operational Tests (FOTs) – This group focused on the results that already have been derived from the ongoing work in various test sites in the U.S. and in Europe. The main questions in the discussion were whether the current experiments are already sufficient to gain insights into larger scale behavior or if additional tests are needed.
- IVC Applications – In this group, the applications’ perspective to IVC was discussed. In the last years, many of the developments have been done looking at lower layer networking problems. This resulted in a number of networking solutions that nicely support specific applications but cannot be integrated to a generalized networking architecture.
- Heterogeneous Networks – Possibly one of the most important and timely working groups focused on the integration of different networking technologies. This is strongly needed to develop integrated IVC solutions and also to overcome early deployment problems like the initially low penetration ratio.

Eventually, all these questions lead to the big question whether vehicular networking can now be shown to improve efficiency and safety on our streets. We are now in an era that completely changes the game in car manufacturing and road traffic management. Computer science is becoming the key element in the design of these systems. It is of utmost importance to bring in expertise from classical computer science (computer networking, simulation and modeling, operating system design) as well as from electrical engineering (digital signal processing, communication networks) as well as experts from the automotive industry and from the intelligent transportation community.

References

- 1 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz, and Lars Wischhof, “Executive Summary – Inter-Vehicular Communication,” in *Dagstuhl Seminar 10402 – Inter-Vehicular Communication*. Schloss Dagstuhl, Wadern, Germany: Schloss Dagstuhl, October 2010. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2011/2929/>
- 2 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz and Lars Wischhof, "Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar," *IEEE Communications Magazine*, vol. 49 (5), pp. 158-164, May 2011.

4.55 Automatic Application Tuning for HPC Architectures

Organizers: Siegfried Benkner, Franz Franchetti, Hans Michael Gerndt, and Jeffrey K. Hollingsworth
Seminar No. 13401

Date: September 29 to October 4, 2014 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.214

© Creative Commons BY 3.0 Unported license

© Siegfried Benkner, Franz Franchetti, Hans Michael Gerndt, and Jeffrey K. Hollingsworth



Participants: Enes Bajrovic, Shajulin Benedict, Siegfried Benkner, Aydın Buluç, Milind Chabbi, I-hsin Chung, Isaias Alberto Compres Urena, Guojing Cong, Thomas Fahringer, Franz Franchetti, Grigori Fursin, Hans Michael Gerndt, Carla Guillen, Torsten Höfler, Jeffrey K. Hollingsworth, Paul D. Hovland, Toshiyuki Imamura, Thomas Karcher, Takahiro Katagiri, Michael Knobloch, Andreas Knüpfer, Jakub Kurzak, Allen D. Malony, Andrea Martinez, Renato Miceli Costa Ribeiro, Robert Mijakovic, Bernd Mohr, Shirley V. Moore, Carmen Novarrete, Georg Ofenbeck, Antonio Pimenta, Barry Rountree, Martin Sandrieser, Robert Schoene, Martin Schulz, Armando Solar-Lezama, Walter F. Tichy, Jesper Larsson Träff, Richard M. Veras, Richard Vuduc, Felix Wolf

Parallel computer systems especially for High Performance Computing are getting increasingly complex. The reasons are manifold. HPC systems today with a peak performance of several petaflops have hundreds of thousands of cores that have to be able to work together efficiently. Those machines have a deep hierarchy, which has to be understood by the programmer to tune his program so that it profits from higher interconnection rates. In addition, to reduce the power consumption of those systems, advanced hard- and software techniques are applied, such as the usage of GPUs that are highly specialized for regular data parallel computations via simple compute cores and high bandwidth to the graphics memory. Another technique is to reduce the clock frequency of processors when appropriate, e.g. when the application or phases of the execution are memory bound. This transforms a homogeneous system into a heterogeneous system, which complicates programming tasks such as load balancing and efficient communication.

The complexity of today's parallel architectures has a significant impact on the performance of parallel applications. Due to the high amount of energy and money being lost because of the low processor utilization, application developers are now investing significant time to tune their codes for the current and emerging systems. This tuning is a cyclic process of gathering data, identifying code regions that can be improved, and tuning those code regions.

There are a growing number of autotuning researchers in Europe, the United States, and Asia. However, there are relatively few opportunities for these researchers to meet together. The unique format of a Dagstuhl seminar provides the opportunity to bring together researchers from around the world that are using different approaches to autotuning.

This workshop brought together those people working on autotuning with people working on performance analysis tools. While the analysis tools indicate performance problems, their

combination with performance tuning might make those tools even more successful. The presentations of experts in both areas will increase the interest and the knowledge of the techniques applied in the other area. It will steer future collaborations and might also lead to concrete ideas for coupling performance analysis and performance tuning tools.

The workshop was driven by the European FP7 project AutoTune that started on October 15th, 2011. It is the goal of AutoTune to implement the Periscope Tuning Framework based on the automatic performance analysis tool Periscope. It will couple Periscope's performance analysis with performance and energy efficiency tuning in an online approach.

Performance Analysis. Performance analysis tools support the programmer in the first two tasks of the tuning cycle. Performance data are gathered during program execution by monitoring the application's execution. Performance data are both summarized and stored as profile data or all details are stored in so called trace files. In addition to application monitoring, performance analysis tools also provide means to analyze and interpret the provided performance data and thus to detect performance problems. The analysis is either supported by graphical display or by annotating the source code.

State of the art performance analysis tools fall into two major classes depending on their monitoring approach: profiling tools and tracing tools. Profiling tools summarize performance data for the overall execution and provide information such as the execution time for code regions, number of cache misses, time spent in MPI routines, and synchronization overhead for OpenMP synchronization constructs. Tracing tools provide information about individual events, generate typically huge trace files and provide means to visually analyze those data to identify bottlenecks in the execution.

Representatives for these two classes are gprof, OMPP and Vampir. Gprof is the GNU Profiler tool. It provides a flat

profile and a callpath profile for the program's functions. The measurements are done by instrumenting the application. OmpP is a profiling tool for OpenMP developed at TUM and the University of Tennessee. It is based on instrumentation with Opari and determines certain overhead categories of parallel regions. In contrast to the previous two tools, Vampir is a commercial trace-based performance analysis tool from Technische Universität Dresden. It provides a powerful visualization of traces and scales to thousands of processors based on a parallel visualization server.

The major research challenges in the development of PA tools are to automate the analysis and to improve the scalability of the tools. Automation of the analysis is important to facilitate the application developer's task. Starting from the formalization of performance properties in the European-American working group APART (<http://www.fz-juelich.de/apart>), automatic performance analysis tools were developed. Paradyne from University of Wisconsin was the first automatic online analysis tool. Its performance consultant guided the search for performance bottlenecks while the application was executing. The most important representatives are SCALASCA and Periscope. SCALASCA is an automatic performance analysis tool developed at Forschungszentrum Jülich and the German Research School on Simulation Sciences. It is based on performance profiles as well as on traces. The automatic trace analysis determines MPI wait time via a parallel trace replay on the application's processors after the application execution terminated.

Periscope is an automatic performance analysis tool for highly parallel applications written in MPI and/or OpenMP currently under development at Technische Universität München. It is a representative for a class of automatic performance analysis tools automating the whole analysis procedure. Unique to Periscope is that it is an online tool and it works in a distributed fashion. This means that the analysis is done while the application is executing (online) and by a set of analysis agents, each searching for performance problems in a subset of the application's processes (distributed). The properties found by Periscope point to code regions that might benefit from further tuning.

Performance Autotuning. The central part of the tuning process is the search for the best combination of code transformations and parameter settings of the execution environment. This creates an enormous search space, which further complicates the whole tuning task. As a result, much research has been dedicated to the area of autotuning in the last years and many different ideas have been gathered. These can be grouped into four categories:

- self-tuning libraries for linear algebra and signal processing like ATLAS, FFTW, OSKI and SPIRAL;
- tools that automatically analyze alternative compiler optimizations and search for their optimal combination;
- autotuners that search a space of application-level parameters that are believed to impact the performance of an application;
- frameworks that try to combine ideas from all the other groups.

The first category contains special purpose libraries that are highly optimized for one specific area. The Automatically Tuned Linear Algebra Software (ATLAS) supports the developers in creating numerical programs. It automatically generates and optimizes the popular Basic Linear Algebra Subroutines (BLAS) kernels for the currently used architecture. Similarly, FFTW is a library for computing the discrete Fourier transform on different systems. Due to the FFTW design, an application using it will perform well on most architectures without modification.

However, the growing diversity of parallel application areas requires a more general autotuning strategy. Thus, substantial research has been done in a different application-independent approach of autotuning. This is based on the automatic search for the right compiler optimizations on the specific platform. Such tools can be separated into two groups according to their methodology: iterative search tools and those using machine learning techniques. There has been much work in the first category. All these tools share the idea of iteratively enabling certain optimizations. They run the compiled program and monitor its performance. Based on the outcome, they decide on the new tuning combination. Due to the huge size of the search space, these tools are relatively slow. There exists an algorithm called combined elimination (CE) that greatly improves the previous search-based methods.

The second branch of compiler-based autotuners applies a different strategy to look for the best optimization settings. They use knowledge about the program's behavior and machine learning techniques to select the optimal combination. This approach is based on an automatically built per-system model, which maps performance counters to good optimization options. This model can then be used with different applications to guide their tuning. Current research work is also targeting the creation of a self-optimizing compiler that automatically learns the best optimization heuristics based on the behavior of the underlying platform.

Among the tools in the third category is the Active Harmony system. It is a runtime parameter optimization tool that helps focus on the application-dependent parameters that are performance critical. The system tries to improve performance during a single execution based on the observed historical performance data. It can be used to tune parameters such as the size of a read-ahead buffer or what algorithm is being used (e.g., heap sort vs. quick sort). As compared with Active Harmony, the work from Nelson uses a different approach that interacts with the programmer to get high-level models of the impact of parameter values. These models are then used by the system to guide the search for optimization parameters. This approach is called model-guided empirical optimization where models and empirical techniques are used in a hybrid approach.

Popular examples for the last group of autotuning tools are the newly released Parallel Active Harmony, and the Autopilot framework. The Parallel Active Harmony is a combination of the Harmony system and the CHILL compiler framework. It is an autotuner for scientific codes that applies a search-based autotuning approach. While monitoring the program performance, the system investigates multiple dynamically generated versions of the detected hot loop nests. The performance of these code segments is then evaluated in parallel on the target architecture and the results are processed by a parallel search algorithm. The best candidate is integrated into the application. The second popular example in this group is the Autopilot. It is an integrated toolkit for performance monitoring and dynamical tuning of heterogeneous computational grids based on closed loop control. It uses distributed sensors to extract qualitative and quantitative performance data from the executing applications. This data is processed by distributed actuators and the preliminary performance benchmark is reported to the application developer.

Energy efficiency autotuning. Multi-Petascale supercomputers consist of more than one hundred thousand processing cores and will consume many MW of electrical power. Energy efficiency will be crucial for both cost and environmental reasons, and may soon become as important as pure peak performance. This is exemplified by the fact that since a few years the TOP500

list (<http://www.top500.org/>) also contains power consumption values. Current procurements for high-end supercomputers show that the cost for electricity and cooling is nearly as high as for the hardware, particularly in countries with high energy costs such as Germany. Power consumption is considered one of the greatest challenges on the road to exascale systems.

Dynamic frequency and voltage scaling provides a mechanism to operate modern processors across a broad range of clock frequencies and voltage levels, allowing to trade off performance vs. energy consumption. Overall frequency scaling ideas are based on Advanced Configuration and Power Interface (ACPI, <http://www.acpi.info/>) specification with Intel's SpeedStep implementation or Cool'n'Quiet by AMD, respectively. Processors like Intel's Sandy Bridge are fully compliant with ACPI. Sets of utilities to exploit these techniques are available, and ideas to use them for complete jobs in multi user HPC clusters have already been described.

Whereas dynamic frequency scaling is commonly used in laptops, the impact and usability in HPC is still quite challenging. For applications using several hundreds or thousands of cores, uncoordinated manipulation of the frequency by some background daemon would introduce a new source of OS jitter. Moreover, changing the processor frequency requires on the order of milliseconds and only yields a benefit if a major part of an application can be run in a given mode continuously. Typically, lowering the CPU frequency can yield a 10% decrease in power consumption while increasing the application runtime by less than 1%. However, the impact of lowering the frequency and voltage on the application performance depends on whether it is CPU, memory, cache or I/O bound. Code regions that are CPU or cache bound can take advantage of higher frequencies, whereas regions that are memory or I/O bound experience only minor performance impacts when reducing the frequency. Therefore it is essential to identify applications and those parts of them that are appropriate for running within a specific power envelope without sacrificing too much performance.

Different metrics for performance, cost, energy, power, cooling and thermal conditions may apply for different usage and optimization scenarios e.g.

- minimizing the energy consumption by reducing the performance of an application by a given percentage
- considering outside temperature conditions, i.e. if it is cold outside and free cooling is applied, an increased power consumption by the compute nodes might be tolerated
- optimizing the total cost of ownership (including baseline investment, power and cooling) for given throughput requirements.

It is quite cumbersome to investigate all these conditions and the various frequency settings manually. Therefore automatic tools are required to automatically identify suitable applications and particular code regions, and finally automatically tune the frequency and power settings to yield optimal results for the desired objectives.

Thematic Sessions. The seminar was organized as a series of thematic sessions. An initial session comprised two overview presentations about performance analysis and measurement tools as well as a general introduction to autotuning, setting the overall context for the seminar. A session on support tools covered code restructuring techniques, testing environments, and performance repositories for autotuning. Two sessions on infrastructures provided insights into frameworks and environments, language support for autotuning as well challenges and requirements in the context of very large-scale systems. A session on energy efficiency tuning gave insight into the challenges and recent developments in optimizing HPC systems and applications with respect to energy consumption. A session on accelerator tuning covered various issues in tuning for GPUs and accelerated parallel systems. A session on techniques covered various topics related to performance-guided tuning, modeling, and scalability. A session on tools covered recent developments in empirical autotuning, semantics support for performance tools and autotuners as well as synthesis of libraries. Various topics related to the tuning of message-passing applications and I/O-related autotuning were covered in a session on MPI and I/O tuning. The session on compiler transformations covered compiler transformations for multi-objective tuning, techniques for tuning irregular applications, as well as on language and compilation support for analysis of semantic graphs.

4.56 Physical-Cyber-Social Computing

Organizers: Amit P. Sheth, Payam Barnaghi, Markus Strohmaier, Ramesh Jain, and Steffen Staab

Seminar No. 13402

Date: September 29 to October 4, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.9.245

© Creative Commons BY 3.0 Unported license

© Amit P. Sheth, Payam Barnaghi, Markus Strohmaier, Ramesh Jain, and Steffen Staab

Participants: Rajendra Akerkar, Harith Alani, Pramod Anantharam, Payam M. Barnaghi, Ciro Cattuto, Benoit Christophe, Edward Curry, Emanuele Della Valle, Schahram Dustdar, Frieder Ganz, Michael Granitzer, Manfred Hauswirth, Laura Hollink, Andreas Hotho, Geert-Jan Houben, Ramesh Jain, Mohan S. Kankanhalli, Artem Katasonov, Claudia Müller-Birn, Axel Polleres, RangaRao Venkatesha Prasad, Matthew Rowe, Amit P. Sheth, Vivek K. Singh, Steffen Staab, Markus Strohmaier, Kerry Taylor, Josiane Xavier Parreira, Koji Zettsu



Miniaturisation, progress with energy issues and cost reductions have resulted in rapid growth in deployment of networked devices and sensing, tightly connecting the physical world with the cyber-world as well as interconnected humans bringing along them virtual social interactions.. The number of devices connected to the Internet already exceeds the number of people on earth and is estimated to grow to 50 billion devices by 2020. The resulting system called Internet of Things (IoT) incorporates a number of technologies including wireless sensor networks, pervasive computing, ambient intelligence, distributed systems and context-aware computing. With growing adoption of smart-phones and social media, citizens or human-in-the-loop sensing and resulting user generated data and data generated by user carried devices have also become key sources of data and information about the physical world and corresponding events. Data from all these sources will result in tremendous volume, large variety and rapid changes (velocity). The combination of cyber-physical and social data can help us to understand events and changes in our surrounding environments better, monitor and control buildings, homes and city infrastructures, provide better healthcare and elderly care services among many other applications. To make efficient use of the physical-cyber-social data, integration and processing of data from various heterogeneous sources is necessary. Providing interoperable information representation and extracting actionable knowledge from deluge of human and machine sensory data are the key issues. We refer to the new computing capabilities needed to exploit all these types of data to enable advanced applications as physical-cyber-social computing.

4.57 Deduction and Arithmetic

Organizers: Nikolaj Bjørner, Reiner Hähnle, Tobias Nipkow, and Christoph Weidenbach
Seminar No. 13411

Date: October 6–11, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.1

© Creative Commons BY 3.0 Unported license

© Nikolaj Bjørner, Reiner Hähnle, Tobias Nipkow, and Christoph Weidenbach



Participants: Wolfgang Ahrendt, Jeremy Avigad, Peter Baumgartner, Bernhard Beckert, Nikolaj Bjørner, Jasmin Blanchette, Richard Bubel, Eva Darulova, Leonardo de Moura, Stephan Falke, Pascal Fontaine, Martin Fränzle, Carsten Fuhs, Jürgen Giesl, Alberto Griggio, Arie Gurfinkel, Reiner Hähnle, James J. Hunt, Dejan Jovanovic, Tim A. King, Konstantin Korovin, Marek Kosta, Jerome Leroux, Assia Mahboubi, Aart Middeldorp, David Monniaux, Wojciech Mostowski, Tobias Nipkow, Grant Olney Passmore, Ruzica Piskac, André Platzer, Enric Rodríguez-Carbonell, Philipp Rümmer, Peter H. Schmitt, Carsten Sinz, Viorica Sofronie-Stokkermans, Christian Sternagel, Thomas Sturm, René Thiemann, Cesare Tinelli, Tjark Weber, Christoph Weidenbach

Arithmetic plays a fundamental role in deduction. Logical constraints over arithmetical properties occur frequently in classical theorems in mathematics, as well as in program analysis and verification. The first automatic theorem prover was an implementation of Presburger Arithmetic in 1954. With the availability of powerful predicate calculus proof procedures some years later, arithmetic would be relegated to the sidelines. Interest in arithmetic revived in the 1980s with the advent of powerful interactive theorem provers that needed and supported arithmetic for their applications. The need for efficient computer aided deduction with support for arithmetic in the area of program analysis and verification recently gave birth to a new technology, so called SMT solvers.

Thus we have three strands of automated deduction: SMT solvers, automated first-order provers, and interactive provers in need of (more) arithmetic.

SMT: SMT (satisfiability modulo theories) solvers distinguish themselves by integrating built-in support for a combination of theories, including prominently the theory of arithmetic. Most often handling arithmetic formulas in isolation is not sufficient. Applications typically use a non-disjoint combination of arithmetic and other theory reasoning. SMT solvers nowadays handle quantifier-free arithmetic well, but are not directly equipped to solve arithmetical formulas with quantifiers. Recent progress on building in quantifier-elimination procedures for linear and non-linear arithmetic have made practical integration of such richer arithmetic deduction viable.

ATP: Research in first-order logic theorem proving used to concentrate on efficient calculi in general and the integration of equational theories in particular. It is obvious that further integration of “richer” arithmetic theories into

first-order logic should be done by rather a combination approach than an integration approach. One major challenge of combining first-order logic calculi with arithmetic procedures is that of compactness/completeness and termination. While Boolean combinations of ground atoms, as they are considered by SMT solvers typically do not cause trouble with respect to those challenges, combining first-order clauses with an arithmetic theory can never result in a compact/complete/terminating calculus, in general. The actual combination typically requires the solution of purely arithmetic problems in order to establish valid inferences and simplifications. These problems are of a specific nature in that the form of the arithmetic formulas and the way they need to be tested require specific variants of the known arithmetic procedures.

ITP: Interactive theorem provers initially came with built-in decision procedures for quantifier-free linear arithmetic. More foundational systems then developed new techniques to implement these decision procedures by reducing them to pure logic, trading efficiency for guaranteed correctness. Aspects of arithmetic reasoning are present in deductive software verification systems: interactive systems combine a number of automatic arithmetic reasoning methods and control them with heuristics that are specific for verification. A challenging application of interactive proof and arithmetic is the Flyspeck project, an effort to formalize Hales’s proof of the Kepler conjecture in an interactive theorem prover.

The Dagstuhl seminar was a timely event that brought together experts in the above subareas of deduction, and in reasoning about arithmetic, to exchange experiences and insights. The research questions pursued and answered included:

- Which arithmetic problems are best solved with which approach?
- How to handle very complex numeric representations such as the IEEE floating-point standard with a high degree of automation?
- Arithmetic in combination with other theories results easily in languages with a very complex decision problem—how can a high degree of automation be obtained nevertheless?
- How can SMT-based reasoning be combined with model-based reasoning?

- What is the best way to incorporate arithmetic simplification available in computer algebra systems into deductive frameworks?
- How can the specific structure of arithmetic problems generated by deduction systems be exploited?

In addition to the technical contributions, the seminar participants attempted in an open discussion session to identify the major trends and open questions around Deduction and Arithmetic.

4.58 Genomic Privacy

Organizers: Kay Hamacher, Jean Pierre Hubaux, and Gene Tsudik
Seminar No. 13412

Date: October 6–9, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.25

© Creative Commons BY 3.0 Unported license

© Jean Louis Raisaro



Participants: Gergely Acs, Erman Ayday, Marina Blanton, Jurgi Camblong, Emiliano De Cristofaro, Zekeriya Erkin, Sky Faber, Jacques Fellay, Kay Hamacher, Urs Hengartner, Zhicong Huang, Jean Pierre Hubaux, Mathias Humbert, Murat Kantarcioglu, Florian Kerschbaum, Regine Kollek, Klaus A. Kuhn, Inald Lagendijk, Bradley Malin, Srdan Marinovic, Satoru Miyano, Andrzej Mizera, Muhammad Naveed, Andreas Peter, Jean-Jacques Quisquater, Jean-Louis Raisaro, Roded Sharan, Vitaly Shmatikov, Amalio Telenti, Gene Tsudik, Xiaofeng Wang

The Dagstuhl seminar 13412 “Genomic Privacy” was a short two-and-a-half-day seminar, the first one on this topic ever, which took place from October 6th to 9th, 2013. The aim was to bring together researchers, from various research areas related to genomic privacy, and to inspire them to exchange theoretical results, practical requirements and ethical and legal implications related to the protection of genomic data. The rise of personalized medicine on the background of available, individual genomic sequences is taken for granted in the biomedical community. Impressive advances in genome sequencing have opened the way to a variety of revolutionary applications in modern healthcare. In particular, the increasing understanding of the human genome, and of its relation to diseases and its response to treatments brings promise of improvements in preventive and personalized healthcare. However, because of the genome’s highly sensitive nature, this progress raises important privacy and ethical concerns that need to be addressed. Indeed, besides carrying information about a person’s genetic condition and his predisposition to specific diseases, the genome also contains information about his relatives. The leakage of such information can open the door to a variety of abuses and threats not yet fully understood. During the seminar, these points were addressed in particular:

- **Expression and Requirements:** What should be protected? For how long? Against whom? Who should be liable? Who would manage cryptographic keys? Anonymity vs. cryptography?
- **Privacy Mechanisms & Regulations:** What privacy enhancing techniques can be envisioned specifically for genomic data? What if some people publish their genome online against the will of their relatives? Which ethical guidelines can be adopted from traditional privacy regulations?
- **Medical Perspective:** Would medical specialists accept to have only a partial view on genomic data? How are epidemio-

logical studies and biobanks affected by legal and/or technical restrictions?

- **Patient Perspective:** What patient’s involvement can be reasonably expected? Can a person’s genomic information be outsourced to some cloud storage service?
- **Economics:** What are the economic implications of genomic privacy; of its neglect?

The seminar fully satisfied the expectations. All participants briefly self-introduced themselves. Some of them were invited by the organizers to give survey talks about their recent research on genomic privacy, thus facilitating and encouraging inter-disciplinary discussions during the morning sessions. There were talks focusing both on the definition of the requirements for the efficient and secure implementation of genomic medicine and on the possible solutions to be addressed. The afternoon sessions were devoted to working groups.

The first speaker Regine Kollek (University of Hamburg, GER), addressed the meaning and context of genomic data, focusing on some of the social and ethical aspects of genomics. She was followed by Brad Malin (Vanderbilt University, US), who provided a summary of the ways (both legal and technical) such data can be protected and raised the question about its worth or if there exists some other practical approach that guarantees flexible genomic data protection plans. Satoru Miyano (University of Tokyo, JP) gave an overview of the requirements in term of storage, computational power and security needed to make “clinical sequencing” become a reality. He described the ongoing program that has been playing a key role in the International Cancer Genome Consortium (ICGC) in Japan. The morning session ended up with a joint talk by Jacques Fellay (EPFL – Lausanne, CH) and Amalio Telenti (Lausanne University Hospital, CH) about the current and future usage of genomic information in clinical settings. They outlined the importance of

defining new threat models, emphasizing that trust is essential in healthcare.

The second day was focused on the possible technical solutions that can be used to ensure genomic privacy. If, on one hand, there are computational expensive cryptographical approaches such as homomorphic encryption or secure multi-party computation that guarantee accuracy at the expense of flexibility and increasing complexity, on the other hand there are also statistical-based solutions such as differential privacy, which are less accurate but more flexible and less expensive in terms of computational and complexity costs. The first speaker of the day, Andreas Peter (University of Twente, NL), described his ongoing work on how to securely outsource genomic sequences in a privacy-preserving way by relying on an oblivious RAM construction. The second talk, by Erman Ayday (EPFL – Lausanne, CH), provided an overview of the activities on genomic privacy in Lausanne. Ayday first focused on how to protect and evaluate genomic privacy in the clinical context, he then showed how to process in a privacy-preserving fashion raw genomic data; and finally he described how to quantify kin genomic privacy. The third speaker of the morning, Vitaly Shmatikov (University of Texas – Austin, US), discussed about how to conduct privacy-preserving exploration in Genome-Wide Association Studies (GWAS). He presented a set of privacy-preserving data mining algorithms

that produce significantly accurate results while guaranteeing differential privacy. Finally, the second-day morning session was closed by Emiliano De Cristofaro's (University College London, GB) survey about how to begin to address privacy-respecting genomic tests by relying on privacy-enhancing techniques based on private set intersection operations.

The final day started with a talk by Xiaofeng Wang (Indiana University – Bloomington, US) about the privacy-preserving sharing and analysis of human genomic data. In particular, he described some techniques for secure outsourcing of genome analysis, and differentially-private pilot data release and data source selection. The remaining part of the morning was devoted to a general discussion about the seminar's outcomes. Due to the seminar and the multi-disciplinary interactions, it became clear that protection of simple genomic sequences is not enough for a full-privacy preserving approach. The organizers, together with the participants, agreed that this problem should be addressed in a sequel Dagstuhl-seminar. Hence, they set up a future work agenda in order to organize again such a fruitful gathering.

We thank Schloss Dagstuhl for the professional and inspiring atmosphere it provides. Such an intense research seminar is possible because Dagstuhl so perfectly meets all researchers' needs.

4.59 Algorithms for Optimization Problems in Planar Graphs

Organizers: Glencora Borradaile, Philip Klein, Dániel Marx, and Claire Mathieu

Seminar No. 13421

Date: October 13–18, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.36

© Creative Commons BY 3.0 Unported license

© Glencora Borradaile, Philip Klein, Dániel Marx, and Claire Mathieu



Participants: MohammadHossein Bateni, Ivona Bezakova, Therese Biedl, Glencora Borradaile, Sergio Cabello, Erin Moriarty Wolf Chambers, Eric Colin de Verdière, Sabine Cornelsen, Arnaud de Mesmay, Frederic Dorn, Alina Ene, Jeff Erickson, Jittat Fakcharoenphol, Kyle Jordan Fox, Petr A. Golovach, Michelangelo Grigni, MohammadTaghi Hajiaghayi, Marcin Kaminski, Philip N. Klein, Yusuke Kobayashi, Nitish Korula, Daniel Lokshtanov, Dániel Marx, Claire Mathieu, Tamara Mchedlidze, Matthias Mnich, Shay Mozes, Matthias Müller-Hannemann, Amir Nayyeri, Rolf Niedermeier, Yahav Nussbaum, Marcin Pilipczuk, Michał Pilipczuk, Peter Rossmanith, Ignaz Rutter, Saket Saurabh, Anastasios Sidiropoulos, Erik Jan van Leeuwen, Oren Weimann, Christian Wulff-Nilsen

Planar graphs, and more generally graphs embedded on surfaces, arise in applications such as road map navigation and logistics, computational topology, graph drawing, and image processing. There has recently been a growing interest in addressing combinatorial optimization problems using algorithms that exploit embeddings on surfaces to achieve provably higher-quality output or provably faster running times. New algorithmic techniques have been discovered that yield dramatic improvements over previously known results. In addition, results have been generalized to apply to other families of graphs: excluded-minor, bounded-genus and bounded-treewidth graphs.

This Dagstuhl seminar brought together researchers who have been working in these areas to present recent research results, consolidate and share understanding of the emerging basic techniques, and collaborate to move past the current barriers.

Polynomial-time solvable problems: There is a long tradition of finding fast algorithms for poly-time problems in planar graphs. In 1956, the first paper on maximum st -flow addressed the case where the network is planar (and s and t are adjacent). In 1976, a linear-time algorithm was given for minimum spanning trees in planar graphs. In 1979, the paper introducing generalized nested dissection gave a fast algorithm for shortest paths in planar graphs with positive and negative lengths. The past couple of decades has witnessed the discovery of fast algorithms for a wide range of polynomial-time problems in planar graphs: variants of max flow, multicommodity flow, variants of shortest paths, Gomory-Hu cut trees, global min-cut, girth, matching, and min-cost flow. It seems, however, there is a long way yet to go; for many promising problems, no planarity-exploiting algorithm is known or there is reason to believe faster algorithms can be obtained.

Approximation schemes: Research on polynomial-time approximation schemes (PTAS) for optimization problems in planar graphs goes back to the pioneering work of Lipton and Tarjan (1977) and Baker (1983), who introduced linear-time algorithms for certain problems in which the constraints were quite local, e.g. maximum-weight independent set and minimum-weight dominating set. For many years, little progress was made on problems with non-local constraints. In the mid-nineties, polynomial-time approximation schemes were developed for the traveling-salesman problem (TSP) in planar graphs, but in these the degree of the polynomial running time depended on the desired accuracy. A decade later, a linear-time approximation scheme was found for TSP. Shortly afterwards, the first polynomial time approximation schemes were found for problems, e.g. Steiner tree, in which the solution was much smaller than the input graph. Since then approximation schemes have been found for several other problems in planar graphs, such as two-connected spanning subgraph, Steiner forest, survivable network design, k -terminal cut, and k -center. Important new techniques have emerged, but we still lack fast approximation schemes for many important problems (e.g. facility location). The area of approximation schemes for planar graphs is ripe for further exploration.

Fixed-parameter tractable algorithms: Another way to cope with computational intractability of some planar graph problems is through the lens of fixed-parameter tractability. The theory of bidimensionality and algorithms exploiting tree decompositions of planar graphs give a general methodology of dealing with planar problems. One way to obtain fixed-parameter tractability results is to show that there is a polynomial-time preprocessing algorithm that creates a “problem kernel” by reducing the size of the instance such that it is bounded by a function of the parameter k . Research

on kernelization for planar graph problems has been a very active topic recently, culminating in a meta-theorem that gives problem kernels for a wide range of problems (2009).

The scientific program of the seminar consisted of 24 talks. Five of these talks were longer (60-90 minute) tutorials overviewing the three main areas of the seminar: Jeff Erickson (“Flows in planar and surface graphs”) and Christian Wulff-Nilsen (“Separators in planar graphs with applications”) covered polynomial-time algorithms; Philip Klein (“Some techniques for approximation schemes on planar graphs”) covered approximation schemes; and

Dániel Marx (“The square-root phenomenon in planar graphs”) and Daniel Lokshtanov (“Kernels for planar graph problems”) covered fixed-parameter tractability. One of the main goals of the seminar was to encourage collaboration between the three communities, and these well-received tutorials were very helpful by introducing the basics of each of these topics. The rest of the talks were 25-minute presentations on recent research of the participants.

The time between lunch and the afternoon coffee break was left open for individual discussions and collaborations in small groups. Two open-problem sessions were organized (on Monday evening and Wednesday evening).

4.60 Nominal Computation Theory

Organizers: Mikołaj Bojańczyk, Bartek Klin, Alexander Kurz, and Andrew M. Pitts
Seminar No. 13422

Date: October 13–16, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.58

© Creative Commons BY 3.0 Unported license

© Mikołaj Bojańczyk, Bartek Klin, Alexander Kurz, and Andrew M. Pitts



Participants: Andreas R. Blass, Mikołaj Bojańczyk, James Cheney, Vincenzo Ciancia, Thomas Colcombet, Roy L. Crole, Anuj Dawar, Jamie Gabbay, Fabio Gadducci, Tomasz Gogacz, Bartek Klin, Alexander Kurz, Sławomir Lasota, Ranko Lazic, Steffen Lösch, Justus Matthesen, Stefan Milius, Ugo Montanari, Andrzej Murawski, Joanna Ochremiak, Daniela Petrisan, Andrew M. Pitts, Luc Segoufin, Alexandra Silva, Ian Stark, Tomoyuki Suzuki, Szymon Toruńczyk, Emilio Tuosto, Nikos Tzevelekos, Christian Urban

The short Dagstuhl seminar 13422 “Nominal Computation Theory” took place from October 13th to 16th, 2013. The topic of the seminar was the theory of nominal sets and their applications to Computer Science. The seminar arose from a recent exciting and unexpected confluence of the following three distinct research directions.

- The research in automata theory on automata over infinite alphabets with applications to querying XML and databases.
- The research in program semantics on nominal sets, with many applications to the syntax and semantics of programming language constructs that involve binding, or localising names.
- The research in concurrency on nominal calculi (π -calculus, etc) with applications to the automatic verification of process specifications.

In each of these three topics, an important role is played by name (or atom) symmetries and permutations, albeit for a priori different reasons. In the first case they arise from the way automata use registers to store letters, in the second case they are used to define the notion freshness, in the third case they are needed to minimize automata. In all three cases there is a connection with mathematical model theory, which is aimed at studying classes of mathematical structures definable by logical theories. The permutations allowed on atomic names can be usefully understood as automorphisms of a relational structure on those names. Model-theoretic notions such as homogeneity, algebraic closure and oligomorphic groups turn out very useful in describing those relational structures on atoms that yield meaningful theories of nominal sets.

The aim of the seminar was to profit from the excitement created by the confluence described above and to explore new directions with a new mix of research communities from computer science and mathematical logic. The main topics of interest

included: automata and complexity theory in nominal sets, verification of nominal automata, symmetry in domain theory, and nominal programming.

The seminar was attended by 30 participants from 8 countries; 20 of them gave presentations, whose abstracts are included in this document. Four of the presentations (A. Pitts, M. Bojańczyk, B. Klin and N. Tzevelekos) were extended tutorials that presented various points of view on the background topics of the meeting. Other speakers presented the current state of the art in the field, with topic varying from mathematical insights into the nature of nominal sets (A. Blass, D. Petrisan), to applications in automata theory (V. Ciancia, T. Colcombet, S. Lasota, T. Suzuki) and computation theory (A. Dawar, S. Toruńczyk), semantics and domain theory (R. Crole, J. Gabbay, S. Loesch, A. Murawski), process calculi (U. Montanari), Petri nets (R. Lazic), logic programming (J. Cheney) and theorem proving (C. Urban).



Fig. 4.7
Gabriele Eickhoff – Grüner-Ausschnitt. Part of the Dagstuhl art collection and donated by: Uwe Assmann, Wolfgang Eisenbrand, Barbara Ryder, Lilli Schimpf, Franz Schimpf, THS Wirtschaftsprüfung, Ute Vollmar, Roland Vollmar, Margret Wilhelm, and participants in Dagstuhl Seminar 00381.

4.61 Real-World Visual Computing

Organizers: Oliver Grau, Marcus A. Magnor, Olga Sorkine-Hornung, and Christian Theobalt
Seminar No. 13431

Date: October 20–25, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.72

© Creative Commons BY 3.0 Unported license

© Marcus Magnor, Oliver Grau, Olga Sorkine-Hornung, and Christian Theobalt



Participants: Philippe Bekaert, Tamy Boubekeur, Edmond Boyer, Gabriel Brostow, Darren Cosker, Carsten Dachsbacher, Robert Dawes, Jean-Michel Dischler, Bernd Eberhardt, Peter Eisert, Paolo Favaro, Dieter W. Fellner, Jan-Michael Frahm, Martin Fuchs, Bastian Goldlücke, Oliver Grau, Volker Helzle, Anna Hilsmann, Adrian Hilton, Martin Jagersand, Jan Kautz, Oliver Klehm, Felix Klose, Andreas Kolb, Hendrik P. A. Lensch, Christian Lipski, Yebin Liu, Céline Loscos, Marcus A. Magnor, Shohei Nobuhara, Sylvain Paris, Fabrizio Pece, Kari Pulli, Bodo Rosenhahn, Holly E. Rushmeier, Hans-Peter Seidel, Alla Sheffer, Philipp Slusallek, Olga Sorkine-Hornung, Alexander Sorkine-Hornung, Ayellet Tal, Christian Theobalt, Stefanie Wuhrer, Ruigang Yang, Remo Ziegler

Dagstuhl seminar 13431 “Real-World Visual Computing” took place October 20–25, 2013. 45 researchers from North America, Asia, and Europe discussed the state-of-the-art, contemporary challenges, and promising future research directions in the areas of acquiring, modeling, editing, and rendering of complex natural scenes and events. The seminar was encompassed an introductory and a closing session, 9 scientific presentation sessions, two book organizational sessions as well as one special session on the Uncanny Valley problem. The seminar brought together junior and senior researchers from computer graphics, computer vision, 3D animation and visual special effects, both from academia and industry, to address the challenges in real-world visual computing. Participants included international experts from Kyoto University, Tsinghua University, University of British Columbia, University of Alberta, University of North Carolina, University of Kentucky, Yale University, Technion – Haifa, Filmakademie Baden-Wuerttemberg, Hochschule der Medien Stuttgart, Disney Research Zurich, BBC Research & Development, Intel Visual Computing Institute, Nvidia Corp., Adobe Systems Inc., metaio GmbH as well as many more research institutions and high-tech companies.

Motivating this seminar was the observation that digital models of real-world entities have become an essential part of innumerable computer graphics applications today. With ever-increasing graphics hardware and software capabilities, however, so does the demand for more and more realistically detailed models. Because the traditional, labor-intensive process of digital model creation by hand threatens to stall further progress in computer graphics, conventional manual modeling approaches are giving way to new approaches that aim at capturing complex digital models directly from the real world. The seminar picked up on recent trends in acquisition hardware for real-world events (e.g., Microsoft Kinect, Lytro light field camera, swarm of smartphone sensors, ...) as well as in visual computing applications (e.g.,

3D movies, Streetview, digital mock-ups, free-viewpoint systems, ...). It brought together experts from academia and industry working on contemporary challenges in image-based techniques, geometry modeling, computational photography and videography, BRDF acquisition, 3D reconstruction, 3D video, motion and performance capture etc. Collectively we fathomed the full potential of real world-based modeling approaches in computer graphics and visual computing.

Over the past decade, computer graphics has evolved into a mainstream area of computer science. Its economic impact and social pervasion range from professional training simulators to interactive entertainment, from movie production to trauma therapy, from geographic information systems to Google Earth. As a result, expectations on computer graphics performance are rising continuously. In fact, thanks to the progress in graphics hardware as well as rendering algorithms, visual realism is today within easy reach of off-the-shelf PCs, laptops, and even handheld devices. With rapidly advancing rendering capabilities, however, in many application areas of computer graphics the modeling process is becoming the limiting factor. Higher visual realism can be achieved only from more detailed and accurate scene descriptions. So far, however, digitally modeling 3D geometry and object texture, surface reflectance characteristics and scene illumination, motion and emotion is a labor-intensive, tedious process performed by highly trained animation specialists. The cost of conventionally creating models of sufficient complexity to engage the full potential of modern GPUs increasingly threatens to stall progress in computer graphics.

To overcome this bottleneck, an increasing number of researchers and engineers worldwide is investigating alternative approaches to create realistic digital models directly from real-world objects and scenes: Google and Microsoft already digitize entire cities using panorama video footage, 3D scanners, and GPS; RTT AG in Munich creates highly realistic digital mock-ups for

the car industry from CAD data and measured surface reflectance characteristics of car paint; at Disney Research, algorithms are being developed to create stereoscopic movies from monocular input; and BBC R&D has developed various 3D sports visualization methods based on analyzing live-broadcast footage.

In recent years, special effects in movies and computer games have reached a new level of complexity. In their aim to construct convincing virtual environments or even virtual actors, VFX companies are more and more relying on techniques to capture models from the real world. Currently available reconstruction tools, however, are still in their infancy. A lot of time is still spent on manual post-processing and modeling. The research community has responded to this trend by investigating new image- and video-based scene reconstruction approaches that can capture richer and more complex models. An example are performance capture methods that estimate more detailed shape and motion models of dynamics scenes than do commercially available systems. Similar methods for reconstruction of entire sets are also currently investigated, but many algorithmic problems remain to be solved.

The trend towards model capture from real world-examples is additionally bolstered by new sensor technologies becoming available at mass-market prices, such as Microsoft's Kinect, time-of-flight 2D depth imagers, or Lytro's Light Field camera. Also the pervasiveness of smartphones containing a camera, GPS, and orientation sensors allows for developing new capturing paradigms of real-world events based on a swarm of networked handheld devices. With the advent of these exciting novel acquisition technologies, investigating how to best integrate these new capture modalities into the computer graphics modeling pipeline, or how to alter traditional modeling to make optimal use of the new capture approaches, has become a top priority in visual computing research.

Researchers working on all of these problems from different direction came together at the seminar to share their experiences

and discuss the scientific challenges. Questions discussed were both theoretical and practical in nature. The seminar participants discussed the contemporary scientific challenges in modeling from the real world and determined which research avenues are likely to be the most promising and interesting ones to pursue over the course of the next years.

Among the questions and issues that have been addressed in the seminar are how to capitalize on new sensors for capture (computational cameras, light field cameras, Time-of-flight sensors, Kinect, omni-visual systems, ...), how to capture different object/scene aspects (geometry, reflectance, texture, material/fabric, illumination, dynamics, ...), how to digitally represent real-world objects/scenes (meshes, voxels, image-based, animation data, ...), how to convincingly & intuitively manipulate real-world models (relighting, motion editing, constrained manipulation, sketch-based, example-based, ...), how to realistically compose/augment and real-time render new scenes (F/X, movie post-production, games, perceptual issues, ...), and how to exploit the immense amount of community image and video data that are captured with handheld devices to build detailed models of the world (buildings, acting/dancing performances, sports events, fish tanks, ...). Also, the challenges arising from the large data sets of real-world models have been addressed. A special session on perceptual issues in animation (the Uncanny Valley problem) set out to identify the most important factors that are still unrealistic in computer animation. As the single most important area, facial animation was identified and some research directions for improvements were discussed.

The overall goal of the seminar to form a lasting, interdisciplinary research community was impressively underlined by the willingness of many seminar participants to work together on an edited book on the topic of the seminar. The book will be published with CRC Press. Completion of the manuscript is scheduled for August 2014.

4.62 Evaluation Methodologies in Information Retrieval

Organizers: Maristella Agosti, Norbert Fuhr, Elaine Toms, and Pertti Vakkari
Seminar No. 13441

Date: October 27 to November 1, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.10.92

© Creative Commons BY 3.0 Unported license

© Maristella Agosti, Norbert Fuhr, Elaine Toms, and Pertti Vakkari



Participants: Maristella Agosti, Omar Alonso, Leif Azzopardi, Nicholas J. Belkin, Ann Blandford, Charles Clarke, Maarten de Rijke, Arjen P. de Vries, Floriana Esposito, Nicola Ferro, Luanne Freund, Norbert Fuhr, Jacek Gwizdzka, Matthias Hagen, Preben Hansen, Jiyin He, Kalervo Järvelin, Hideo Joho, Jaap Kamps, Noriko Kando, Evangelos Kanoulas, Diane Kelly, Birger Larsen, Dirk Lewandowski, Christina Lioma, Thomas Mandl, Peter Mutschke, Ragnar Nordlie, Heather O'Brien, Doug Oard, Vivien Petras, Martin Potthast, Soo Young Rieh, Gianmaria Silvello, Paul Thomas, Elaine Toms, Vu Tran, Pertti Vakkari, C .J. Keith van Rijsbergen, Robert Villa, Max L. Wilson, Christa Womser-Hacker

Evaluation of information retrieval (IR) systems has a long tradition. However, the test-collection based evaluation paradigm is of limited value for assessing today's IR applications, since it fails to address major aspects of the IR process. Thus there is a need for new evaluation approaches, which was the focus of this seminar.

Before the event, each participant was asked to identify one to five crucial issues in IR evaluation methodology. Pertti Vakkari presented a summary of this homework, pointing out that there are five major themes deemed relevant by the participants: 1) Evaluation frameworks, 2) Whole session evaluation and evaluation over sessions, 3) Evaluation criteria: from relevance to utility, 4) User modeling, and 5) Methodology and metrics.

Based on the evaluation model proposed in Saracevic & Covi [1], the seminar started with four introductory talks covering major areas of IR evaluation: Nick Belkin gave a survey over "Framework(s) for Evaluation (of whole-session) IR", addressing the system components to be evaluated and the context to be considered. In his presentation "Modeling User Behavior for Information Retrieval Evaluation", Charlie Clarke described efforts for improving system-oriented evaluation through explicit models of user behavior. Kal Järvelin talked about "Criteria in User-oriented Information Retrieval Evaluation", characterizing them as different types of experimental variables and distinguishing between output- and (task-)outcome related criteria. "Evaluation Measures in Information Retrieval" by Norbert Fuhr outlined the steps necessary for defining a new metric and the underlying assumptions, calling for empiric foundation and theoretic soundness. Diane Kelly presented problematic issues related to "Methodology in IR Evaluation", such as the relationship between observation variables and criteria, the design of questionnaires, the difference between explanatory and predictive research and the appropriateness of statistical methods when dealing with big data. The round of introductory talks was concluded with

Maristella Agosti's presentation "Future in Information Retrieval Evaluation", where she summarized challenges identified in three recent workshops in this area.

For the rest of the week, the participants then formed working groups described in the following.

"From Searching to Learning" focused on the learning as search outcome and the need for systems supporting this process. Learning may occur at two different levels, namely the content level and the search competence level. There is a need for understanding of the learning process, its relationship to the searcher's work task, the role of the system, and the development of appropriate evaluation methods. Approaches may address different aspects of the problem, such as the system, the interaction, the content, the user and the process. For evaluation, the framework from Ingwersen and Jarvelin [2] suggests criteria and measures at the levels of information retrieval, information seeking, the work task and the social-organizational and culture level.

"Social Media" allow users to create and share content, with a strong focus on personal connections. While web search engines are still the primary starting point for many information seeking activities, information access activities are shifting to more personalized services taking into account social data. This trend leads to new IR-related research issues, such as e.g. utility, privacy, the influence of diverse cultural backgrounds, data quality, authority, content ownership, and social recommendations. Traditional assumptions about information seeking will have to be revised, especially since social media may play a role in a broad range of information spaces, ranging from everyday life and popular culture to professional environments like journalism and research literature.

"Graph Search and Beyond" starts from the observation that an increasing amount of information on the Web is structured in terms of entities and relationships, thus forming a graph, which,

in turn allows for answering more complex information needs. For handling these, search engines should support incremental structured query input and dynamic structured result set exploration. Thus, in contrast to the classical search engine result page, graph search calls for an incremental query exploration page, where entries represent the answers themselves (in the form of entities, relationships and sub-graphs). The new possibilities of querying and result presentation call for the development of adequate evaluation methods

“Reliability and Validity” is considered as the most central issue in IR evaluation, especially in the current situation where there is increasing discussion in the research community about reproducibility and generalizability of experimental results. Thus, this working group decided to start the preparation of a book on best practices in IR evaluation, which will cover the following aspects: Basic definitions and concepts, reliability and validity in experimentation, reporting out experiments, failure analysis, definition of new measures and methods, guidelines for reviewing experimental papers.

“Domain Specific Information Retrieval” in specific domains like e.g. in cultural heritage, patents and medical collections is not only characterized through the specifics of the content, but also through the typical context(s) in which this information is accessed and used, which requires specific functionalities that go beyond the simple search interaction. Also, context often plays an important role, and thus should be considered by the information system. However, there is a lack of appropriate evaluation methods for considering contexts and new functions.

“Task-Based IR” typically refers to research focusing on the task or goal motivating a person to invoke an IR system, thus calling for systems being able to recognize the nature of the task and to support the accompanying search process. As task types, we can distinguish between motivating tasks, seeking tasks, and search tasks. Task-based IR approaches should be able to model

people as well as the process, and to distinguish between the (task-related) outcome and the (system) output.

“Searching for Fun” refers to the interaction with an information system without a specific search objective, like e.g. online window shopping, watching pictures or movies, or reading online. This type of activity requires different evaluation criteria, e.g. with regard to stopping behavior, dwell time and novelty. Also, it is important to distinguish between system criteria and user criteria, where the latter may be subdivided into process criteria and outcome criteria. A major problem in this area is the design of user studies, especially since the starting points (e.g. casual or leisure needs) are difficult to create under experimental conditions. A number of further issues was also identified.

The working group “The Significance of Search, Support for Complex Tasks, and Searcher-aware Information Access Systems” addressed three loosely related challenges. The first topic addresses the definition of IR in the light of the dramatic changes during the last two decades, and the limited impact of our research. The second topic is the development of tools supporting more complex tasks, and their evaluation. Finally, information systems should become more informed about the searcher and the progress in user’s task.

“Interaction, Measures and Models” discussed the need for a common framework for user interaction models and associated evaluation measures, especially as a means for achieving a higher degree of reliability in interactive IR experiments. This would allow for evaluating the effect of the interaction and the interface on performance. A possible solution could consist of three components, namely an interaction model, a gain model and a cost model.

Finally, many of the attendees were planning to continue to collaborate on the topics addressed during the seminar since the fruitful discussions were a useful base for future cooperation.

References

- 1 Tefko Saracevic, Lisa Covi (2000). Challenges for digital library evaluation. In D. H. Kraft (Ed.), *Knowledge Innovations: Celebrating Our Heritage, Designing Our Future. Proceedings of the 63rd Annual Meeting of the American Society for Information Science. Washington, D.C.: American Society for Information Science.* pp. 341–350.
- 2 Peter Ingwersen, Kalervo Järvelin (2005). *The Turn: Integration of Information Seeking and Retrieval. In Context. Dordrecht, NL: Springer.* ISBN 1-4020-3850-X

4.63 Computational Audio Analysis

Organizers: Meinard Müller, Shrikanth S. Narayanan, and Björn Schuller
Seminar No. 13451

Date: November 3–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.1

© Creative Commons BY 3.0 Unported license

© Meinard Müller, Shrikanth S. Narayanan, and Björn Schuller



Participants: Xavier Anguera, Jon Barker, Stephan Baumann, Murtaza Bulut, Carlos Busso, Nick Campbell, Laurence Devillers, Jonathan Driedger, Bernd Edler, Anna Esposito, Sebastian Ewert, Cédric Févotte, Jort Gemmeke, Franz Graf, Martin Heckmann, Dorothea Kolossa, Gernot Kubin, Frank Kurth, Sungbok Lee, Florian Metze, Roger K. Moore, Emily Mower Provost, Meinard Müller, Shrikanth S. Narayanan, Nobutaka Ono, Bryan Pardo, Alexandros Potamianos, Bhiksha Raj, Gaël Richard, Mark Sandler, Björn Schuller, Joan Serrà, Rita Singh, Paris Smaragdis, Stefano Squartini, Shiva Sundaram, Khiat Truong, Christian Uhle, Emmanuel Vincent, Tuomas Virtanen

With the rapid growth and omnipresence of digitized multimedia data, the processing, analysis, and understanding of such data by means of automated methods has become a central issue in computer science and associated areas of research. As for the acoustic domain, audio analysis has traditionally been focused on data related to speech with the goal to recognize and transcribe the spoken words. In this seminar, we considered current and future audio analysis tasks that go beyond the classical speech recognition scenario. For example, we looked at the computational analysis of speech with regard to the speakers' traits (e. g., gender, age, height, cultural and social background), physical conditions (e. g., sleepiness, alcohol intoxication, health state), or emotion-related and affective states (e. g., stress, interest, confidence, frustration). So, rather than recognizing *what* is being said, the goal is to find out *how* and *by whom* it is being said. Besides speech, there is a rich variety of sounds such as music recordings, animal sounds, environmental sounds, and combinations thereof. Just as for the speech domain, we discussed how to decompose and classify the content of complex sound mixtures with the objective to infer semantically meaningful information.

When dealing with specific audio domains such as speech or music, it is crucial to properly understand and apply the appropriate domain-specific properties, be they acoustic, linguistic, or musical. Furthermore, data-driven learning techniques that exploit the availability of carefully annotated audio material have successfully been used for recognition and classification tasks. In this seminar, we discussed issues that arise when dealing with rather vague categories as in emotion recognition or when considering general audio sources such as environmental sounds. In such scenarios, model assumptions are often violated, or it becomes impossible to define explicit representations or models. Furthermore, for non-standard audio material, annotated datasets are hardly available. Also, data-driven methods that are used in

speech recognition are (often) not directly applicable in this context; instead semi-supervised or unsupervised learning techniques can be a promising approach to remedy these issues. Another central topic of this seminar was concerned with the problem of source separation. In the real world, acoustic data is very complex typically consisting of a superposition of overlapping speech, music, and general sound sources. Therefore, efficient source separation techniques are required that allow for splitting up, re-synthesizing, analyzing, and classifying the individual sources—a problem that, for general audio signals, is yet not well understood.

In this executive summary, we give a short overview of the main topics addressed in this seminar. We start by briefly describing the background of the participants and the overall organization. We then give an overview of the presentations of the participants and the results obtained from the different working groups. Finally, we reflect on the most important aspects of this seminar and conclude with future implications.

■ Participants, Interaction, Activities

In our seminar, we had 41 participants, who came from various countries around the world including North America (10 participants), Japan (1 participant), and Europe (Austria, Belgium, Finland, France, Germany, Greece, Italy, Netherlands, Spain, United Kingdom). Most of the participants came to Dagstuhl for the first time and expressed enthusiasm about the open and retreat-like atmosphere. Besides its international character, the seminar was also highly interdisciplinary. While most of the participating researchers are working in the fields of signal processing and machine learning, we have had participants with a background in cognition, human computer interaction, music, linguistics, and other fields. This made the seminar very special in having many cross-disciplinary intersections and pro-

voicing discussions as well as numerous social activities including common music making.

■ Overall Organization and Schedule

Dagstuhl seminars are known for having a high degree of flexibility and interactivity, which allow participants to discuss ideas and to raise questions rather than to present research results. Following this tradition, we fixed the schedule during the seminar asking for spontaneous contributions with future-oriented content, thus avoiding a conference-like atmosphere, where the focus is on past research achievements. The first two days were used to let people introduce themselves, present scientific problems they are particularly interested in and express their expectations and wishes for the seminar. In addition, we have had six initial stimulus talks, where specific participants were asked to address some burning questions on speech, music, and sound processing from a more meta point of view. Rather than being usual presentations, most of these stimulus talks seamlessly moved towards an open discussion of the plenum. Based on this input, the second day concluded with a brainstorming session, where we identified central topics covering the participants' interests and discussed the schedule and format of the subsequent days. To discuss these topics, we split up into five groups, each group discussing one of the topics in greater depth in parallel sessions on Wednesday morning. The results and conclusions of these group meetings were then presented to the plenum on Thursday morning, which resulted in vivid discussions. Continuing the previous activities, further parallel group meetings were held on Thursday afternoon, the results of which being presented on Friday morning. Finally, asking each participant to give a short (written) statement of what he or she understands by the seminar's overall topic "Computational Audio Analysis," we had a very entertaining and stimulating session by going through and discussing all these statements one by one. In summary, having a mixture of different presentation styles and group meetings gave all participants the opportunity for presenting and discussing their ideas, while avoiding a monotonous conference-like atmosphere.

■ Main Topics

We discussed various topics that addressed the challenges when dealing with mixtures of general and non-standard acoustic data. A particular focus was put on data representations and analysis techniques including audio signal processing, machine learning, and probabilistic models. After a joint brainstorming session, we agreed on discussing five central topics which fitted in the overall theme of the seminar and reflected the participants' interests. We now give a brief summary of these topics, which were addressed in the parallel group meetings and resulting panel discussions.

1. The "*Small Data*" group looked at audio analysis and classification scenarios where only few labeled examples or small amounts of (training) data are available. In such scenarios, machine learning techniques that depend on large amounts of (training) data ("*Big Data*") are not applicable. Various strategies including model-based as well as semi- and unsupervised approaches were discussed.
2. The "*Source Separation*" group addressed the task of decomposing a given sound mixture into elementary sources, which is not only a fundamental problem in audio processing, but also constitutes an intellectual and interdisciplinary challenge. Besides questioning the way the source separation problem is often posed, the need of concrete application

scenarios as well as the objective of suitable evaluation metrics were discussed.

3. The "*Interaction and Affect*" group discussed the question on how to generate and interpret signals that express interactions between different agents. One main conclusion was that one requires more flexible models that better adapts to the temporal and situational context as well as to the agents' roles, behaviors and traits.
4. The "*Knowledge Representation*" group addressed the issue of how knowledge can be used to define and derive sound units that can be used as elementary building blocks for a wide range of applications. Based on deep neural network techniques, the group discussed how database information and other meta-data can be better exploited and integrated using feed-forward as well as recurrent architectures.
5. The "*Unsupervised Learning*" group looked at the problem on how to learn the structure of data without reference to external objectives. Besides issues on learning meaningful elementary units, the need of considering hierarchies of abstractions and multi-layer characterizations was discussed.

Besides an extensive discussion of these five main topics, we have had many further contributions and smaller discussions on issues that concern natural human machine communication, human centered audio processing, computational paralinguistics, sound processing in everyday environments, acoustic monitoring, informed source separation, and audio structure analysis.

■ Conclusions

In our seminar, we addressed central issues on how to process audio material of various types and degrees of complexity. In view of the richness and multitude of acoustic data, one requires representations and machine learning techniques that allow for capturing and coupling various sources of information. Therefore, unsupervised and semi-supervised learning procedures are needed in scenarios where only very few examples and poor training resources are available. Also, source separation techniques are needed, which yield meaningful audio decomposition results even when having only limited knowledge on the type of audio. Another central issue of this seminar was how to bring in the human into the audio processing pipeline. On the one hand, we discussed how we can learn from the way human process and perceive sounds. On the other hand, we addressed the issue on extracting human-related parameters such as affective and paralinguistic information from sound sources. These discussions showed that understanding and processing complex sound mixtures using computational tools poses many challenging research problems yet to be solved.

The Dagstuhl seminar gave us the opportunity for discussing such issues in an inspiring and retreat-like atmosphere. The generation of novel, technically oriented scientific contributions was not the focus of the seminar. Naturally, many of the contributions and discussions were on a rather abstract level, laying the foundations for future projects and collaborations. Thus, the main impact of the seminar is likely to take place in the medium to long term. Some more immediate results, such as plans to share research data and software, also arose from the discussions. As measurable outputs from the seminar, we expect to see several joint papers and applications for funding. Beside the scientific aspect, the social aspect of our seminar was just as important. We had an interdisciplinary, international, and very interactive group of researchers, consisting of leaders and future leaders in our field. Most of our participants visited Dagstuhl for the first time and enthusiastically praised the open and inspiring atmosphere. The group dynamics were excellent

with many personal exchanges and common activities. Some scientists mentioned their appreciation of having the opportunity for prolonged discussions with researchers from neighboring research fields—something which is often impossible during conference-like events.

In conclusion, our expectations of the seminar were not only met but exceeded, in particular with respect to networking and community building. Last but not least, we heartily thank the Dagstuhl board for allowing us to organize this seminar, the Dagstuhl office for their great support in the organization process, and the entire Dagstuhl staff for their excellent services during the seminar.



Fig. 4.8

Sabine Eisenbrand – Sandbild 5. Part of the Dagstuhl art collection and donated by participants in Dagstuhl Seminar 13451.

4.64 Proxemics in Human-Computer Interaction

Organizers: Saul Greenberg, Kasper Hornbæk, Aaron Quigley, Harald Reiterer, Roman Rädle
Seminar No. 13452

Date: November 3–8, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.29

© Creative Commons BY 3.0 Unported license

© Saul Greenberg, Kasper Hornbæk, Aaron Quigley, Harald Reiterer, and Roman Rädle



Participants: Elisabeth André, Jakob E. Bardram, Susanne Böll, Sebastian Boring, Andreas Butz, Raimund Dachselt, Andreas Dippon, Jakub Dostal, Saul Greenberg, Kasper Hornbæk, Jörn Hurtienne, Petra Isenberg, Mikkel R. Jakobsen, Hans-Christian Jetter, Nicolai Marquardt, Fabrice Matulic, Florian Floyd Mueller, Max Nicosia, Kenton O'Hara, Thomas Pederson, Aaron Quigley, Roman Rädle, Harald Reiterer, Stacey D. Scott, Henrik Sorensen, Jürgen Steimle, Sophie Stellmach, Jo Vermeulen, Daniel Vogel

■ Introduction

Over time, people encounter different dimensions of proxemics in everyday life, such as in face-to-face communication while discussing ongoing work with colleagues, in an elevator with strangers as private space is suspended, or at home with their families. In disciplines like architecture and interior design, knowledge about proxemics has been used for decades to model use of space for face-to-face interactions, urban planning, and environmental design. In human-computer interaction (HCI) and human-robot interaction (HRI), the use of proxemics is fairly new, and both disciplines recently began employing proxemics and related theories and models (e.g., Hall's theory of proxemics in his book, "The Hidden Dimension" [2]) to design new interaction concepts that act on proxemics features. Several recent designs explore the use of human body position, orientation, and movement for implicit interaction with large displays, supporting collaboration, and to control and communicate with robots. This research is facilitated by the operationalization of proxemics for ubiquitous computing [16], toolkits to track proxemics [7–9], and new paradigms such as reality-based interaction (RBI) [4] or Blended Interaction [6] that take a fresh look at the role of the user's body and the environment in HCI. However, work on understanding how proxemics can be used for HCI (and HRI) has only just begun (e.g., Proxemic Interactions [1]).

■ Goals and Structure

In the seminar, we used Greenberg et al.'s dimensions on Proxemic Interactions [1] and Pedersen et al.'s Egocentric Interaction Paradigm [11] as starting points. These theories are based on findings regarding how humans perceive proxemics; therefore, they might be incomplete, particularly since human perception is much more subtle, gradual, and less discrete than

illustrated in Hall's reaction bubbles (proxemic zones [2]). In addition, these discrete zones cope with only the physical features (perception of interpersonal distance). Other features, such as psychological and psychophysical features, have not yet been considered in HCI. However, these features are perceptible by human sensors (olfaction, equilibrioception, and thermoception). Current theories neither give guidelines nor provide sufficient methods for "good" or "bad" designs for systems employing proxemics.

We thought the time was right for bringing researchers with different backgrounds and experiences together to map out the important questions that remain unanswered and to generate ideas for developing an agenda for future research on proxemics in HCI.

The structure of the seminar was based on the four pillars technology, application, vision, and theory that were equally exposed in seminar activities. The forum held 29 attendees with multidisciplinary backgrounds from research institutes in Canada, Denmark, England, Switzerland, Australia, France, Belgium, and Germany. We achieved productive and critical reflections and prospects on proxemics in HCI by letting experts from their respective fields work on a shared vision and theory. We selected the attendees to ensure an equal distribution of expertise across the four pillars.

The diversified program allowed attendees to introduce themselves and their work in brief presentations and offered one impulse keynote given by Saul Greenberg and Nicolai Marquardt. Greenberg and Marquardt coined the term Proxemic Interactions and decisively influenced the application of proxemics in HCI. We also provided ample time for discussions, breakout sessions, and creative work addressing concepts such as:

- Intelligibility of Proxemic Interactions
- Users' options to opt-in or opt-out
- The "dark side" of Proxemic Interactions
- The meaning of physical space

- How image schemas [3] can be used to brainstorm innovative proxemic systems
- Ad-Hoc proxemics
- Including everyday entities in proxemic systems

Throughout the entire seminar, attendees were encouraged to write down their questions, ideas, and comments. These materials were collected and posted to one of the four pillars on a pin board for the purpose of inspiring breakout groups and ad lib collaboration. The breakout session proposed by the group centered around open problems and challenges within proxemic interactions, which was then discussed in each session.

■ Technology

In recent years, emerging technology has changed the interaction between human and computer. For instance, smartphones and tablets have entered our daily life. More of such novel post-WIMP¹² technologies will be available in the foreseeable future and ultimately define how we interact in physical spaces. Interaction might take place across device boundaries on (multiple) public [15], large and private, mobile, and tangible displays [13]. It might involve collaboration of co-located users around interactive tabletops [7], in front of large vertical screens [5], or on rollout displays [14]. It might be based on non-traditional, post-WIMP interaction styles, such as pen-based [10], multi-touch, and tangible user interfaces. Or, it might provide new forms of functionality beyond the traditional WIMP model of applications by tracking users' spatial location and movements for navigation within large, digital information spaces [12]. Attendees discussed existing technologies that allow people-to-people, people-to-object, and object-to-object proxemics relations tracking, as well as improvements on tracking reliability using sensor fusion.

■ Application

Seminar attendees discussed the “light” and “dark” side of Proxemic Interactions. Until now, research has focused on the benefits of these interactions; however, they bear risks. We all can imagine how advertisement would change if it becomes possible to show customized ads according to our online shopping profiles while we are walking on public streets or in shopping malls. During the seminar, participants discussed what types of applications would best showcase the benefit of proxemics and avoid the risks that arise when systems are able to track and identify people. Part of this discussion included brainstorming opt-in or opt-out functions for proxemics-aware systems so that users can remain in control of these systems.

■ Vision

In its past, HCI has benefited from ambitious visions of future interaction such as Apple's Knowledge Navigator or Mark Weiser's “A day in the life of Sal” [16]. Although visions are not always helpful and can lead in wrong directions, we believe that a new overarching vision of future Proxemic Interactions can help inspire ongoing research and thrive in coming generations. This vision is intended to inform researchers, designers, and laymen alike. For researchers, a vision can serve to illustrate research goals, trigger new research directions, and create awareness for as yet un-reflected assumptions in our field. For designers, visions help to present concepts and technologies as a part of a believable scenario – and not only in the isolation of conference papers. Furthermore, visions serve to fascinate and inspire laymen, who prefer to learn about future technologies from narrations instead of purely technical publications. The seminar aimed at creating a unified vision of Proxemic Interactions based on the individual contributions and experiences of the seminar attendees. Current and past visions have been discussed in plenum and breakout groups.

■ Theory

In the light of the countless variants and dynamics of post-WIMP interaction, traditional collections of design guidelines or “golden rules” cannot provide enough guidance about “good” or “bad” designs. Instead, we need better theories and models of human cognition to be able to understand and classify designs of Proxemic Interactions and to predict their appropriateness. We wanted to understand how physical, psychological and psychophysical features collate and can be transferred into a coherent theory of proxemics in HCI and how to give guidelines or provide sufficient methods for “good” or “bad” designs. Therefore, we had to:

1. Better understand proxemics in HCI to develop such methods
2. Discuss the open question: to what extent can proxemics leverage or constrain human-computer interaction?

■ Conclusion

The Dagstuhl Seminar 13452 offered a fantastic forum for established researchers and practitioners at a comfortable place. We framed and discussed research questions and worked together on a unifying theory for Proxemics in Human-Computer Interaction. Applications for Proxemic Interactions were sketched out and critically reflected in the light of the “dark side” of proxemics. We also discussed how we can learn from related fields and how they can profit from proxemics in HCI.

The seminar can be seen as a good starting point to identify the role of Proxemics in Human-Computer Interaction. However, it still remains an open research area and its place in HCI needs to be better understood.

¹² WIMP stands as an acronym for Windows, Icons, Menus, Pointers

■ References

- 1 Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R. and Wang, M. 2011. Proxemic Interactions: The New Ubicomp? *interactions*. 18, January (2011), 42–50.
- 2 Hall, Edward, T. 1966. *The Hidden Dimension*. Doubleday.
- 3 Hurtienne, J., Israel, J.H. and Weber, K. 2008. Cooking up real world business applications combining physicality, digitality, and image schemas. *In Proc. of TEI'08*, New York, New York, USA, Feb. 2008, 239.
- 4 Jacob, R.J.K., Girouard, A., Hirshfield, L.M., Horn, M.S., Shaer, O., Solovey, E.T. and Zigelbaum, J. 2008. Reality-based interaction: a framework for post-WIMP interfaces. *In Proc. of CHI'08*, New York, NY, USA, 2008, 201–210.
- 5 Jakobsen, M. and Hornbæk, K. 2012. Proximity and physical navigation in collaborative work with a multi-touch wall-display. *In Proc. of CHI EA'12*, New York, NY, USA, 2012, 2519–2524.
- 6 Jetter, H.-C., Reiterer, H. and Geyer, F. 2013. Blended Interaction: understanding natural human-computer interaction in post-WIMP interactive spaces. *Personal and Ubiquitous Computing*. (Oct. 2013).
- 7 Klinkhammer, D., Nitsche, M., Specht, M. and Reiterer, H. 2011. Adaptive personal territories for co-located tabletop interaction in a museum setting. *In Proc. of ITS'11*, New York, NY, USA, Nov. 2011, 107.
- 8 Marquardt, N., Diaz-Marino, R., Boring, S. and Greenberg, S. 2011. The proximity toolkit: prototyping proxemic interactions in ubiquitous computing ecologies. *In Proc. of UIST '11*, New York, NY, USA, 2011, 315–326.
- 9 Marquardt, N., Hinckley, K. and Greenberg, S. 2012. Cross-device interaction via micro-mobility and f-formations. *In Proc. of UIST'12*, New York, NY, USA, Oct. 2012, 13.
- 10 Matulic, F. and Norrie, M.C. 2013. Pen and touch gestural environment for document editing on interactive tabletops. *In Proc. of ITS'13*, New York, NY, USA, Oct. 2013, 41–50.
- 11 Pederson, T. 2012. Proximity as Key Property in the Egocentric Interaction Paradigm. http://hci.uni-konstanz.de/proxemics/wp-content/uploads/2012/10/proxemics2012_Pedersen_et_al.pdf.
- 12 Rädle, R., Jetter, H.-C., Butscher, S. and Reiterer, H. 2013. The effect of egocentric body movements on users' navigation performance and spatial memory in zoomable user interfaces. *In Proc. of ITS'13*, New York, NY, USA, 2013, 23–32.
- 13 Spindler, M., Stellmach, S. and Dachselt, R. 2009. PaperLens: Advanced Magic Lens Interaction Above the Tabletop. *In Proc. of ITS'09*, New York, NY, USA, Nov. 2009, 69–76.
- 14 Steimle, J. and Olberding, S. 2012. When mobile phones expand into handheld tabletops. *In Proc. of CHI EA'12*, New York, NY, USA, May 2012, 271–280.
- 15 Vogel, D. and Balakrishnan, R. 2004. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. *In Proc. of UIST'04*, New York, NY, USA, 2004, 137–146.
- 16 Weiser, M. 1999. The computer for the 21st century. *ACM SIGMOBILE Mobile Computing and Communications Review*. 3, 3 (1999), 3–11.

4.65 Electronic Markets and Auctions

Organizers: Yishay Mansour, Benny Moldovanu, Noam Nisan, and Berthold Vöcking
Seminar No. 13461

Date: November 10–15, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.58

© Creative Commons BY 3.0 Unported license

© Yishay Mansour and Noam Nisan

Participants: Yossi Azar, Moshe Babaioff, Felix Brandt, Yang Cai, Shahar Dobzinski, Khaled Elbassioni, Michal Feldman, Amos Fiat, Felix Fischer, Yannai A. Gonczarowski, Nick Gravin, Tobias Harks, Paul Harrenstein, Avinatan Hassidim, Monika Henzinger, Martin Hoefer, Zhiyi Huang, Elias Koutsoupias, Annamaria Kovacs, Piotr Krysta, Stefano Leonardi, Yishay Mansour, Paul Milgrom, Rudolf Müller, Noam Nisan, Sigal Oren, Britta Peis, Heiko Röglin, Rahul Savani, Yaron Singer, Balasubramanian Sivan, Alexander Skopalik, Éva Tardos, Andreas Tönnis, Rob van Stee, Shai Vardi, Bernhard von Stengel, S. Matthew Weinberg, Christopher A. Wilkens



The relatively young field of Algorithmic Game Theory sets a goal of providing a computational understanding of game theory models. The research in the field has many focal points, including exploring the quality of equilibria, computation of equilibria, algorithmic mechanism design, as well as analyzing computer science related games and gaining an economics perspective for many important optimization problems.

While it is still too early for the evaluation the long term contribution of Algorithmic Game Theory to the field of Economics, in general, and to Game Theory in particular, we would like to highlight some successful contributions. The efficient computational aspects are a clear contribution, and this is also coupled with the understanding that sub-optimal solutions can have various degrees of sub-optimality. By using approximation algorithms approaches traditional in Theoretical Computer Science, the sub-optimality can be quantify in a very rigorous and clear way. The study of concrete convergence rates, rather than convergence in the limit, has proved to be highly successful here, as well. Finally, the extensive study of discrete models, especially combinatorial auctions, has been an area where computer science has made significant contributions.

The economic field of Mechanism Design asks how to design mechanisms that will implement some desired social choice function under rational behavior of the participants. This field is at the forefront of economics research, and its goal is to gain a better understanding of designing mechanisms that considers

the incentives of participant. This is in general viewed as part of market design, and micro-economics

One of the central areas of Algorithmic Game Theory is *Algorithmic Mechanism Design*. This field is relevant to designing distributed computer systems, suggested that mechanism design should also consider the algorithmic issues involved beyond the strategic ones commonly studied in economics. The seminar concentrated on Algorithmic Game Theory, with an emphasis on the sub-field of Algorithmic Mechanism Design.

The central application of Mechanism Design is the implementation of auctions and markets, and similarly the central application of algorithmic mechanism design is the implementation of complex computerized auctions and markets. As markets and auctions are increasingly implemented over computer networks, and as they are getting more sophisticated, much theoretical research has gone into the design of complex auctions and markets. Issues that need to be treated include computational ones, strategic ones, and communication ones. A central application is, so called, combinatorial auctions, which aim to concurrently sell many related items.

This seminar had researchers discussing basic research questions that lie behind the growing challenges in electronic markets and auctions. The seminar took a broad view of these challenges, focusing on foundational issues, taking a wide perspective, from the high-level issues of Algorithmic Game Theory through the Algorithmic Mechanism Design aspects, to basic challenges of electronic markets and auction.

4.66 Computational Models of Language Meaning in Context

Organizers: Hans Kamp, Alessandro Lenci, and James Pustejovsky

Seminar No. 13462

Date: November 10–15, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.79

© Creative Commons BY 3.0 Unported license

© Hans Kamp, Alessandro Lenci, and James Pustejovsky



Participants: Nicholas Asher, Marco Baroni, Peter A. Cariani, Stephen Clark, Ann Copestake, Ido Dagan, Katrin Erk, Stefan Evert, Patrick W. Hanks, Graeme Hirst, Jerry R. Hobbs, Hans Kamp, Lauri Karttunen, Alessandro Lenci, Sebastian Löbner, Louise McNally, Sebastian Padó, Massimo Poesio, James Pustejovsky, Anna Rumshisky, Hinrich Schütze, Mark Steedman, Suzanne Stevenson, Tim van de Cruys, Jan van Eijck, Dominic Widdows, Annie Zaenen, Alessandra Zarcone

The term distributional semantics qualifies a rich family of computational methods sharing the assumption that the statistical distribution of words in context plays a key role in characterizing their semantic behavior. Distributional semantic models, such as LSA, HAL, etc., represent the meaning of a content word in terms of a distributed vector recording its pattern of co-occurrences (sometimes, in specific syntactic relations) with other content words within a corpus. Different types of semantic tasks and phenomena are then modeled in terms of linear algebra operations on distributional vectors. Distributional semantic models provide a quantitative correlate to the notion of semantic similarity, and are able to address various lexical semantic tasks, such as synonym identification, semantic classification, selectional preference modeling, and so forth.

Distributional semantics has become increasingly popular in Natural Language Processing. Its attractiveness lies in the fact that distributional representations do not require manual supervision and reduce the the a priori stipulations in semantic modeling. Moreover, distributional models generally outperform other types of formal lexical representations, such as for instance semantic networks. Many researchers have also strongly argued for the psychological validity of distributional semantic representations. Corpus-derived measures of semantic similarity have been assessed in a variety of psychological tasks ranging from similarity judgments to simulations of semantic and associative priming, showing a high correlation with human behavioral data.

Despite its successes, no single distributional semantic model meets all requirements posed by formal semantics or linguistic theory, nor do they cater for all aspects of meaning that are important to philosophers or cognitive scientists. In fact, the distributional paradigm raises the question of the extent to which semantic properties can be reduced to combinatorial relations.

Many central aspects of natural language semantics are left out of the picture in distributional semantics, such as predication, compositionality, lexical inferences, quantification and anaphora, just to quote a few. A central question about distributional models is whether and how distributional vectors can also be used in the compositional construction of meaning for constituents larger than words, and ultimately for sentences or discourses – the traditional domains of denotation-based formal semantics. Being able to model key aspects of semantic composition and associated semantic entailments represents a crucial condition for distributional model to provide a more general model of meaning. Conversely, we may wonder whether distributional representations can help to model those aspects of meaning that notoriously challenge semantic compositionality, such as semantic context-sensitivity, polysemy, predicate coercion, pragmatically-induced reference and presupposition.

The main question is whether the current limits of distributional semantics represent contingent shortcomings of existing models – hopefully to be overcome by future research –, or instead they point to intrinsic inadequacies of vector-based representations to address key aspects of natural language semantics. To this end, there were five themes addressed by the participants:

1. The problems in conventional semantic models that distributional semantics claims to be able to solve;
2. The promise of distributional semantics linking to multimodal representations
3. The current limitations of distributional semantics theories to account for linguistic compositionality;
4. The absence of any robust first-order models of inference for distributional semantics;
5. The integration of distributional semantic principles and techniques into a broader dynamic model theoretic framework.



Fig. 4.9
Peter Amici – Drei Flugzeuge und zwölf Landschaften. Part of the Dagstuhl art collection and donated by: Jon Callas, Maxime Crochemore, Angelika Mueller v. Brochowski, Reinhard Wilhelm, Lenore D. Zuck, and participants in Dagstuhl Seminars 13382, 13511 and 14031.

4.67 Synchronous Programming

Organizers: Albert Benveniste, Stephen A. Edwards, Alain Girault, and Klaus Schneider
Seminar No. 13471

Date: November 17–22, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.117

© Creative Commons BY 3.0 Unported license

© Stephen A. Edwards, Alain Girault, and Klaus Schneider



Participants: Joaquin Aguado, Mihail Asavoae, Yu Bai, Fernando Barros, Sanjoy K. Baruah, Guillaume Baudart, Vagelis Bebelis, Gérard Berry, Nicolas Berthier, Timothy Bourke, David Broman, Benoît Caillaud, Albert Cohen, Willem-Paul de Roever, Robert de Simone, Gwenaël Delaval, Stephen A. Edwards, Peter A. Fritzon, Manuel Gesell, Dan R. Ghica, Alain Girault, Adrien Guatto, Nicolas Halbwegs, Jun Inoue, Xian Li, Gerald Lüttgen, Antoine Madet, Frédéric Mallet, Louis Mandel, Florence Maraninchi, Eleftherios Matsikoudis, Michael Mandler, Christian Motika, Valentin Perrelle, John Plaice, Dumitru Potop-Butucaru, Marc Pouzet, Pascal Raymond, Ingo Sander, Klaus Schneider, Alex Shafarenko, Chris Shaver, Wei-Tsun Sun, Reinhard von Hanxleden

■ Model-based Design of Embedded Systems

In general, the *development of embedded systems* is a challenging task: Concerning the hardware platforms, developers have to cope with tight resource constraints, *heterogeneous and application-specific hardware architectures*, virtual prototypes, and many other difficulties during the design phases. Concerning the software side, several concurrent tasks are executed on the available hardware, either with or without the help of special operating systems, sometimes statically or dynamically scheduled to the available *hardware platforms*, and sometimes tightly coupled with the hardware platforms themselves (implementing memory barriers etc). Finally, many non-functional aspects have to be considered as well like the energy consumption, the *reliability*, and most important the *prediction of the worst-case computation times*. As many embedded systems are *real-time systems*, it is not sufficient to perform the right computations; in addition, the results have to be available at the right point of time to achieve the desired functionality. Besides, the direct interaction with other systems that often have a continuous behavior requires to consider *cyber-physical systems*. Since many embedded systems are used in safety-critical applications, incorrect or delayed behaviors are unacceptable, so that *formal verification* is often applied. Since, moreover, the development costs have to be minimized, new design flows that allow the development of safe and flexible embedded systems are of high interest.

For these reasons, *model-based design flows* became popular where one starts with an abstract model of the embedded system. Many languages are discussed for such model-based approaches, but most of them are based on only a few models of computation. A *model of computation* thereby defines *which, when and why an action of the system takes place* taking into account the *timeliness*, the *causality*, and the *concurrency* of the computations. Classic

models of computation are *dataflow process networks*, where computations can take place as soon as sufficient input data is available, *synchronous systems*, which are triggered by clocks, *discrete-event based systems*, where each process is sensitive to the occurrence of a set of certain events, and *cyber-physical systems* whose behavior consists of discrete and continuous transitions (the latter are determined by differential equations).

It is not surprising that all models of computation have their advantages and disadvantages. For example, dataflow process networks can be naturally mapped to distributed systems and have a robust form of asynchronous concurrency provided that the nodes implement continuous functions (as required for Kahn networks). Synchronous systems are the perfect choice for implementing deterministic systems with predictable real-time behaviors on platforms having a local control (like clocks in digital hardware circuits). Discrete-event based systems are ideal for efficiently simulating systems, since the events directly trigger the next actions.

Many years of research were necessary to understand the above mentioned models of computation in depth to be able to develop corresponding programming languages, compilers and verification techniques. The synchronous programming community made substantial progress in this area: Today, the synchronous programming languages have precise formal semantics which are supported by efficient compiler techniques. Moreover, synchronous languages provide high-level descriptions of real-time embedded systems so that all relevant requirements for a model-based design flow are fulfilled. There are also graphical versions of these textual languages, notably Safe State Machines (developed from Argos and SyncCharts), and there are commercial versions like SCADE. The SCADE tool provides a code generator certified against DO 178-B, which makes it particularly attractive for the aircraft sector.

Quoting Benveniste et al.: *Today, synchronous languages have been established as a technology of choice for modeling, specifying, validating, and implementing real-time embedded applications. The paradigm of synchrony has emerged as an engineer-friendly design method based on mathematically sound tools* [Proceedings of the IEEE, January 2003].

■ Open Problems

Despite the incredible progress made in the past, even the combination of the classic synchronous languages Esterel, Lustre, and Signal is not yet fully understood. All these languages are based on the abstraction of physical time to a logical time, where each logical step of time may consist of finitely many executions of actions that are – at least in the programming model – executed in zero time. Such a logical step of the computation matches naturally with an interaction of a reactive system with its environment. However, looking at the details, one can observe that the semantics differ: for example, Lustre and Signal are not based on a single clock like Esterel, and while Esterel's and Lustre's semantics are operational and can therefore be defined by least fixpoints, Signal is rather declarative and requires a more complicated analysis before code generation.

Since different models of computation have different advantages and disadvantages, their *combination* becomes more and more important. This does also imply the *translation and communication between models of computations*. For example, so-called globally asynchronous, locally synchronous (GALS) systems have been developed, mixing both asynchronous and synchronous computations. For model-based designs starting from synchronous languages, special forms of synchronous systems have been defined in terms of the (weakly) endochronous systems. Intuitively, endochronous systems are synchronous systems that can determine from which input ports the values are expected for the next reaction step (and therefore they can derive the absence of other inputs, and they do not need the explicit knowledge of absence). For this reason, one can integrate endochronous systems in an asynchronous environment without destroying their synchronous behaviors.

Similar techniques are used for generating distributed systems from high-level descriptions (like synchronous programs) which lead, e.g., also to first approaches to multithreaded code generation from synchronous languages, which becomes more important due to the advent of multicore processors in embedded system design. More progress is needed and will likely be available in the near future in combining these different forms of discrete models of computations.

The combination of synchronous, endochronous, or asynchronous discrete systems with continuous behaviors to describe cyber-physical systems is still in its infancies. Of course, there are many languages for modeling, simulating, and even formally verifying these systems, but most of these languages lack of a formal semantics, and essentially none of them lends itself for a model-based design like synchronous languages. The generalization of the concepts of synchronous systems to polychronous systems, and even further to cyber-physical systems will be a challenge for future research.

■ Results of the Seminar

The major goal of the seminar was therefore to allow researchers and practitioners in the field of models of computation and model-based design to discuss their different approaches. Desired results are new combinations of these techniques to form new language concepts and design flows that are able to choose the best suited language for particular components and that allow engineers the sound integration of synchronous and asynchronous, discrete and continuous, or event- and time-triggered systems. Besides this, still more research is required for further developing compilation techniques for future manycore processors, and even to develop special processors like the PRET architectures to obtain better estimated time bounds for the execution of programs.

The seminar proposed here aims at addressing all of these questions, building on a strong and active community and expanding its scope into relevant related fields, by inviting researchers prominent in model-based design, embedded real-time systems, mixed system modeling, models of computation, and distributed systems. The seminar was held in the tradition of the Synchronous Programming (SYNCHRON) workshops that are used as the yearly meeting place for the community in this exciting field. The SYNCHRON workshops started in 1994 at Schloss Dagstuhl, and we were proud to celebrate the 20th edition of the workshop from November 18–22, 2013 again in Schloss Dagstuhl.

During its 20 years of existence, the workshop has significantly evolved: its scope has grown to expand to 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 have taken the field in new directions. We carefully selected the potential persons to be invited in that senior and junior researchers of the different branches mentioned above will participate the seminar.

This year, we had 44 participants where 23 came from France, 10 from Germany, 5 from the USA, 2 from Sweden, 2 from UK, one from Portugal and one even from Australia. The seminar had 33 presentations of about 45 minutes length with very active discussions¹³. The presentations can be clustered in typical research areas around synchronous languages like

- synchronous and asynchronous models of computation
- hybrid systems
- causality and other program analyses
- compilation techniques
- predictable software and hardware architectures

It was a pleasure to see that the synchronous programming community is still very active in these research fields and that even after 20 years of research, there are still more and more interesting and fruitful results to be discovered. The following sections contains short abstracts of the presentations of the seminar, and further documents were provided by many participants on the seminar's webpage.

¹³ See <http://www.dagstuhl.de/schedules/13471.pdf> for the schedule.

4.68 Global Measurement Framework

Organizers: Philip Eardley, Marco Mellia, Jörg Ott, Jürgen Schönwälder, and Henning Schulzrinne

Schulzrinne

Seminar No. 13472

Date: November 17–20, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.144

© Creative Commons BY 3.0 Unported license

© Philip Eardley, Marco Mellia, Jörg Ott, Jürgen Schönwälder, and Henning Schulzrinne



Participants: Saba Ahsan, Vaibhav Bajpai, Arthur W. Berger, Ernst Biersack, Trevor Burbridge, Fabian E. Bustamante, Pedro Casas, Benoit Claise, Sam Crawford, Philip Eardley, Daniel Karrenberg, Mirja Kühlewind, Abdelkader Lahmadi, Jukka Manner, Marco Mellia, Al Morton, Jörg Ott, Fabio Ricciato, Dario Rossi, Ramin Sadre, Jürgen Schönwälder, Henning Schulzrinne, Andrea Soppera, Anna Sperotto, Burkhard Stiller, Tivadar Szemethy, Brian Trammell

The Internet has a history of unexpected and often unpredictable behaviors due to manifold interactions of thousands of networks, and billions of components and devices and users. The resulting complexity requires measurements to understand how the network is performing, to observe how it is evolving, and to determine where failures or degradations occur. Especially with constantly evolving applications and their interaction paradigms, new phenomena occur and need to be factored into operations and management: one example is the substantial effort going into defining interfaces to assist peer-to-peer applications so that the amount of cross-ISP traffic is reduced. Measurements thus form an integral part of network operator tool sets to keep the net up and running.

But measurements are equally important for the research community to understand network traffic as well as protocol and application dynamics and their evolution. And they assist in quantifying application and (access) network performance and thus provide a tool for end users and regulators to monitor operators and their service level agreements. Tools such as speedtest.net have become widely used for individual measurements and basic ISP rating. Measurement service providers such as SamKnows or RIPE offer networks of probes, i.e., separate devices or embedded software on access routers, for continuous background measurements at the end users. These help ISPs and regulators in their work. Standards bodies such as the IETF and the Broadband Forum have established working groups to define a global measurement architecture and common interfaces and to extend the set of metrics describing communication properties.

This Dagstuhl seminar brought together researchers from industry, academia, and regulators to discuss the state of the art in measurements and their exploitation, measurement and analysis techniques, privacy and anonymization, and to contribute to a common understanding in a number of areas, including:

- improving the expressiveness of measurement metrics (and

develop appropriate new ones) beyond throughput, loss rate, and RTT so that the actual application-specific user *quality of experience* can be assessed;

- expanding the reach, scale, and diversity of measurements and the corresponding data analysis to obtain a more comprehensive view on the performance of networks and applications;
- structuring the otherwise mostly disconnected measurement activities to allow interfacing between them and/or providing defined access methods to them, for both carrying out measurements and accessing measurement results (offline and in real-time);
- providing ways to better instrument and more broadly utilize measurement infrastructure, inside operators, for end users, and at third parties.

Because the means for taking steps towards achieving the above goals was on learning about and from each other and developing joint perspectives, the seminar chose an extremely interactive organization comprising three elements:

1. Individual presentations were limited to an initial round of introductions (1 slide each) covering a set of questions for the participants to get know each others background and interests.
2. Panel discussions (with ample involvement of the “audience”) set the stage for the discussion topics of the day.
3. Extensive group work to dive into a number of topics and also for presenting and discussing the group outcome on the next day.

A side effect of this organization is that there were virtually no individual talks and hence no talk abstracts were collected.

We focused on two complementary aspects of a global measurement framework: 1) *creating* a global measurement framework and 2) *using* such a framework. Both were introduced by panels, with a lot of discussion contributing to these overviews, as described in the respective introduction to the following two sections.




Fig. 4.10

Gabriele Stirl – Maoam. Part of the Dagstuhl art collection and donated by: Holger Schlingloff and participants in Dagstuhl Seminar 06411.

4.69 Unleashing Operational Process Mining

Organizers: Rafael Accorsi, Ernesto Damiani, and Wil van der Aalst
Seminar No. 13481

Date: November 24–29, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.11.154

© Creative Commons BY 3.0 Unported license

© Rafael Accorsi, Ernesto Damiani, and Wil van der Aalst



Participants: Rafael Accorsi, Antonio Caforio, Josep Carmona, Paolo Ceravolo, Jan Claes, Jonathan Cook, Ernesto Damiani, Massimiliano de Leoni, Benoit Depaire, Simon N. Foley, Luciano Garcia-Banuelos, Christian Günther, Anna A. Kalenkova, Akhil Kumar, Geetika T. Lakshmanan, Teemu Lehto, Marcello Leida, Maria Leitner, Fabrizio Maria Maggi, Ronny S. Mans, Alexey A. Mitsyuk, Jorge Munoz-Gama, Zbigniew Paszkiewicz, Alessandro Provetto, Hajo A. Reijers, Joel Tiago Ribeiro, Roland Rieke, Stefanie Rinderle-Ma, Anne Rozinat, Ricardo Seguel, Marcos Sepulveda Fernandez, Sergey A. Shershakov, Pnina Soffer, Minseok Song, Alessandro Sperduti, Thomas Stocker, Wil van der Aalst, Boudewijn van Dongen, Frank van Geffen, Eric Verbeek, Thomas Vogelgesang, Jianmin Wang, Barbara Weber, Ton Weijters

Society shifted from being predominantly “analog” to “digital” in just a few years. This has had an incredible impact on the way we do business and communicate. Gartner uses the phrase “The Nexus of Forces” to refer to the convergence and mutual reinforcement of four interdependent trends: social, mobile, cloud, and information. The term “Big Data” is often used to refer to the incredible growth of data in recent years. However, the ultimate goal is not to collect more data, but to turn data into real value. This means that data should be used to improve existing products, processes and services, or enable new ones.

Event data are the most important source of information. Events may take place inside a machine (e.g., an X-ray machine or baggage handling system), inside an enterprise information system (e.g., an order placed by a customer), inside a hospital (e.g., the analysis of a blood sample), inside a social network (e.g., exchanging e-mails or twitter messages), inside a transportation system (e.g., checking in, buying a ticket, or passing through a toll booth), etc.

Process mining aims to *discover, monitor and improve real processes by extracting knowledge from event logs* readily available in today’s information systems¹⁴. The starting point for process mining is an *event log*. Each event in such a log refers to an *activity* (i.e., a well-defined step in some process) and is related to a particular *case* (i.e., a *process instance*). The events belonging to a case are *ordered* and can be seen as one “run” of the process. Event logs may store additional information about events. In fact, whenever possible, process mining techniques use extra information such as the *resource* (i.e., person or device) executing or initiating the activity, the *timestamp* of the event, or *data elements* recorded with the event (e.g., the size of an order).

Event logs can be used to conduct three types of process

mining. The first type of process mining is *discovery*. A discovery technique takes an event log and produces a model without using any a-priori information. Process discovery is the most prominent process mining technique. For many organizations it is surprising to see that existing techniques are indeed able to discover real processes merely based on example behaviors stored in event logs. The second type of process mining is *conformance*. Here, an existing process model is compared with an event log of the same process. Conformance checking can be used to check if reality, as recorded in the log, conforms to the model and vice versa. The third type of process mining is *enhancement*. Here, the idea is to extend or improve an existing process model thereby using information about the actual process recorded in some event log. Whereas conformance checking measures the alignment between model and reality, this third type of process mining aims at changing or extending the a-priori model. For instance, by using timestamps in the event log one can extend the model to show bottlenecks, service levels, and throughput times.

Process mining algorithms have been implemented in various academic and commercial systems. The corresponding tools are being increasingly relevant in industry and have proven to be essential means to meet business goals. ProM is the de facto standard platform for process mining in the academic world. Examples of commercial tools are Disco (Fluxicon), Perceptive Process Mining (before Futura Reflect and BPM|one), QPR ProcessAnalyzer, ARIS Process Performance Manager, Celonis Discovery, Interstage Process Discovery (Fujitsu), Discovery Analyst (StereoLOGIC), and XMAlyzer (XMPro). Representatives of ProM community and the first three commercial vendors participated in Dagstuhl Seminar 13481 “Unleashing Operational Process Mining”.

¹⁴ Process Mining: Discovery, Conformance and Enhancement of Business Processes by W.M.P. van der Aalst, Springer Verlag, 2011 (ISBN 978-3-642-19344-6).

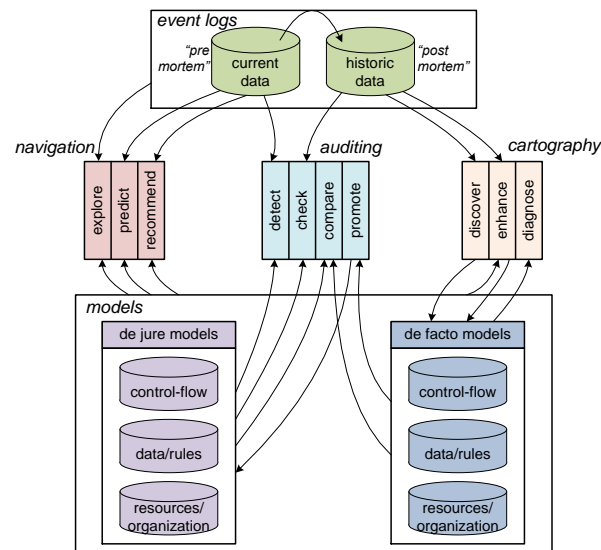


Fig. 4.11

Overview of the different process mining tasks. Taken from "Process Mining: Discovery, Conformance and Enhancement of Business Processes."

The Dagstuhl Seminar was co-organized with the *IEEE Task Force on Process Mining* (see <http://www.win.tue.nl/ieeetfpm/>). The goal of this Task Force is to promote the research, development, education and understanding of process mining. Sixty organizations and over one hundred experts have joined forces in the IEEE Task Force on Process Mining.

Next to some introductory talks (e.g., an overview of the process mining field by Wil van der Aalst), 31 talks were given by the participants. The talks covered the entire process mining spectrum, including:

- from theory to applications,
- from methodological to tool-oriented,
- from data quality to new analysis techniques,
- from big data to semi-structured data,
- from discovery to conformance,
- from health-care to security, and
- from off-line to online.

It was remarkable to see that all participants (including the academics) were very motivated to solve real-life problems and considered increasing the adoption of process mining as one of the key priorities, thereby justifying the title and spirit of the seminar, namely "Unleashing the Power of Process Mining". This does not imply that there are not many foundational research challenges. For example, the increasing amounts of event data are creating many new challenges and new questions have emerged. Such issues were discussed both during the sessions and on informal meetings during the breaks and at the evening.

Half of the program was devoted to discussions on a set of predefined themes. These topics were extracted based on a questionnaire filled out by all participants before the seminar.

1. Process mining of multi-perspective models (Chair: Akhil Kumar)
2. Data quality and data preparation (Chair: Frank van Geffen)
3. Process discovery: Playing with the representational bias (Josep Carmona)
4. Evaluation of process mining algorithms: benchmark data sets and conformance metrics (Chair: Boudewijn van Dongen)

5. Advanced topics in process discovery: on-the-fly and distributed process discovery (Chair: Alessandro Sperduti)
6. Process mining and Big Data (Chair: Marcello Leida)
7. Process mining in Healthcare (Chair: Pnina Soffer)
8. Security and privacy issues in large process data sets (Chair: Simon Foley, replacing Günter Müller)
9. Conformance checking for security, compliance and auditing (Chair: Massimiliano De Leoni, replacing Marco Montali)
10. How to sell process mining? (Chair: Anne Rozinat)
11. What is the ideal tool for an expert user? (Chair: Benoit Depaire)
12. What is the ideal tool for a casual business user? (Chair: Teemu Lehto)

The chairs did an excellent job in guiding the discussions. After the each discussion participants had a better understanding of the challenges that process mining is facing. This definitely include many research challenges, but also challenges related to boosting the adoption of process mining in industry.

The social program was rich and vivid, including an exclusion to Trier's Christmas market, a night walk to ruins, table football, table tennis, and late night discussions.

A tangible output of the seminar is a special issue of IEEE Transactions on Services Computing based on the seminar. This special issue has the title "Processes Meet Big Data" and will be based on contributions from participants of this seminar (also open to others). This special issue of IEEE Transaction on Service-Oriented Computing is intended to create an international forum for presenting innovative developments of process monitoring, analysis and mining over service-oriented architectures, aimed at handling "big logs" and use them effectively for discovery, dash-boarding and mining. The ultimate objective is to identify the promising research avenues, report the main results and promote the visibility and relevance of this area.

Overall, the seminar was very successful. Most participants encouraged the organizers to organize another Dagstuhl Seminar on process mining. Several suggestions were given for such a future seminar, e.g., providing event logs for competitions and complementary types of analysis before or during the seminar. These recommendations were subject of the discussion sessions, whose summaries can be found below.

4.70 Forensic Computing

Organizers: Felix C. Freiling, Gerrit Hornung, and Radim Polčák
Seminar No. 13482

Date: November 24–29, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.11.193
 © Creative Commons BY 3.0 Unported license
 © Felix C. Freiling, Radim Polčák and Gerrit Hornung



Participants: Andreas Dewald, Jan Ellermann, Hannes Federrath, Felix C. Freiling, Michael Gruhn, Christian Hawellek, Dennis Heinson, Dominik Herrmann, Gerrit Hornung, Sven Kälber, Stefan Kiltz, Volker Krummel, Radim Polčák, Thomas Schreck, Erich Schweighofer, Tobias Singelstein, Vaclav Stupka, Tatiana Tropina, Nicolas von zur Mühlen, York Yannikos, Riha Zdenek

After a brief introduction by the organizers, the seminar started off with a sequence of 3 slide/5 minute talks by all participants stating their research interests, their background and their expectations towards the seminar. In the afternoon, three motivation talks by Felix Freiling (“What is forensic computing?”), Gerrit Hornung (“The fundamental rights dimension of individual and mass surveillance”) and Radim Polčák (“Experiences from drafting the cybersecurity act in CZ”) paved the way for a common understanding of the open questions in the area and the relation of forensic computing to computer security law.

The rest of the afternoon questions and expectations were collected and grouped using moderation cards. The result was a spectrum of five areas of interest that we termed as follows:

1. technical possibilities for evidence collection
2. digital evidence: admissibility, spoofing, integrity protection
3. open source intelligence
4. investigations vs. privacy
5. offensive countermeasures
6. transborder/cloud evidence collection

For immediate discussion on Tuesday the participants voted for their favorite topics. As a result, three discussion groups were formed for the next day: digital evidence (topic 2), investigations vs. privacy (topic 4) and offensive countermeasures (topic 5). Topic 1 was to be handled by an overview talk by Andreas Dewald on the following day.

Tuesday morning started with a talk by Andreas Dewald on technically unavoidable evidence and was followed by a multimedia presentation about cold boot and hot re-plug attacks. After this technical introduction work in the discussion groups took place until the afternoon, when the collected results of the discussion groups were presented in a plenary session. As a highlight, the group on offensive countermeasures presented a taxonomy of 5 categories of offensive countermeasures that were specific enough for both law and computer science to investigate.

Wednesday morning commenced with a talk about the work of Interpol by Jan Ellermann (“Data protection as an asset in Europol’s fight against cybercrime”). It was followed by a presentation of current research by Dominik Herrmann about the usage of fingerprinting in network forensics (“Fingerprinting Techniques for Network Forensics”). The round of talks was concluded by an introduction to the law of evidence in criminal procedural law by Tobias Singelstein (“Basics zum Beweisrecht im Strafverfahren”).

The afternoon was spent on a pleasant hike to a nearby village where the Dagstuhl office had organized delicious traditional coffee and cake. On the way back to Schloss Dagstuhl a group of adventurers again, as in 2011, separated from the main party to explore the woods around Wadern. However, unlike 2011, they managed to return to Dagstuhl in time without major difficulties.

Thursday was started with a talk by Dennis Heinson on investigations in enterprises (“Internal Investigations, IT Forensics and Law”). Afterwards two new discussion groups were formed, partly based on the areas of interest collected on Monday, and commenced discussing the topics of (1) internal investigations and (2) transborder/cloud issues. In the afternoon, the results of these groups were collected in a plenary session during which especially the transborder issues caused a heated and insightful discussion.

Friday morning hosted a series of three talks from computer science, law and practice by Christian Hawellek (on techniques for modeling surveillance), Stefan Kiltz (“Forensically Sound Data for Digitised Forensics on the Example of Locksmith Forensics”) and – last but not least – Erich Schweighofer (“Surveillance of US-surveillance”).

■ Conclusion

In summary, the participants (and the organizers) enjoyed the week in Dagstuhl. In particular, the chance to get to know

many new people from both the technical and the legal side of forensic computing was appreciated. From the viewpoint of the organizers, several points appear worth mentioning which we wish to document here.

First of all, it became clear to all participants that forensic computing is still in the process of maturing. The legal regulations as well as the technical instruments used in forensic computing are evolving quickly and it needs a joint effort by both communities to make progress. In our opinion, the seminar was much better than the preceding seminar in 2011, mainly because the lawyers were more interested in technical details and the technical people presented their “special secret instruments” in an understandable way. The seminar showed that fruitful discussions between both sides are possible, that lawyers can be cool as well and that there exist at least some lawyers with advanced technical understanding. For the technical people it was insightful to get a basic feeling on how the interpretation of law works and to see that there are quite a lot of gray legal areas. After all, forensic expertise is just one bit of evidence in court, and it may not be the most important one. And there are actually many, many data protection problems out there that will need to be handled within the field of forensic computing.

Overall, it was again a challenge to gather interested people in Dagstuhl. Dagstuhl seminars are well-known in computer science, but not in law, and it is well-known that practitioners, which are common in forensic computing (prosecutors, defenders, police, expert witnesses), with their tight time schedules can hardly afford to come to Dagstuhl for an entire week, especially from overseas. This is a problem which will remain and explains why – again – the seminar was dominated by German speaking participants.

The topic of forensic computing, however, is also gaining importance in the academic community, and at Dagstuhl: In February 2014, a seminar on “Digital Evidence and Forensic Readiness” (Dagstuhl Seminar 14092) will take place, opening the possibility for several of the participants to meet and discuss again, albeit with a slightly sharpened focus. In case another general seminar like this would take place, the topic of mutual understanding can be placed into focus even stronger. This could be achieved by distributing introductory papers from “the other side” in advance or by giving introductory tutorials in forensic techniques at the seminar. In the end, the seminar left us with more open questions than we had at the beginning. But at least this was to be expected.

4.71 Computational Mass Spectrometry

Organizers: Rudolf Aebersold, Oliver Kohlbacher, and Olga Vitek
Seminar No. 13491

Date: December 1–6, 2013 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.3.12.1
 © Creative Commons BY 3.0 Unported license
 © All participants of Dagstuhl Seminar 13491



Participants: Rudolf Aebersold, Theodore Alexandrov, Dario Amodei, Bernd Bodenmiller, Sebastian Böcker, Karsten Boldt, Daniel R. Boutz, Julia Burkhart, Manfred Claassen, John Cottrell, Eric Deutsch, Joshua Elias, David Fenyö, Anne-Claude Gingras, Henning Hermjakob, Lukas Käll, Sangtae Kim, Oliver Kohlbacher, Theresa Kristl, Bernhard Küster, Henry Lam, Wolf D. Lehmann, Kathryn Lilley, Michal Linial, Mike MacCoss, Brendan MacLean, Alexander Makarov, Lennart Martens, Sara Nasso, Alexey Nesvizhskii, Steffen Neumann, William Stafford Noble, Paola Picotti, Knut Reinert, Bernhard Renard, Hannes Röst, Stephen Tate, Andreas Tholey, Henning Urlaub, Olga Vitek, Christian von Mering, Susan T. Weintraub, Witold E. Wolski, René Zahedi

■ Motivation

Mass Spectrometry (MS) is an analytical technique of immense versatility. Detection of explosives at airports, urine tests for doping in sports, tests for cancer biomarkers in a clinic – all these rely on mass spectrometry as the key analytical technique. During the last decade, technological advances have resulted in a flood of mass spectrometric data (high-resolution mass spectrometry, mass spectrometry coupled to high-performance liquid chromatography – HPLC-MS). The publication of the first human genome in 2001 was a key event that enabled the explosive development of proteomics, which led to the conception of the Human Proteome Project in 2010. Today, mass spectrometric techniques are an indispensable tool in the life sciences. Their development, however, is more and more hampered by the lack of computational tools for the analysis of the data. Modern instrumentation can easily produce data sets of hundreds of gigabytes from an individual sample. Most experimental groups are no longer able to deal with both the amount and the inherent complexity of these data. Computer science has the necessary tools to address these problems. It is thus necessary to intensify collaboration between the three key communities involved: life scientists applying MS; analytical chemists and engineers developing the instruments; computer scientists, bioinformaticians and statisticians developing algorithms and software for data analysis.

■ Goals

The seminar 'Computational Mass Spectrometry' is a follow-up seminar to the successful Dagstuhl seminars on 'Computational Proteomics (05471 and 08101)'. The different title was chosen to reflect the growing scope of computational mass spectrometry: from proteomics to metabolomic, lipidomics, and glycomics.

The goal of the seminar was thus to assess the state of the art for the field of computational mass spectrometry as a whole and to identify the challenges the field will be facing for the years to come. To this end we put together a list of participants covering both computational and experimental aspects of mass spectrometry from industry and academia from around the world. The result of these discussions should then be summarized in a joint status paper.

■ Results

The seminar was very productive and led to a number of tangible outcomes summarized below.

■ The Big Challenges

Not unexpectedly, it turned out to be difficult to identify the *big challenges* of the coming years and views on this differed quite a bit. After lengthy discussions, we were able to categorize the challenges. We are currently in the process of finalizing the draft of a paper on these challenges for computational mass spectrometry, which is supposed to be submitted by end of March 2014. The paper is a joint work of all the participants and will document the current state of the field. The challenges identified were the following:

■ Challenges of computational and statistical interpretation of mass spectra

■ Identification

Identification of analytes is still a challenge. In proteomics, the identification of post-translational modifications and of different proteoforms pose problems. Also the identification of non-tryptic peptides (peptidomics, MHC ligands) are interesting problems. Estimation of false-discovery rates based on target-decoy approaches has been criticized, but

there is still a distinct lack of established alternatives. With the increasing interest in small-molecule mass spectrometry, the identification of metabolites, glycans, and lipids is increasingly becoming an issue and the algorithmic support for this is currently still lacking.

■ *Quantification*

Quantification faces challenges due to the – still-growing – diversity of experimental methods for analyte quantifications that necessitate a permanent development of new computational approaches. There are also more fundamental, statistical problems, for example, inferring the absence of an analyte based on the absence of a signal. Quantification is also expected to contribute to the understanding of protein complexes and their stoichiometry.

Challenges arising from new experimental frontiers

■ *Data-independent acquisition*

The recent developments of data-independent acquisition techniques resulted in a set of entirely new computational challenges due to the different structure of the underlying data.

■ *Imaging*

Imaging mass spectrometry has become mature on the experimental side. The analysis of spatially resolved MS data, however, poses entirely new problems for computational mass spectrometry with increased complexity and data volume.

■ *Single-cell mass spectrometry*

Multi-parameter single cell mass spectrometry enables the characterization of rare and heterogeneous cell populations and prevents the typical averaging across a whole tissue/cell population. The key challenge will be the development of new computational tools able to define biologically meaningful cell types and then model the dynamic behaviour of the biological processes.

■ *Top-down proteomics*

Despite its obvious advantages of top-down approaches for functional proteomics, isoform identification and related topics, the approach suffers from unmet challenges on the computational side. Methods for mass spectrum deconvolution need to be improved and algorithms for the identification of multiple PTM sites are required.

Challenges of extracting maximal information from datasets

■ *Democratization of data*

Public availability of large datasets enables novel types of studies in computational mass spectrometry (data mining). The standardized deposition in and reliable repositories handling this data is still a major problem that needs to be addressed.

■ *Integration of MS data with different technologies*

Increasingly, computational biologists face data from multiple omics technologies. Integrating data from computational mass spectrometry across omics levels (genomics with transcriptomics, transcriptomics with proteomics, proteomics with metabolomics) poses a difficult data integration challenge, but will be essential for a more comprehensive view of the biological systems under study.

■ *Visualization of heterogeneous data sets*

The amount, structure and complexity of large-scale mass spectrometric data turns out to be a challenging issue. While some end-users of these methods tend to be interested in a final, aggregated result of a complex data analysis pipeline, it is often essential to analyze the data conveniently down to the raw spectra. Tools navigating these data sets on all levels are currently not yet available.

■ **Community Building**

It was felt among participants that computational mass spectrometry is lacking a structured community. Researchers in computational mass spectrometry come from diverse backgrounds: statistics, computer science, analytical sciences, biology, or medicine. Traditionally they are thus organized in different scientific societies, for example the International Society on Computational Biology (ISCB), the American Society of Mass Spectrometry (ASMS), the Human Proteome Organization (HUPO), the Metabolomics Society, and of course various national societies. Many participants attend both computational and experimental conferences in the area of mass spectrometry organized by these different organizations. Participants suggested to form subgroups for computational mass spectrometry in different societies. At the same time, in order to avoid duplication of structures and efforts, it was planned to share these subgroups across the different societies and establish joint chairs of these groups, organize joint workshops, and coordinate educational activities.

After the Dagstuhl seminar we contacted ISCB and HUPO to discuss the formation of these subgroups. After intensive discussion with the societies, HUPO and ISCB both agreed to this plan. A HUPO subgroup CompMS on computational mass spectrometry was formed. In parallel, ISCB agreed to form a *Community of Special Interest* (CoSI) CompMS. Both subgroups share a joint structure. A joint steering committee (Steering Committee (Oliver Kohlbacher, Olga Vitek, Shoba Ranganathan, Henning Hermjakob, and Ruedi Aebersold) has been established to guide both groups through their formation period. The groups have set up a joint mailing list, a website, and are currently planning initial kick-off meetings as satellite workshops to ISMB 2014 (Boston) and HUPO 2014 (Madrid).

■ **Teaching Initiative**

Recognizing the great need for educational materials for various audiences (bioinformaticians, biologists, computer scientists) some participants initiated an initiative to put these materials together as online courses. Discussions of this initiative have come quite far. It is currently planned to come up with a core curriculum for mass spectrometry. This core curriculum will be open for discussion within the computational mass spectrometry community. After the contents of the core curriculum has been established, tutorial papers will be solicited for the various modules of the curriculum. These papers will refer to each other, will use a coherent vocabulary and notation and will appear as a paper collection online in PLoS Computational Biology (edited by Theodore Alexandrov). Additional materials will be included, for example, online courses and lecture videos. An initial tutorial workshop is currently in planning to kickstart the further development of the curriculum.

■ **Reviewing Guidelines**

A working group discussed the problems that computational papers face in the reviewing process. The main driver for this discussion was expediting the review process, specifically in terms of reducing the number of review cycles. It is worth noting that the Journal of Proteome Research (JPR), published by the American Chemical Society (ACS), presents a special case since this journal is the only one in the field that does not have a regular mechanism for the reviewers to see the comments of the other reviewers and the corresponding responses of the authors after each round of review. The proposal initially on the table was to share all reviews among reviewers and invite comments and changes before the first editorial decision is made for the first round of review. This system is, for instance, already in place via EasyChair (software used for the RECOMB meetings, but not for proteomics journals). After

discussion, it was decided that it would clearly be beneficial if the JPR distributed all reviews among reviewers after each stage of revision. But it was felt that it would only be necessary to collect comments and feedback from the reviewers (based on sending them all reviews) before the editor reached an initial decision in cases where there was substantive disagreement among reviewers on critical points. These ad hoc communications can be handled

in a semi-manual way within the existing manuscript management systems used by the proteomics journals, with the added benefit of maintaining an audit trail for the process. The reviewing guidelines developed by the participants in Dagstuhl are currently being discussed by with the editorial boards of different journals (currently J. Proteome Res and Mol. Cell. Prot.).



Fig. 4.12

Maria Krause – Konzentration. Part of the Dagstuhl art collection and donated by Maria Krause.

4.72 Geosensor Networks: Bridging Algorithms and Applications

Organizers: Matt Duckham, Stefan Dulman, Jörg-Rüdiger Sack, and Monika Sester

Seminar No. 13492

Date: December 1–6, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.12.17

© Creative Commons BY 3.0 Unported license

© Matt Duckham, Stefan Dulman, Jörg-Rüdiger Sack, and Monika Sester



Participants: Annalisa Appice, Jacob Beal, Susanne Bleisch, Nikolaus Correll, René Doursat, Matt Duckham, Stefan Dulman, Daniel Fitzner, Allison Kealy, Thomas Kirste, Steve Liang, Marco Mamei, Sabrina Merkel, Zoltan Papp, Edzer Pebesma, Bernd Resch, Kay Römer, Jörg-Rüdiger Sack, Hedda R. Schmidtke, Monika Sester, Martijn Warnier

The aims of Dagstuhl Seminar 13492, “Geosensor Networks: Bridging Algorithms and Applications,” were to advance research into, and application of geosensor networks by enhancing interdisciplinary and cross-domain collaboration. The premise of the Seminar was that the potential for useful, practical applications of geosensor networks (wireless sensor networks tasked with monitoring changes in geographic space) are being held back by the enormous diversity of applications and expertise connected with different facets of geosensor networks. The result today is many niche solutions to specific problems, where what is needed are a few general solutions to broader problems.

More specifically, the diversity of concepts, approaches, and tools used in connection with geosensor networks is inhibiting more rapid and fruitful research and development. Examples of this diversity include:

- *Ontologies, representations, and models:* The diversity of uses of geographic information leads inexorably to diversity in the different ontologies, representations, and models commonly used to conceptualize that information.
- *Algorithms and data structures:* The plethora of models, algorithms, data structures, and architectures that exist in the literature are frequently incompatible, founded on divergent assumptions and inconsistent approaches.
- *Applications:* While the potential applications of geosensor networks are legion, today we lack agreement on a set of applications that together encompass and illustrate the bulk of issues faced by all applications.
- *Benchmarks, tools, and technologies:* Perhaps more than any other single issue, a lack of consensus on core benchmark data sets, problems, simulation systems, and software tools inhibits convergence in research and application.

■ Participants

To begin to address this diversity, and bridge the gap between theory and application, the Seminar participants represented a broad spectrum of disciplines, including computer science, geographic information science, computational geometry, statistics, artificial intelligence, pervasive computing. The seminar had strong groundings in previous Dagstuhl Seminars, including Seminars 10491 and 12512 on Representation, Analysis and Visualization of Moving Objects, and Seminar 06361 on Computing Media and Languages for Space-Oriented Computation. However, we were very pleased that the Seminar also attracted a significant proportion of newer Dagstuhl attendees: more than half the attendees had attended at most one Dagstuhl seminar before, with around one quarter of Seminar participants attending their first ever Dagstuhl.

Bringing together this diversity of backgrounds, expertise, and experiences was central to the core aim of building bridges between related fields, and was central to the success of the Seminar.

■ Format

The seminar was structured around three complementary perspectives: models and algorithms; benchmarking and applications; and teaching and curricula. The objective of the models and algorithms perspective was to survey, catalog, and compare the ontologies, representations, algorithms, and data structures that are fundamental to computing in geosensor networks. Through the benchmarks and application perspective, the seminar aimed to improve comparability and compatibility in models and algorithms, as well as connect existing models and algorithms more directly to practical uses. Focal applications included emergency response, intelligent transportation, smart materials, and environmental monitoring. As a capstone, the teaching and curricula

perspective aims to distill and collate the collected expertise in models, algorithms, benchmarks, and applications into a coherent body of knowledge: a “library” of core concepts and techniques for computation with and application of geosensor networks.

The seminar focused less on presenting individual lectures than on achieving its objectives above through collaborative discussions and activities. The organizers invited three speakers with diverse backgrounds to give longer talks (40 minutes) and lead the subsequent discussions. The three speakers were René Doursat (CNRS, Paris, discussing organic computing), Thomas Kirste (Universität Rostock, discussing situational awareness and intention recognition using sensed data), and Edzer Pebesma (Universität Münster, on spatial data analysis with sensor data). As well as providing an introduction to the breadth and depth of ideas related to the field, the speakers were able to inspire the participants and spark many subsequent discussions.

The majority of the seminar then focused on workshop-style discussion and break-out groups. In this way the seminar aimed to elicit answers to the question of what are the essential elements of computing with geosensor network. The aim was to advance the field through consensus on priorities as well as providing opportunities for new innovations to emerge from new collaborations. The working groups’ discussions and conclusions are summarized in this report (Section 5). Broadly, the working groups’ focuses included spatial computing (e.g., self-organization and smart materials); applications of sensor networks (e.g., developing countries and big data); social issues (e.g., privacy); education (e.g., teaching resources and curriculum); and data and benchmarking.

However, even though the primary focus was on discussion and collaboration, the program still allowed time for short focus talks from participants (up to 10 minutes for senior researchers, or up to 15 minutes for junior researchers—researchers were able to self-select as to whether they regarded themselves as junior or senior). A summary of the focus talks given by participants is also contained in this report. All the speakers were asked to address one of the three Seminar perspectives (models and algorithms; benchmarking and applications; and teaching and curricula), as can be seen from section 4.

■ Outcomes

The participants were highly satisfied with the quality of the seminar. Many and diverse research results were presented during the Seminar, surveyed in the following sections. As with many Dagstuhl Seminars, the new collaborations and results of those

collaborations are ongoing. However, amongst the key findings and ongoing collaborations, we highlight:

- Considerable progress has been made in recent years in the areas of (decentralized) spatial computing. This includes advances in the bottom-up design of distributed and decentralized algorithms. However, in contrast top-down aggregate programming techniques offer an important advantage over more conventional decentralized programming, in that they are substantially less complex for developers to use. Whatever future advances in this area may hold, the Seminar participants were agreed in that decentralization is a means, but never an ends. A focus on the behavior of a distributed geosensor network as a whole, rather than the rules required to generate that behavior, should always be the focus. In this respect, the focus on emergent behaviors found in spatial computing would seem ideally suited.
- One particular area of progress at the seminar was in teaching and curriculum. Despite the wide range of expertise and academic backgrounds of the participants, the Seminar exposed the considerable commonality and agreement around the fundamental concepts behind geosensor networks. This convergence was evident in the recent development of tools (such as the Proto aggregate programming language) and the publication of text books on the subject of geosensor networks.
- In contrast, one area of particular difficulty was in benchmarking and data sets. The availability of data sets is unquestionably increasing, as evidenced by several different data sets that were made available at, and through the preparatory work by participants in advance of the Seminar. However, the wide diversity of requirements for data sets across different applications continues to defy standardization or convergence on a small set of benchmarks. Issues such as validation and ground truthing; requirements for massive data sets with thousands or millions of sensors; metadata and provenance; and privacy issues were all various inhibitors to the development of a small set of benchmark data sets.
- Finally, the Seminar highlighted the numerous practical, social, and environmental challenges that still remain to truly bridging the gap between theoretical algorithms and practical applications, such as cost and deployment strategies, privacy, and environmental pollution. Although these issues remain largely unsolved longer-term research problems, they are already the explicit focus of several new collaborations that have directly resulted from this Seminar.

4.73 Approaches and Applications of Inductive Programming

Organizers: Sumit Gulwani, Emanuel Kitzelmann, and Ute Schmid
Seminar No. 13502

Date: December 8–11, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.12.43

© Creative Commons BY 3.0 Unported license

© Sumit Gulwani, Emanuel Kitzelmann, and Ute Schmid



Participants: Umair Zafrulla Ahmed, Sam Bayless, Tarek R. Besold, Luc De Raedt, Eyal Dechter, Alexey Grigoryev, Sumit Gulwani, Mike Hansen, Robert J. Henderson, José Hernández-Orallo, Petra Hofstedt, Frank Jäkel, Susumu Katayama, Dileep Kini, Emanuel Kitzelmann, Mark Marron, Fernando Martínez-Plumed, Martin Möhrmann, Stephen H. Muggleton, Daniel Perelman, Iurii Perov, Ruzica Piskac, Oleksandr Polozov, Marco Ragni, Ute Schmid, George M. Sergievsky, Michael Siebers, Rishabh Singh, Armando Solar-Lezama, Janis Voigtländer, David White, Eran Yahav, Benjamin Zorn

Inductive programming (IP) research addresses the problem of learning programs from incomplete specifications, such as input/output examples, traces, or constraints. In general, program synthesis is a topic of interest to researchers in artificial intelligence as well as in programming research since the 1960s [2]. On the one hand, this research aims at relieving programmers from the tedious task of explicit coding on the other hand it helps to uncover the complex cognitive processes involved in programming as a special domain of complex problem solving. From the beginning, there were two main directions of research – deductive knowledge based approaches and inductive machine learning based approaches. Due to the progress in machine learning, over the last decades the inductive approach currently seems to be the more promising.

Researchers working on the topic of IP are distributed over different communities, especially inductive logic programming (ILP) [6, 12], evolutionary programming [13], functional programming [5, 10, 15], grammar inference [1], and programming languages and verification [7]. Furthermore, domain specific IP techniques are developed for end-user programming [4, 9] and in the context of intelligent tutoring in the domain of programming [8]. In cognitive science, researchers concerned with general principles of human inductive reasoning have constructed computer models for inductive generalization which also have some relation to IP [3, 16].

In general, approaches can be classified by (1) the strategy of program construction which can be example-driven or generate-and-test driven; (2) the implicit or explicit restriction bias which can be Horn clauses, functional programs, or domain specific languages possibly with further constraints given as meta-interpreters, templates or program schemes; (3) the possibility to consider background knowledge.

IP research had its first boost in the 1970s in the context of learning Lisp programs from examples. Due to only limited

progress, this direction of research decayed and in the 1990s was newly addressed in the context of ILP and evolutionary programming. Again, after first promising results, disappointment set in [11, 14]. However, over the last years, a new revival in IP research can be observed in different communities and promising results, for example in the domain of enduser programming, give rise to new expectations.

Therefore, in the Dagstuhl Seminar AAIP we brought together researchers from these different communities as well as researchers of related fields. The possibility to discuss and evaluate approaches from different perspectives helped to (a) gaining better insights in general mechanisms underlying inductive programming algorithms, (b) identifying commonalities between induction algorithms and empirical knowledge about cognitive characteristics of the induction of complex rules, and (c) open up new areas for applications for inductive programming in enduser programming, support tools for example driven programming, and architectures for cognitive systems.

The presentations covered several aspects of inductive programming and were grouped in the topic sessions

- Inductive Programming Systems and Algorithms (with an introductory talk by Stephen Muggleton),
- Enduser Programming (with an introductory talk by Sumit Gulwani),
- Intelligent Tutoring and Grading,
- Cognitive Aspects of Induction (with an introductory talk by José Hernández-Orallo),
- Combining Inductive Programming with Declarative Programming and with Other Approaches to Program Synthesis (with an introductory talk by Luc de Raedt).

In an initial discussion round three focus topics were identified and further discussed in working groups

- Comparing Inductive Logic and Inductive Functional Programming as well as other Approaches to Program Synthesis,
- Potential New Areas of Applications and Challenges for Inductive Programming,
- Benchmarks and Metrics.

■ Concluding Remarks and Future Plans

In the final panel discussion the results of the seminar as well as future plans were identified. Participants stated that they learned a lot about different inductive programming techniques

and tools to try. The general opinion was that it was very inspiring to have researchers from different backgrounds. To facilitate mutual understanding it was proposed to give introductory lectures, define the vocabulary of the different groups, collect a reading list, and identify common benchmark problems.

To progress in establishing inductive programming as a specific area of research it was proposed to write a Wikipedia page, and to collect introductory literature from the different areas covered in the seminar. Furthermore, plans for joint publications and joint grant proposals were made.

This seminar was highly productive and everybody hoped that there will be a follow-up in the near future.

■ References

- 1 D. Angluin and C. H. Smith. Inductive inference: theory and methods. *Computing Surveys*, 15(3):237–269, September 1983.
- 2 A. W. Biermann, G. Guiho, and Y. Kodratoff, editors. *Automatic Program Construction Techniques*. Macmillan, New York, 1984.
- 3 P. A. Carpenter, M. A. Just, and P. Shell. What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices test. *Psychological Review*, 97:404–431, 1990.
- 4 A. Cypher. EAGER: Programming repetitive tasks by example. In *Human Factors in Computing Systems, Proc. of CHI'91*, pages 33–39. ACM Press, 1991.
- 5 C. Ferri-Ramírez, José Hernández-Orallo, and M. José Ramírez-Quintana. Incremental learning of functional logic programs. In *FLOPS '01: Proceedings of the 5th International Symposium on Functional and Logic Programming*, pages 233–247, London, UK, 2001. Springer-Verlag.
- 6 P. Flener and S. Yilmaz. Inductive synthesis of recursive logic programs: Achievements and prospects. *Journal of Logic Programming*, 41(2–3):141–195, 1999.
- 7 Sumit Gulwani. Synthesis from examples: Interaction models and algorithms. In *SYNASC*, pages 8–14, 2012.
- 8 Sumit Gulwani. Example-based learning in computer-aided stem education. *To appear in Commun. ACM*, 2014.
- 9 Sumit Gulwani, William R. Harris, and Rishabh Singh. Spreadsheet data manipulation using examples. *Commun. ACM*, 55(8):97–105, 2012.
- 10 E. Kitzelmann and U. Schmid. Inductive synthesis of functional programs: An explanation based generalization approach. *Journal of Machine Learning Research*, 7(Feb):429–454, 2006.
- 11 Tessa Lau. Why programming-by-demonstration systems fail: Lessons learned for usable ai. *AI Magazine*, 30(4):65–67, 2009.
- 12 S. Muggleton and L. De Raedt. Inductive logic programming: Theory and methods. *Journal of Logic Programming, Special Issue on 10 Years of Logic Programming*, 19-20:629–679, 1994.
- 13 Roland Olsson. Inductive functional programming using incremental program transformation. *Artificial Intelligence*, 74(1):55–83, March 1995.
- 14 J.R. Quinlan and R.M. Cameron-Jones. FOIL: A midterm report. In P. B. Brazdil, editor, *Proceedings of the ECML'93*, number 667 in LNCS, pages 3–20. Springer, 1993.
- 15 U. Schmid and F. Wysotzki. Induction of recursive program schemes. In *Proc. 10th European Conference on Machine Learning (ECML-98)*, volume 1398 of LNAI, pages 214–225. Springer, 1998.
- 16 Ute Schmid and Emanuel Kitzelmann. Inductive rule learning on the knowledge level. *Cognitive Systems Research*, 12(3):237–248, 2011.

4.74 Software Engineering for Self-Adaptive Systems: Assurances

Organizers: Rogerio de Lemos, David Garlan, Carlo Ghezzi, and Holger Giese

Seminar No. 13511

Date: December 15–19, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.12.67

© Creative Commons BY 3.0 Unported license

© Rogerio de Lemos, David Garlan, Carlo Ghezzi, and Holger Giese



Participants: Jesper Andersson, Chris Bailey, Luciano Baresi, Benoit Baudry, Nelly Bencomo, Yuriy Brun, Radu Calinescu, Javier Cámara, Myra B. Cohen, Vittorio Cortellessa, Bojan Cukic, Carlos Eduardo da Silva, Rogerio de Lemos, Antonio Filieri, Carlo Ghezzi, Holger Giese, Alessandra Gorla, Vincenzo Grassi, Lars Grunske, Paola Inverardi, Jean-Marc Jezequel, Zhi Jin, Gabor Karsai, Philip Koopman, Seok-Won Lee, Alberto Leva, Marin Litoiu, Sam Malek, Raffaella Mirandola, Marco Mori, Hausi A. Müller, Mauro Pezzè, Romain Rouvoy, Cecilia Mary Fischer Rubira, Eric Rutten, Bradley Schmerl, Mary Shaw, Giordano Tamburrelli, Gabriel Tamura, Norha Milena Villegas Machado, Thomas Vogel, Danny Weyns, Franco Zambonelli

Repairing faults, or performing upgrades on different kinds of software systems have been tasks traditionally performed as a maintenance activity conducted off-line. However, as software systems become central to support everyday activities and face increasing dependability requirements, even as they have increased levels of complexity and uncertainty in their operational environments, there is a critical need to improve their resilience, optimize their performance, and at the same time, reduce their development and operational costs. This situation has led to the development of systems able to reconfigure their structure and modify their behaviour at run-time in order to improve their operation, recover from failures, and adapt to changes with little or no human intervention. These kinds of systems typically operate using an explicit representation of their own structure, behaviour and goals, and appear in the literature under different designations (e.g., self-adaptive, self-healing, self-managed, self-*, autonomic, etc.).

In spite of recent and important advances in the area, one key aspect of self-adaptive systems that poses important challenges yet to be tackled in depth is assurances: that is, providing evidence that systems satisfy their functional and non-functional requirements during operation. Specifically, assurances involve not only system dependability, but also resilience with respect to changes that may occur in the system, its environment, or its

goals. The provision of assurances for self-adaptive systems, which should be done tandem with their development, operation and evolution, is difficult since run-time changes (e.g., resource variability) introduce a high degree of uncertainty that is atypical in more conventional systems.

This Dagstuhl Seminar has focused on the topic of obtaining assurances for self-adaptive software systems. Self-adaptive systems has been studied independently within different research areas of software engineering, including requirements engineering, modelling, architecture and middleware, event-based, component-based and knowledge-based systems, testing, verification and validation, as well as software maintenance and evolution [1, 2]. On the other hand, the topic of assurances for software-based systems has been widely investigated by the dependability community, in particular when considered in the context of safety-critical systems. For these types of systems there is the need to build coherent arguments showing that the system is able to comply with strict functional and non-functional requirements, which are often dictated by safety standards and general safety guidelines. The major challenge when combining self-adaptability and dependability is how to obtain assurances regarding the uncertainty of changes that may affect the system, its environment or its goals.

References

- 1 B. H. Cheng, H. Giese, P. Inverardi, J. Magee, and R. de Lemos. 08031 Abstracts Collection. In *Software Engineering for Self-Adaptive Systems*, Dagstuhl Seminar Proceedings, no. 08031, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Germany, 2008.
- 2 R. de Lemos, H. Giese, H. Müller, and M. Shaw. 10431 Report. In *Dagstuhl Seminar Proceedings*, no. 10431, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Germany., 2011.

4.75 Social Issues in Computational Transportation Science

Organizers: Glenn Geers, Monika Sester, Stephan Winter, and Ouri E. Wolfson
Seminar No. 13512

Date: December 15–19, 2013 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.3.12.97

© Creative Commons BY 3.0 Unported license

© Glenn Geers, Monika Sester, Stephan Winter, and Ouri E. Wolfson

Participants: Francesco Ciari, Caitlin Cottrill, Sergio Di Martino, Jan Fabian Ehmke, Rob Fitzpatrick, Stefan Funke, Glenn Geers, Sergio Iarri, Benjamin Kickhoefer, Franziska Klügl, Alexandros Labrinidis, Patrick Laube, Steve Liang, Thomas Liebig, Seng Wai Loke, Dirk Christian Mattfeld, Harvey J. Miller, Walied Othman, Nicole Ronald, Peter Sanders, Monika Sester, Sabine Störandt, Sabine Timpf, Stephan Winter, Ouri E. Wolfson, Bo Xu, Yang Yue



The Dagstuhl Seminar “Social Issues in Computational Transportation Science” (13512) brought together researchers working in various areas contributing to Computational Transportation Science (CTS). CTS is an emerging discipline that combines computer science and engineering with the modeling, planning, social, and economic aspects of transportation. It is the discipline behind intelligent transportation systems (ITS), i.e., emerging from the convergence of ICT and transportation. The discipline studies how to improve the safety, mobility, and sustainability of transportation systems by taking advantage of information technologies and ubiquitous computing.

After a first Dagstuhl Seminar on CTS in 2010 (10121), in this seminar we focused on the social computing aspect of CTS, reflecting on the potential of many recent developments in transportation, such as social networks, crowdsourcing of spatial data, persuasive technologies, and behavioural economics in social computing.

In fact, the seminar (which was a day shorter because of Christmas) consisted of three parts: a number of tutorials and short talks, a competition for the best application challenge in CTS, and a joint sketch of an introductory course on CTS. An excursion to the Christmas Market in Trier rounded up the week.

The CTS application challenge was inspired by above mentioned social aspects, such as incentives to change travel behaviour, data integration / analytics required to feed these incentives, multimodal integrated door-to-door travel, autonomous vehicles, automated crowdsourcing for travel statistics, or smart solutions for the parking problem. In case you are curious which team won the best proposal award, their proposal is online¹⁵.

The sketch of an introductory course on CTS clearly profited from the broad variety of expertise at the seminar. Everybody was

learning from the sketches of modules contributed by others, to a degree that we all wished we could take this course in full length.

Overall, this report collects material that wants to be taken into practice. We hope that we inspire teams all over the world to contribute ICT expertise for more sustainable mobility choices, and perhaps add to the development of curricula in this area.

¹⁵ <https://sites.google.com/site/karmobility/home>

5

Öffentlichkeitsarbeit

Public Relations and Outreach

Pressemitteilungen und Medienarbeit

5.1

Press Releases and Media Work

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Darstellung des Konzepts von Schloss Dagstuhl kann dabei ebenfalls berücksichtigt werden. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl¹⁶ abrufbar.

Durch Unterstützung des Saarländischen Rundfunks steht Schloss Dagstuhl ein professionelles Reporterset zur Verfügung, welches Rundfunkjournalisten erlaubt, mit Seminarteilnehmern Interviews in digitaler verlustfreier Audioqualität zu führen.

Schloss Dagstuhl hat sich im Allgemeinen zur Anlaufstelle für Journalisten etabliert, die über bestimmte Informatikthemen aber auch über Schloss Dagstuhl berichten möchten. Um junge Journalisten und Volontäre zu ermutigen, über anspruchsvolle Informatikthemen zu berichten, bietet Schloss Dagstuhl jährlich einen Workshop Wissenschaftsjournalismus an. In 2013 fand er statt vom 9. bis 12. Juni parallel zu dem Dagstuhl-Seminar 13241 „Virtual Realities.“ Als Dozenten für den Workshop konnten Tim Schröder (Wissenschaftsjournalist und Medientrainer, Oldenburg) und Gordon Bolduan (Presse Sprecher des Exzellenz-Cluster „Multimodal Computing and Interaction“ an der Universität des Saarlandes) gewonnen werden. Alle Teilnehmer als auch die Dozenten waren höchst zufrieden mit den Inhalten und Ergebnissen des Workshops. Weitere Informationen sind auf der Webseite des Workshops¹⁷ abrufbar.

In 2013 wurden mehrere Seminar-bezogene Pressemeldungen herausgegeben, siehe Fig. 5.1. Beiträge in den Medien über Seminare und sonstige Aktivitäten von Schloss Dagstuhl sind über das Internetportal von Schloss Dagstuhl nach Jahren geordnet recherchierbar. Die Pressemeldungen dienen oftmals als Initiator für Berichterstattungen durch die Medien. So wurden auch in 2013 zahlreiche journalistische Beiträge produziert (Fachpresse, Radio, TV), die einige komplexe Seminarthemen allgemein verständlich aufbereitet haben.

Schloss Dagstuhl verbreitet Neuigkeiten rund um sein Programm über soziale Netzwerkdienste wie Twitter und LinkedIn. Über Twitter-Nutzer @dagstuhl werden primär Programmankündigungen an aktuell ca. 450 Abonnenten verbreitet. Zunehmend nutzen aber auch Seminarteilnehmer den Dienst, um ihre Eindrücke vom Seminar mitzuteilen. Darüber hinaus werden über den Twitter-Nutzer @dblp_org Informationen über die Bibliographiedatenbank dblp verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ unterhalten (mit derzeit etwa 590 Mitgliedern), mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Weiterhin werden dort interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben.

Regular press releases showcase and disseminate information about current informatics topics in a comprehensible manner and clarify the concept behind Schloss Dagstuhl. Press releases and media reports that come to the center’s attention are available on the Schloss Dagstuhl website¹⁶.

Thanks to the support of the Saarländischer Rundfunk, Schloss Dagstuhl has access to professional reporting equipment that enables broadcast journalists to conduct interviews with seminar participants in digital lossless audio quality.

Schloss Dagstuhl has become a port of call for journalists seeking to report on specific informatics topics and/or on Schloss Dagstuhl itself. In order to encourage young journalists and trainees to report on complex informatics topics, Schloss Dagstuhl offers an annual workshop on science journalism. In 2013, the workshop took place on June 9–12 in parallel with Dagstuhl Seminar 13241, “Virtual Realities”. Trainers included Tim Schröder from Oldenburg (scientific writer and media trainer) and Gordon Bolduan (press relations officer at the Cluster of Excellence “Multimodal Computing and Interaction” at Saarland University). Participants as well as trainers and referees were very satisfied with the workshop. See the event webpage¹⁷ for further details.

In 2013, the center issued several seminar-related press releases as listed in Fig. 5.1. Media reports on seminars and other activities of Schloss Dagstuhl are available on the Schloss Dagstuhl website. The press releases are often picked up by the media. In 2013 journalists prepared a number of media reports (press, radio, TV) on complex topics from the seminars, making them intelligible to the general public.

News on the program of Schloss Dagstuhl are disseminated via social networks such as Twitter and LinkedIn. The Twitter handle @dagstuhl is used primarily to disseminate program announcements to about 450 followers, but is increasingly used also by Dagstuhl Seminar participants to share their impressions. Additionally, information about the dblp computer science bibliography is sent using the Twitter account @dblp_org. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with about 590 members), which supports the networking of participants in Dagstuhl Seminars. Interesting news items pertaining to Schloss Dagstuhl are also disseminated. Additionally, interesting news about Schloss Dagstuhl are announced there.

¹⁶ <http://www.dagstuhl.de/de/ueber-dagstuhl/presse/>

Fortbildung

5.2

Educational Training

5

Schloss Dagstuhl engagiert sich im schulischen Bereich durch Organisation einer jährlichen Lehrerfortbildung, die sich an Informatiklehrer im Saarland und Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Das Interesse an dieser Fortbildung stieg seit dem Beginn in 1991 stetig an und die 23. „Lehrerfortbildung in Informatik“, die vom 11. bis 13. Dezember 2013 stattfand, führte mehr Teilnehmer zusammen als jemals zuvor. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch häufen sich Anfragen zur Teilnahme von Lehrern aus anderen Bundesländern. Mehr Informationen zur Veranstaltung 2013 gibt es auf der Webseite der Veranstaltung¹⁸.

Schloss Dagstuhl holds an annual teacher training workshop specifically designed for teachers of secondary students working in the Saarland or the Rhineland Palatinate. The workshop is organized together with the Landesinstitut Pädagogik und Medien (LPM), Saarland, and the Pädagogisches Landesinstitut Rheinland-Pfalz (PL). Interest in the workshop has risen steadily since the program began in 1991 and the 23th annual Dagstuhl Teacher Training Workshop, held at Schloss Dagstuhl on December 11–13, 2013, attracted more participants than ever before. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, Schloss Dagstuhl does receive requests for participation from teachers of other federal states. Details about the workshop in 2013 are available at the event webpage¹⁸.

¹⁷ <http://www.dagstuhl.de/13242>

¹⁸ <http://www.dagstuhl.de/13503>

Wie riesige Netzwerke und komplexe Daten gemeinsam visualisiert werden Information Visualization – Towards Multivariate Network Visualization Dagstuhl Seminar 13201 http://www.dagstuhl.de/13201
Was ist Crowdsourcing? Crowdsourcing: From Theory to Practice and Long-Term Perspectives Dagstuhl Seminar 13361 http://www.dagstuhl.de/13361
Mehr Sicherheit im Straßenverkehr durch kommunizierende Autos Inter-Vehicular Communication – Quo Vadis Dagstuhl Seminar 13392 http://www.dagstuhl.de/13392
Können Computersysteme Emotionen erkennen? Computational Audio Analysis Dagstuhl Seminar 13451 http://www.dagstuhl.de/13451
Wie kann man elektronische Märkte und Auktionen verbessern? Electronic Markets and Auctions Dagstuhl Seminar 13461 http://www.dagstuhl.de/13461

Fig. 5.1
Seminar-related press releases in 2013.

6 **Dagstuhl Publishing** *Dagstuhl Publishing*

Portfolio

6.1

Portfolio

Die Open-Access-Verlagsdienstleistungen von Schloss Dagstuhl werden in der Wissenschaftsgemeinde gut aufgenommen. Im Portfolio des Angebots gibt es zum einen Publikationsserien, die sich auf Veranstaltungen beziehen, die auf Schloss Dagstuhl abgehalten wurden (*Dagstuhl Reports*, *Dagstuhl Manifestos*, *Dagstuhl Follow-Ups*), zum anderen Serien, die Konferenzen und Workshops außerhalb von Schloss Dagstuhl bedienen. Zudem wird seit 2013 die wissenschaftliche Zeitschrift *LITES* veröffentlicht.

■ Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*¹⁹ dokumentiert, was eine Zitation der Seminare im wissenschaftlichen Kontext ermöglicht. Zudem erlaubt es auch denjenigen Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Seminaren und Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 12.5) agiert als Herausbergremium für die Reihe. Um umfassende Zusammenstellungen von begutachteten Artikeln auf Basis eines Dagstuhl-Seminars oder -Perspektiven-Workshops zu ermöglichen, wurde die Buchreihe *Dagstuhl Follow-Ups* (siehe unten) gegründet.

In 2013 wurden für alle 75 Seminare und Perspektiven-Workshops ein Bericht in der Reihe *Dagstuhl Reports* veröffentlicht. An dieser Stelle bedanken wir uns ganz herzlich bei den Organisatoren und Kollektoren für die erfolgreiche Zusammenarbeit.

■ Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*²⁰ die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 12.5) fungiert hier ebenfalls als Herausbergremium. Die Ausgabe für 2013 enthält zwei Manifestos, siehe Fig. 6.1.

■ Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*²¹ ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert. Für jedes Buch ist ein gesonderter Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausbergremium verantwortlich ist) begutachtet und freigegeben werden muss. Die Bücher erscheinen nur online. In 2013 konnten drei Bücher veröffentlicht werden, siehe Fig. 6.2.

The scientific community appreciates the Open Access publishing services offered by Schloss Dagstuhl. The portfolio covers series related to events at Schloss Dagstuhl (*Dagstuhl Reports*, *Dagstuhl Manifestos*, *Dagstuhl Follow-Ups*) and series for conferences and workshops held outside of Schloss Dagstuhl (*OASICs* and *LIPICs*). The scholarly journal *LITES* has been running since 2013, when it was launched.

■ Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*¹⁹ which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

The periodical started with the first seminars of January 2011 and publishes in monthly issues reports on seminars and workshops that took place on a given month. The content is not peer-reviewed. The Scientific Directorate (see Fig. 12.5) acts as editorial board. For comprehensive collections of peer-reviewed articles developed on the basis of a Dagstuhl Seminar or Perspectives Workshop, we offer seminar organizers the possibility of publishing a volume in our book series *Dagstuhl Follow-Ups* (see below).

All of the 75 Dagstuhl Seminars and Dagstuhl Perspectives Workshops that took place in 2013 have published a report. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

■ Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*²⁰ in Open Access manner. The Scientific Directorate (see Fig. 12.5) acts as the editorial board of the journal. The 2013 volume includes two Dagstuhl Manifestos; see Fig. 6.1.

■ Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*²¹ book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Perspectives Workshop. Each book needs a separate proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). The books are published online only. In 2013, three volumes were published; see Fig. 6.2.

¹⁹ <http://drops.dagstuhl.de/dagrep>

²⁰ <http://drops.dagstuhl.de/dagman>

²¹ <http://drops.dagstuhl.de/dfu>

■ OASlcs: OpenAccess Series in Informatics

Die *OASlcs*-Reihe²² veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausbergremium, siehe Fig. 6.3, diskutiert sorgfältig alle eingehenden Anträge, um ausschließlich qualitativ hochwertige sowie professionell durchgeführte Veranstaltungen in die Reihe aufzunehmen und um gegebenenfalls Empfehlungen zur Verbesserung der Veranstaltungsstruktur zu geben.

Es wurden sieben Bände in 2013 von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 6.4.

■ LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe²³ veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausbergremium besteht aus einschlägig bekannten Wissenschaftlern und wird von Pascal Weil als Hauptherausgeber geleitet, siehe Fig. 6.5. In 2013 wurden Tagungsbände von sechs großen Konferenzen veröffentlicht: CSL, FSTTCS, RTA, STACS – und erstmalig – TYPES und TQC. Die FSTTCS-Konferenz wurde in 2013 nach fünfjähriger Laufzeit erneut evaluiert und für weitere fünf Jahre (2014–2018) akzeptiert. Zudem wurde ein Antrag der beiden Konferenzen APPROX/RANDOM (welche als gemeinsame Veranstaltung durchgeführt werden) angenommen für den Veröffentlichungszeitraum 2014 bis 2018.

Im Herausbergremium ist Vinay Viswanathan ausgeschieden. Catuscia Palamidessi wurde als neues Mitglied gewählt.

■ OASlcs: OpenAccess Series in Informatics

The *OASlcs* series²² aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 6.3, discusses carefully all submitted proposals to ensure that only significant and professionally organized events are added to the series and that – if applicable – suggestions are given for improving the structure of the event.

In 2013, Dagstuhl published seven *OASlcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 6.4.

■ LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series²³ publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers supervises the conferences that are accepted for LIPIcs; Pascal Weil acts as editor-in-chief. See also Fig. 6.5. In 2013, the series published the proceedings of six major conferences: CSL, FSTTCS, RTA, STACS, and – for the first time – TYPES and TQC. FSTTCS was re-evaluated by the LIPIcs editorial board and accepted for another five-year period (2014–2018). Furthermore, a proposal of the two conferences APPROX/RANDOM (organized as a single event) was accepted for the years 2014 to 2018.

Vinay Viswanathan retired from the editorial board. As a new member, Catuscia Palamidessi was elected.

²² <http://drops.dagstuhl.de/oasics>

²³ <http://drops.dagstuhl.de/lipics>

Machine Learning Methods for Computer Security
<http://dx.doi.org/10.4230/DagMan.3.1.1>
 based on Dagstuhl Perspectives Workshop 12371

ICT for Bridging Biology and Medicine
<http://dx.doi.org/10.4230/DagMan.3.1.31>
 based on Dagstuhl Perspectives Workshop 13342

Fig. 6.1

Manifestos published in 2013 in the journal *Dagstuhl Manifestos*.

Vol. 4 | Normative Multi-Agent Systems | Editors: Giulia Andrichetto, Guido Governatori, Pablo Noriega, Leendert W. N. van der Torre
<http://www.dagstuhl.de/dagpub/978-3-939897-51-4>

Vol. 5 | Data Exchange, Integration, and Streams | Editors: Phokion G. Kolaitis, Maurizio Lenzerini, Nicole Schweikardt
<http://www.dagstuhl.de/dagpub/978-3-939897-61-3>

Vol. 6 | Artificial and Computational Intelligence in Games | Editors: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, Julian Togelius
<http://www.dagstuhl.de/dagpub/978-3-939897-56-9>

Fig. 6.2

Dagstuhl Follow-Ups volumes published in 2013.

■ LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*²⁴ veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die sich für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 6.7), begutachtet alle eingereichten Arbeiten.

Im Gegensatz zu anderen Zeitschriften im Bereich eingebetteter Systeme, steht bei *LITES* eine moderate Veröffentlichungsgebühr (article-processing charge, APC) sowie ein schnelles Begutachtungsverfahren (innerhalb eines Jahres ab Einreichung) im Vordergrund. Die APC von 100€ ist momentan für den Zeitraum 2013–2015 sichergestellt Dank finanzieller Unterstützung von Google und der Klaus Tschira Stiftung.

In 2013 wurden bereits erste Beiträge eingereicht und begutachtet. Die Veröffentlichung dieser Artikel ist für das Frühjahr 2014 geplant.

■ LITES: Leibniz Transactions on Embedded Systems

The *LITES*²⁴ journal publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was established in 2012 and started operating in early 2013. A broad team of experienced researchers, acting as editorial board (see Fig. 6.7), reviews all submitted contributions.

In contrast to existing journals on embedded computer systems, *LITES* charges only a moderate article-processing charge (APC) and aims at efficient reviewing procedures to ensure that articles are published within one year of submission. The APC of 100€ is guaranteed for the 2013–2015 period thanks to support from sponsors like Google and the Klaus Tschira Stiftung.

LITES received several submissions in 2013, which were immediately reviewed. The publication of these articles is planned for early 2014.

²⁴ <http://drops.dagstuhl.de/lites>

Prof. Dr. Daniel Cremers TU Munich, Germany
Prof. Dr. Barbara Hammer Bielefeld University, Germany
Prof. Dr. Marc Langheinrich University of Lugano, Switzerland
Prof. Dr. Dorothea Wagner Karlsruhe Institute of Technology, Germany Editor-in-Chief

Fig. 6.3
OASlcs Editorial Board.

Vol. 29 SLATE'13 2 nd Symposium on Languages, Applications and Technologies http://www.dagstuhl.de/dagpub/978-3-939897-52-1
Vol. 30 WCET'13 13 th International Workshop on Worst-Case Execution Time Analysis http://www.dagstuhl.de/dagpub/978-3-939897-54-5
Vol. 31 FSFMA'13 1 st French Singaporean Workshop on Formal Methods and Applications http://www.dagstuhl.de/dagpub/978-3-939897-56-9
Vol. 32 CMN'13 2013 Workshop on Computational Models of Narrative http://www.dagstuhl.de/dagpub/978-3-939897-57-6
Vol. 33 ATMOS'13 13 th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems http://www.dagstuhl.de/dagpub/978-3-939897-58-3
Vol. 34 GCB'13 German Conference on Bioinformatics 2013 http://www.dagstuhl.de/dagpub/978-3-939897-59-0
Vol. 35 ICCSW'13 2013 Imperial College Computing Student Workshop http://www.dagstuhl.de/dagpub/978-3-939897-63-7

Fig. 6.4
OASlcs volumes published in 2013.

Infrastruktur

6.2

Infrastructure

6

■ Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 6.8. Die Bände aus der Reihe *LIPICs* werden beim Conference Proceedings Citation Index (CPCI), welcher vom Medienkonzern Thomson Reuters unterhalten wird, eingereicht; zudem werden diese seitens SCOPUS in deren Katalog aufgenommen. Die Reihen *LIPICs* und *OASICs* sind zudem im Directory of Open Access Journals (DOAJ) gelistet, siehe Fig. 6.8.

■ LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*²⁵ ein Online-Repositorium ins Leben gerufen, um Open Access-

■ Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 6.8. The *LIPICs* volumes are submitted to the Conference Proceedings Citation Index (CPCI), maintained by the Thomson Reuters media group; additionally, SCOPUS is integrating them into their catalog. The *LIPICs* and *OASICs* series are also listed in the Directory of Open Access Journals (DOAJ), see Fig. 6.8

■ LeibnizOpen

The Leibniz Association has established the *Leibniz-Open*²⁵ repository to promote the open-access publica-

Prof. Dr. Susanne Albers Technical University Munich, Germany
Prof. Dr. Chris Hankin Imperial College London, United Kingdom
Prof. Deepak Kapur PhD University of New Mexico, US
Prof. Michael Mitzenmacher PhD Harvard University, US
Prof. Madhavan Mukund PhD Chennai Mathematical Institute, India
Dr. Catuscia Palamidessi INRIA, France <i>tenure started on June 1, 2013</i>
Prof. Dr. Wolfgang Thomas RWTH Aachen, Germany
Dr. Vinay Viswanathan Chennai Mathematical Institute, India LimberLink Technologies Pvt. Ltd. <i>tenure ended on May 31, 2013</i>
Pascal Weil PhD CNRS, France University Bordeaux, France Editor-in-Chief
Prof. Dr. Dr. h. c. Reinhard Wilhelm Saarland University, Germany Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Germany

Fig. 6.5
LIPICs Editorial Board.

Vol. 19 TYPES'11 18 th International Workshop on Types for Proofs and Programs http://www.dagstuhl.de/dagpub/978-3-939897-49-1
Vol. 20 STACS'13 30 th International Symposium on Theoretical Aspects of Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-50-7
Vol. 21 RTA'13 24 th International Conference on Rewriting Techniques and Applications http://www.dagstuhl.de/dagpub/978-3-939897-53-8
Vol. 22 TQC'13 8 th Conference on the Theory of Quantum Computation, Communication and Cryptography http://www.dagstuhl.de/dagpub/978-3-939897-55-2
Vol. 23 CSL'13 Computer Science Logic 2013 http://www.dagstuhl.de/dagpub/978-3-939897-60-6
Vol. 24 FSTTCS'13 IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-64-4

Fig. 6.6
LIPICs volumes published in 2013.

Veröffentlichungen von Leibniz-Instituten und deren Wissenschaftlern zu unterstützen und sichtbar zu machen. Schloss Dagstuhl liefert alle Artikel aus den Reihen *Dagstuhl Reports* und *Dagstuhl Manifestos* an das Repository und stärkt dadurch Forschungsergebnisse aus der Informatik innerhalb dieses multidisziplinären Repositoriums.

■ AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Potentiale durch Open Access!“ initiiert und koordiniert. Der Workshop fand am 31. Januar und 1. Februar 2013 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt. Der Workshop wurde von ca. 45 Teilnehmer aus den Verlagsabteilungen von ungefähr 20

tions of Leibniz institutes and their researchers. Schloss Dagstuhl submits all articles from the *Dagstuhl Reports* and *Dagstuhl Manifestos* series to the repository, thereby strengthening informatics-related research in this multi-disciplinary repository.

■ Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Potentiale durch Open Access!” was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop took place at the Leibniz Association headquarters in Berlin from January 31 to February 1, 2013, and brought together approximately 45 professionals in charge of publishing activities at about 20 Leibniz institutes. For 2014 a

²⁵ <http://www.leibnizopen.de/>

Prof. Dr. Alan Burns University of York, UK Editor-in-Chief
Prof. Sang Lyul Min PhD Seoul National University, South Korea Subject area: Architecture, platforms
Prof. Dr. Marco di Natale Scuola Superiore Santa Anna, Italy Subject area: Automotive applications
Dr. Virginie Wiels ONERA, France Subject area: Avionics applications
Prof. Karl-Erik Arzen PhD Lund University, Sweden Subject area: Control
Prof. Steve Goddard PhD University of Nebraska-Lincoln, US Subject area: Cyber-physical systems
Prof. Dr. Axel Jantsch Royal Institute of Technology Stockholm, Sweden Subject area: Distributed embedded systems and networks
Prof. Bashir Al Hashimi University of Southampton, UK Subject area: Energy-efficiency
Prof. Mateo Valero PhD Technical University of Catalonia Subject area: High-performance embedded systems
Prof. Dr. Martin Fränzle Carl von Ossietzky University Oldenburg, Germany Subject area: Hybrid systems
Prof. Dr. Samarjit Chakraborty Technical University Munich, Germany Subject area: Multimedia applications
Prof. Dr. Gernot Heiser University of New South Wales, Australia Subject area: Operating systems
Prof. Dr. Lothar Thiele ETH Zürich, Switzerland Subject area: Performance and wireless sensor networks
Dr. Neil Audsley University of York, UK Subject area: Real time
Prof. Sanjoy Baruah PhD University of North Carolina at Chapel Hill, US Subject area: Scheduling
Prof. Dr. Florence Maraninchi University of Grenoble, France Verimag Lab, France Subject area: Verification, formal methods, model-based design

Fig. 6.7

LITES Editorial Board.

Leibniz-Instituten wahrgenommen. Für 2014 ist eine Nachfolgeveranstaltung mit dem Titel „Erfolgreiches Journal-Management: Qualität und Reputation“ in Planung.

■ AG Open Access der Schwerpunktinitiative „Digitale Information“

Die Allianz der deutschen Wissenschaftsorganisationen, zu der neben der Max-Planck-Gesellschaft, der Helmholtz-Gemeinschaft, sowie weiteren Organisationen auch die Leibniz-Gemeinschaft gehört, hat eine Schwerpunktinitiative „Digitale Information“ ins Leben gerufen, bei der auch das Thema *Open Access* als Handlungsfeld vertreten ist. Mit Dr. Marc Herbstritt wurde seitens der Leibniz-Gemeinschaft ab Juli 2013 ein Mitglied des wissenschaftlichen Stabs von Schloss Dagstuhl in die Arbeitsgruppe „Open Access“²⁶ berufen

Die Mitarbeit in dieser Arbeitsgruppe erlaubt, Anforderungen aus dem Wissenschaftsumfeld der Informatik auf politischer Ebene einzubringen. Zudem erleichtert es den Austausch und die Abstimmung fortlaufender Prozesse vor dem Hintergrund der weiterhin dynamischen Umgestaltung der Publikationslandschaft hin zu Open Access.

■ Technisches Back-end: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)²⁷ werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative²⁸ berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer grossen Leserschaft zur Verfügung. Als technische Grundlage dient eine adaptierte Version des OPUS-Systems.²⁹

²⁶ http://www.allianzinitiative.de/de/handlungsfelder/open_access/

²⁷ <http://www.dagstuhl.de/drops>

follow-up workshop entitled “Erfolgreiches Journal-Management: Qualität und Reputation” is planned.

■ Open Access Working Group of the Priority Initiative “Digital Information”

The Alliance of German Science Organizations, to which – among others – the Max Planck Society, the Helmholtz Association and also the Leibniz Association belong to, has established a priority initiative “Digital Information” where *Open Access* is handled as a core activity. Since July 2013, Dagstuhl scientific staff member Dr. Marc Herbstritt has collaborated with this working group as the delegated representative of the Leibniz Association.²⁶

Such collaboration offers an opportunity to highlight the scientific requirements of the computer science discipline on a political level. Additionally, it enables and simplifies the exchange and calibration of ongoing changes in the publishing landscape towards Open Access.

■ Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)²⁷. The general guidelines of the Dublin Core initiative²⁸ applicable to online publications are adhered to, meaning that all the requisite metadata of each publication is stored, thus ensuring availability in the long term. This enables the online publications to be cited by and accessible to a wide readership. The technical basis for this is an adapted version of the OPUS system.²⁹

DBLP
Dagstuhl Reports http://dblp.uni-trier.de/db/journals/dagstuhl-reports/
Dagstuhl Manifestos http://dblp.uni-trier.de/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups http://dblp.uni-trier.de/db/series/dfu/
OASlcs http://dblp.uni-trier.de/db/series/oasics/
LIPlcs http://dblp.uni-trier.de/db/series/lipics/
DOAJ
OASlcs http://doaj.org/toc/759bc28bcc174e25a1c571e9e29f9632
LIPlcs http://doaj.org/toc/160b1ba80f8a46278ac8c92722c898c6

Fig. 6.8
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

■ Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)³⁰ zur (digitalen) Langzeitarchivierung eingereicht.

■ Mirroring

Um dem Verlust von Daten vorzubeugen, werden seit 2010 zwei Kooperationen zur Spiegelung (Mirroring) von Inhalten des Publikationsservers DROPS gepflegt:

- io-port.net: Das unter Leitung des FIZ Karlsruhe, Leibniz-Institut für Informationsinfrastruktur, organisierte Informatik-Publikations-Portal io-port.net spiegelt alle Bände der LIPIcs-Reihe.³¹ In 2011 wurde die bestehende Verbindung durch eine gemeinsame Kooperationserklärung gefestigt.
- SunSite Central Europe: Der Sun-Server-Park, der an der RWTH Aachen unter Leitung von Prof. Dr. Matthias Jarke betrieben wird, bietet eine Heimat für zahlreiche Software-Archive als auch Publikationen. Der gesamte DROPS-Bestand wird nun in regelmäßigen Abständen auf der SunSite Aachen gespiegelt.³²

■ Long-term Archiving

All publications are submitted to the German National Library (D-NB)³⁰ for (digital) long-term archiving.

■ Mirroring

In order to prevent data loss, two cooperative ventures were initiated in 2010 for mirroring the content of the DROPS publication server:

- io-port.net: The informatics publication portal organized under the auspices of io-port.net, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, mirrors all volumes of the LIPIcs series³¹. In 2011, the existing affiliation was consolidated by a memorandum of understanding.
- SunSite Central Europe: The Sun server park, located at the Aachen University of Technology and operated under the guidance of Prof. Dr. Matthias Jarke, is home to numerous software archives and publications. All the DROPS assets are now mirrored at regular intervals on the Aachen SunSite.³²

²⁸ <http://dublincore.org/>

²⁹ <http://elib.uni-stuttgart.de/opus/doku/about.php>

³⁰ http://www.dnb.de/DE/Netzpublikationen/Langzeitarchivierung/langzeitarchivierung_node.html

³¹ <http://www.io-port.net> (→ Digital Library → LIPIcs)

³² <http://vesta.informatik.rwth-aachen.de/Dagstuhl/>

7 Bibliographiedatenbank dblp *dblp computer science bibliography*

Schloss Dagstuhl und dblp

7.1

Schloss Dagstuhl and dblp

Die Zusammenarbeit zwischen Schloss Dagstuhl und der an der Universität Trier entwickelten Bibliographiedatenbank dblp besteht bereits seit Ende 2010. Ursprünglich durch ein Projekt im Leibniz-Wettbewerb gefördert wird das Engagement von Schloss Dagstuhl für dblp seit Juni 2013 von Dagstuhl direkt finanziert. Die Finanzierung wird zudem bereits seit November 2010 durch eine großzügige Spende der Klaus-Tschira-Stiftung unterstützt. Bereits seit 2012 steht nun auch unter dblp.dagstuhl.de ein eigener dblp-Webservice unter der Domain von Schloss Dagstuhl bereit und ergänzt damit das dblp-Angebot der Universität Trier unter dblp.uni-trier.de. Das Kooperationsabkommen zwischen Schloss Dagstuhl und der Universität Trier wurde 2013 um zunächst weitere drei Jahre verlängert.

Im Zuge der Konsolidierung der Zusammenarbeit wurden unter dem Dach von Schloss Dagstuhl zweieinhalb Mitarbeiterstellen im wissenschaftlichen Stab geschaffen, die hauptamtlich für die Betreuung und Weiterentwicklung von dblp abgestellt sind. Schon jetzt können sich die Erfolge des Engagement von Schloss Dagstuhl sehen lassen: Seit Beginn der Zusammenarbeit Anfang November 2010 ist der Datenbestand von dblp um über eine Millionen neuer Publikationseinträge gewachsen; das entspricht einem Wachstum der Datenbasis um über 67 % in nur etwas über drei Jahren. Die ursprüngliche Aufnahmequote von etwa 100 000 neuen Publikationseinträgen pro Jahr konnte auf über 325 000 Publikationseinträge pro Jahr gesteigert werden und hat damit bereits alle initialen Erwartungen übertroffen. Ein dblp-Beirat zur wissenschaftliche Aufsicht, bestehend aus namhaften Persönlichkeiten aus den

The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – exists since end of 2010. The commitment of Schloss Dagstuhl to dblp was initially funded by a project of the Leibniz Competition. Since June 2013, it has been funded directly by Schloss Dagstuhl. Since November 2010, Schloss Dagstuhl’s dblp team has also been supported by a generous donation from the Klaus Tschira Foundation. Schloss Dagstuhl’s own dblp web service at dblp.dagstuhl.de was established in 2012 and complements the dblp service available at the University of Trier at dblp.uni-trier.de. In 2013, the cooperation agreement between Schloss Dagstuhl and the University of Trier was renewed for another three years.

As part of the consolidation of this cooperation, two and a half staff positions – assigned full-time to the support and development of dblp – were created in the scientific staff of Schloss Dagstuhl. Even after this short period of time, the results of Schloss Dagstuhl’s commitment to dblp are already quite impressive: the dblp data set has grown by over one million new publication records since the beginning of cooperation in early November 2010. This corresponds to a total growth of the database by over 67 % in just over three years. The original data acquisition rate of about 100,000 new publication entries per year has increased to more than 325,000 publication entries per year, thus already exceeding all initial expectations. A dblp advisory board, consisting of distinguished researchers from the various disciplines of computer science, has been established and has since supported dblp with its expertise.

dblp-Beirat dblp Advisory Board
Prof. Dr. Hannah Bast University of Freiburg, Germany Chair
Prof. Dr. Andreas Butz Ludwig Maximilians University Munich, Germany
Prof. Dr.-Ing. Rüdiger Dillmann Karlsruhe Institute of Technology, Germany
Prof. Oliver Günther, Ph.D. University of Potsdam, Germany
Prof. Dr. Hans-Peter Lenhof Saarland University, Germany
Prof. Dr. Mila Majster-Cederbaum University of Mannheim, Germany
Prof. Dr. Dietmar Saupe University of Konstanz, Germany
Prof. Dr. Dr. h.c. Otto Spaniol RWTH Aachen, Germany
Prof. Dr.-Ing. Jürgen Teich University of Erlangen-Nuremberg, Germany
Prof. Dr. Dr. h.c. Reinhard Wilhelm Saarland University, Germany

Fig. 7.1
dblp Advisory Board.

verschiedenen Disziplinen der Informatik, wurde etabliert und unterstützt seitdem dblp mit seiner Expertise.

Gemessen über alle dblp-Server wurde der Webserver 2013 monatlich von ca. 100 000 individuellen Anwendern aus aller Welt genutzt, die täglich über eine Millionen Seiten abrufen. Damit ist dblp einer der weltweit wichtigsten bibliographischen Dienste für wissenschaftliche Literatur in der Informatik. Ziel von Schloss Dagstuhl ist es, dblp als das zentrale, an den Bedürfnissen der Forschenden und an Forschung Interessierten ausgerichtete Werkzeug zum Nachweis, zur Dokumentation und zur Recherche von Publikationseinheiten in der Informatik zu verstetigen und weiter zu entwickeln.

Fortentwicklung der Datenbasis

7.2

Die Literaturdatenbank dblp indiziert Publikation an Hand vollständiger Inhaltsverzeichnisse von Konferenzbänden oder Journalausgaben. Mit Hilfe einer eigens entwickelten Software zur Datenextraktion werden Metadaten von Verlagswebseiten ausgelesen und zur weiteren Bearbeitung vorbereitet. Die Metadaten werden anschließend vom dblp-Team redaktionell bearbeitet: Eventuelle Fehler werden korrigiert, mehrdeutige und ungenaue Angaben werden verbessert. Diese Datenpflege wird zwar von Hilfssoftware unterstützt, erfolgt aber vornehmlich händisch durch den jeweiligen Mitarbeiter.

Neben den selbst von Verlagswebseiten akquirierten Daten bezieht dblp auch Datenlieferungen direkt von einigen wichtigen Verlagen in der Informatik. Gegen Ende 2013 bestanden solche bilateralen Datenabkommen mit Springer Science+Business Media, dem Institute of Electrical and Electronics Engineers (IEEE), IOS Press und natürlich Dagstuhl Publishing. Seit Frühjahr 2014 ist zudem die USENIX Association als Partner hinzu gekommen.

Im Jahr 2013 wurde die Datenbank auf diese Weise um 326 009 Publikationseinträge erweitert und indizierte gegen Ende 2013 inzwischen 2 487 077 Datensätze. 133 neue Journale und 214 neue Konferenzserien wurden in die Datenbank aufgenommen. Diese Aufnahmequote verstetigt damit das bereits hohe Niveau des Vorjahres.

Einen Überblick über die Entwicklung der Datenakquise kann Fig. 7.2 und Fig. 7.3 entnommen werden.

Technische Weiterentwicklung

7.3

■ Weiterentwicklung der Daten-Wrapper

Die von dblp eigens entwickelten Software zur Datenextraktion von den Verlagswebseiten, die so genannten *Daten-Wrapper*, wurden auch 2013 kontinuierlich weiterentwickelt. Im Laufe des Jahres wurde die Software an über 40 weitere Websysteme angepasst und konnte somit gegen Ende 2013 die Metadaten von insgesamt 124 verschiedenen Systemen erfassen. Vor allem die Akquise von

Measured over all dblp mirror sites, the dblp web service is used by about 100,000 individual users from all over the world, requesting more than one million dblp pages per month. Hence, dblp is one of the world's leading bibliographic services for scientific literature in computer science. The goal of the commitment of Schloss Dagstuhl is to consolidate and improve dblp as the central tool for reference, documentation, and exploration of the research publications in computer science, which is oriented towards the needs of the computer science research community.

Development of the Data Stock

The dblp computer science bibliography indexes publications on the basis of complete tables of contents of conference proceedings and journal volumes. Using dblp's own data harvesting software, bibliographic metadata is extracted from a publisher's website and prepared as intermediate data packages for further processing. This intermediate data is manually curated by the members of the dblp team: eventual errors in the metadata are corrected, and ambiguous or incomplete data items are improved. This curation step is assisted by auxiliary software scripts, but all decisions are made intellectually by the team member.

Additionally, dblp receives data deliveries from a number of major publishers of computer science publications. By end of 2013, bilateral data delivery agreements have been established with Springer Science+Business Media, the Institute of Electrical and Electronics Engineers (IEEE), IOS Press, and Dagstuhl Publishing. Since early 2014, the USENIX Association has also joined as a data delivery partner.

In 2013, a total of 326,009 new publication records were added to dblp. Hence, by end of 2013 there were a total number of 2,487,077 publication records. 133 new journals and 214 new conference series were added to the bibliography. These figures mirror the already high number of new data items from 2012.

Fig. 7.2 and Fig. 7.3 provide an overview of the development of dblp's data acquisition.

Technical Improvements

■ Improvement of the data wrappers

In order to extract bibliographic metadata from publisher websites, dblp has developed its own data harvesting software, the so-called data wrappers. These data wrappers have been continuously improved and expanded through the course of 2013 by adopting the data wrappers to over 40 new publisher sites, up to a total of 124 website systems by the end of 2013. One focus of 2013 was the automated

Metadaten mittlerer und kleinerer Verlage stand dabei 2013 verstärkt im Vordergrund. Einen Schwerpunkt bildeten dabei die Metadaten von Open Access-Schnittstellen und gängigen OpenSource-Journalmanagementsystemen (insb. dem Open Journal System, OJS).

In Zusammenarbeit mit dem DFG-geförderten Projekt „Smart Harvesting“³³ von Universität Trier und GESIS – Leibniz Institut für Sozialwissenschaften wurde zudem die Entwicklung eines Daten-Wrappers vorangebracht, der nicht mehr rein regelbasiert arbeitet. Das Wrapper-Konzept soll zudem vom Anwendungsgebiet der Verlagswebseiten auf die automatisierte Datenextraktion aus inhomogenen PDF-Inhaltsverzeichnissen ausgeweitet werden. Die verbesserten Wrapper befinden sich derzeit noch in der Prototyp-Phase und sind noch nicht im Produktiveinsatz.

extraction of metadata from medium and smaller sized publishers. In particular, open access data interfaces and open source journal management systems (such as the Open Journal System, OJS) received an increased amount of attention.

In corporation with the DFG funded project “Smart Harvesting”³³ of the University of Trier and GESIS – Leibniz-Institute for the Social Sciences, a new concept for data wrappers is explored that do not rely on a solely rule-based logic. In addition, the concept of wrappers is expanded to other data sources, such as extracting metadata from inhomogeneous PDF files. The new generation of data wrappers is currently in a prototype stage and they are not yet ready for use in a live production environment.

³³ *Smart Harvesting: Verbesserung des Open Access-Zugangs; Erhöhung der Qualität der Metadaten*, DFG Fördernummer WA 1267/2-1
Smart Harvesting: Better Open Access; Improvement of Metadata Quality, DFG grant WA 1267/2-1

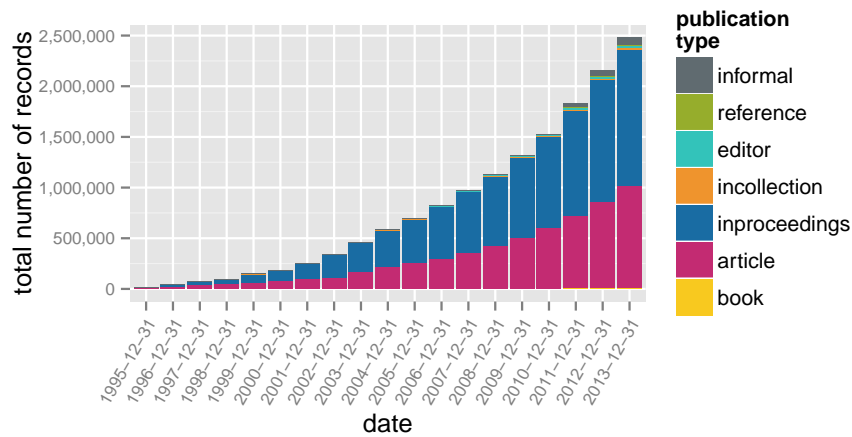


Fig. 7.2 Total number of records in dblp by year and type.

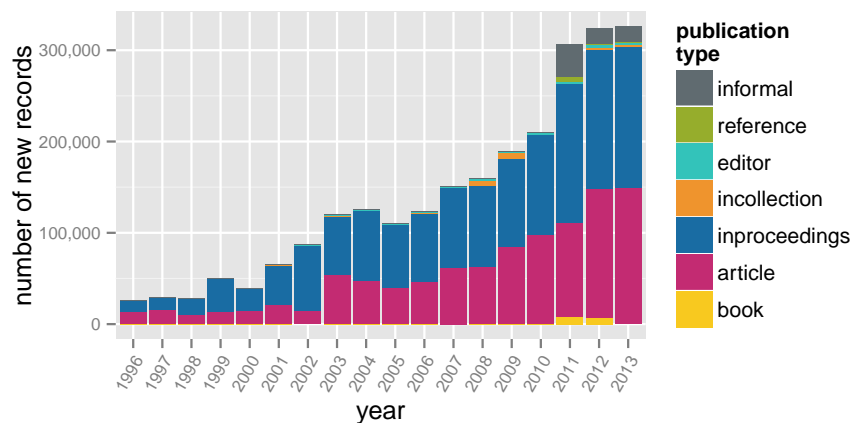


Fig. 7.3 New records added to dblp by year and type.

■ Überarbeitung des dblp-Webdienstes

Als reines Online-Angebot ist die Funktionalität und die Bedienbarkeit des Webservices für dblp von zentraler Bedeutung. Daher stellte die Verbesserung und Weiterentwicklung des dblp-Webservices einen wesentlichen Punkt der Entwicklung von dblp dar.

Nachdem bereits im Jahr 2012 das graphische Layout der Nutzerschnittstelle grundlegend überarbeitet und an die Möglichkeiten eines zeitgemäßen CSS/HTML5-Layoutes angepasst wurde, wurde 2013 verstärkt an der Modernisierung der Funktionalitäten gearbeitet. So wurde unter anderem in Zusammenarbeit mit Prof. Dr. Hannah Bast (Universität Freiburg) die bisher extern betriebenen *CompleteSearch*-Suchmaschine als Modul in den dblp-Service integriert und erlaubt so die Verbesserung und Erweiterung bestehender Such- und Filteroptionen. Die Integration basiert dabei auf einer neuen, Service-orientierten Architektur, die sich Quer durch alle Komponenten des dblp-Service zieht und insbesondere auf zukünftige Erweiterbarkeit ausgerichtet ist.

In gleicher Weise wie die graphische Nutzerschnittstelle neu ausgerichtet wurde, wurden auch die von dblp angebotenen API-Services grundlegend überarbeitet und um wesentliche Komponenten wie BibTeX-Listenexporte, RSS-Feeds und DOI- bzw. ISBN-Lookups erweitert. Des Weiteren wurde mit der Etablierung eines RDF/XML-basierten Datensicht die Anbindung an das semantische Web mit der Bereitstellung von Linked Open Data in die Wege geleitet. Mit Hilfe dieser neuen Infrastruktur können die dblp-Daten nun mit den Werkzeugen des semantischen Webs (z. B. SPARQL-Browsern) erschlossen werden.

Nach Abschluss der letzten Testphase wird die neue Service-Architektur Anfang 2014 zunächst auf dem Server dblp.dagstuhl.de sowie später auch auf den Servern in Trier schrittweise in Betrieb gehen.

Erhebung zu wissenschaftlichen Publikationsorganen in der deutschen Informatik

7.4

2013 wurde von Schloss Dagstuhl in Zusammenarbeit mit der Gesellschaft für Informatik e.V. (GI) und dem Fakultätentag Informatik (FTI) eine groß angelegte Studie über die Veröffentlichungspraxis der deutschen Informatikerinnen und Informatikern durchgeführt. Alle Mitglieder der GI waren eingeladen, die Journale und Konferenzreihen, in denen sie zuletzt publiziert haben, zu benennen sowie bezüglich wissenschaftlicher Qualitätsaspekte und thematischer Ausrichtung einzuschätzen. Die Ergebnisse der Umfrage werden helfen, eine Einschätzung internationaler Informatik-Konferenzen und Fachzeitschriften aus deutscher Perspektive zu entwickeln und somit an den Bedürfnissen der Forschenden ausgerichtete Prioritäten für die zukünftige Fortentwicklung der dblp-Datenbasis zu setzen. Ein weiteres Ziel ist es, eine genauere Übersicht über die Abdeckung der verschiedenen Informatik-Teildis-

■ Redesign of the dblp web service

Since dblp is an online-only service, the utility and accessibility of the dblp website is of crucial importance. In 2013, the improvement of the dblp web infrastructure was a main focus of the dblp team.

The graphical layout of the user interface has already been revised and adapted to the possibilities of a contemporary CSS/HTML5-Layouts in 2012. In 2013, the focus was on improving the functionalities of the web service. In corporation with Prof. Dr. Hannah Bast (University of Freiburg), the formerly externally operated search engine *CompleteSearch* was integrated as a module to the dblp service. The new search engine allows for the improvement and expansion of existing search and filter options. This integration is based on a new service-oriented architecture of the dblp service that is geared towards expandability.

At the same time, the API services offered by dblp have been thoroughly revised and expanded to include essential components, such as BibTeX list exports, RSS feeds, and DOI or ISBN look-ups. Furthermore, an RDF/XML-based data view was implemented to allow dblp to be directly linked to the semantic web by providing linked open data. This new infrastructure will allow dblp data to be explored using semantic web tools (e.g., using a SPARQL browser).

After a final test phase, the new service architecture will be available on dblp.dagstuhl.de in early 2014. Later on, the new service will also be available in Trier.

Survey on Scientific Publication Venues in German Computer Science

In 2013, a large-scale survey on the publication practices among German computer scientist was conducted by Schloss Dagstuhl in cooperation with the Gesellschaft für Informatik e.V. (GI) and the Fakultätentag Informatik (FTI). All members of the GI had been invited to give their assessment on a number of scientific quality aspects and the thematic focus of the recent publication venues they published in. This survey will help to develop a rating of international computer science conferences and journals from a German perspective, and to assign priorities for possible further data inclusions to dblp according to the needs of the computer science research community. Another goal is to understand the coverage of dblp within the different computer science sub-disciplines. This survey, a main focus of the dblp team in 2013, was supervised by an editorial board (see Fig. 7.4).

ziplinen innerhalb von dblp zu gewinnen. Die Umfrage war im Jahr 2013 ein Schwerpunkt der Arbeit des dblp-Teams. Die Umfrage wurde von einem Editorial Board wissenschaftlich beaufsichtigt (siehe Fig. 7.4).

Die Erhebung wurde als freiwillige Online-Umfrage³⁴ durchgeführt. Dazu wurden 18 034 GI-Mitglieder per Post persönlich eingeladen und mit eindeutigen, anonymisierten Zugangsdaten ausgestattet. Insgesamt 2 356 der eingeladenen Mitglieder haben an der Umfrage teilgenommen. Von diesen Personen haben 2 009 Befragte die Umfrage in vollem Umfang bearbeitet; 347 Fragesätze wurden nur unvollständig beantwortet und mussten daher bei der Bewertung aussortiert werden. Die Rücklaufquote an vollständig erfassten Befragungen liegt damit bei etwa 11 %.

Im Rahmen der Befragung wurden insgesamt 1 421 verschiedene Konferenzserien und 844 verschiedene Journale angegeben. Dies entspricht in etwa dem Verhältnis zwischen Konferenzen und Journalen, welches auch im dblp-Datensatz zu beobachten ist. Eine detaillierte Aufschlüsselung der gegebenen Reihen zeigt, dass 71 % der genannten Konferenzen und 63 % der genannten Journale bereits in dblp gelistet waren. Eingeschränkt auf alle Reihen, die wenigstens drei Mal von verschiedenen Befragten genannt wurden, zeigt sich sogar eine Abdeckung von 94 % bzw. 97 %. Dies weist bereits auf eine offenbar hervorragende Abdeckung von dblp innerhalb der Kerninformatik hin. Eine genauere Betrachtung der Reihen, die nur von einzelnen Befragten genannt wurden, zeigt zudem, dass deren überwältigende Mehrheit nicht zum Zuständigkeitsbereich von dblp gehören.

Die angegebenen Reihen wurden zudem von den Befragten an Hand einer Reihe von wissenschaftlichen Qualitätskriterien wie Originalität und Methodik der Arbeiten, Einzigartigkeit der Reihe innerhalb der Disziplin und

The survey was implemented as a voluntary online questionnaire.³⁴ 18,034 GI members were invited to participate, and unique and anonymized login credentials were sent to each GI member by mail. A total of 2,356 invitees chose to participate in the survey. Of the answers given, 2,009 questionnaires were answered from start to finish, while the answering of the remaining 347 questionnaires had stopped at some point halfway through, and thus, those incomplete answers had to be discarded from the survey. The figure of complete answers corresponds to a response rate of approximately 11 %.

Throughout the survey, a total of 1,421 distinct conference series and 844 distinct journals were listed and evaluated by the participants. This corresponds to the ratio between conference and journal publications in computer science that can also be found in the dblp data set. The survey answers show that 71 % of all given conference series and 63 % of all journals were already indexed by dblp. However, if we limit this statistic to only consider the venues that were evaluated by three or more participants, then we obtain a coverage of 94 % and 97 %, respectively. This implies that dblp already has a very solid coverage of the central venues of computer science. A closer examination of venues that were given only once in the survey also shows that those conferences and journals, for the most part, do not belong to the scope of dblp.

Each given conference and journal was also rated by the participants with respect to an array of quality criteria such as originality, methodology, uniqueness to the community, among others. Unfortunately, on average there were only 2.8 answers given on every distinct venue, which raises doubts as to the statistical relevance of the aggregated results of the quality criteria. The editorial board is currently discussing how these answers can be utilized and

Editorial Board editorial board
Prof. Dr. Hannah Bast dblp advisory board
Prof. Dr.-Ing. Peter Liggesmeyer Gesellschaft für Informatik e.V. (GI)
Prof. Dr. Rüdiger Reischuk Fakultätentag Informatik (FTI)
Prof. Dr. Dr. h.c. Reinhard Wilhelm Schloss Dagstuhl
Dr. Michael Wagner dblp team
Beratenden Mitglieder Consultative members
Prof. Dr. Uwe Brinkschulte
Prof. Dr. Oliver Deussen
Prof. Dr. Gregor Engels
Prof. Dr.-Ing. Dr. h.c. Manfred Nagl
Prof. Dr. Erhard Rahm

Fig. 7.4

Editorial board of the survey on scientific publication venues.

anderen Aspekten beurteilt. Unglücklicherweise wurden einzelne Reihen im Durchschnitt von nur 2,8 Befragten bewertet, weswegen diese Daten in der Breite nur zu statistisch zweifelhaften Ergebnisse führen dürften. Das Editorial Board ist derzeit noch mit der Diskussion um die Nachverwendung der Ergebnisse sowie der Veröffentlichung der anonymisierten Rohdaten befasst.

interpreted, as well as if and how the anonymized raw data should be made publicly available.

Ausblick

7.5

Outlook

■ Arbeitsplan 2014

Für das Jahr 2014 stehen folgende Fortentwicklungen auf dem Arbeitsplan von dblp:

- quantitative Konsolidierung der Anzahl an jährlichen Neuaufnahmen
- teamorientierte Restrukturierung der internen Arbeitsabläufe der Datenakquise
- Aufbau von Strukturen für ein qualitätssicherndes Auswahlverfahren
- Erweiterung des dblp-Katalogs um Monographien und Dissertationen
- semantische Anreicherung der dblp-Daten um insbesondere Informationen zur Institutszugehörigkeit von Personen
- Untersuchung zur Machbarkeit ergänzender Serviceangebote für Dagstuhl-Seminare
- Vernetzung mit weiteren Infrastrukturinstituten der Leibniz-Gemeinschaft
- Gewinnung von Drittmitteln und Aktivierung alternativer Einnahmequellen

■ Zusammenarbeit von dblp, Zentralblatt MATH und HITS

Die Urheberschaft wissenschaftlicher Publikationen eindeutig zu erkennen und zuzuordnen ist eine der großen Herausforderungen bibliographischer Datendienste. Die Forschung kennt dies Problem in seiner allgemeinen Form als das Problem der „Entity-Resolution“ oder der „Autoren-namen-Disambiguierung“, und es stellt ein wichtiges Forschungsthema im Bereich der linguistischen Datenverarbeitung dar. In einem gemeinsamen Projekt wollen sich die Bibliographiedatenbank dblp, das Zentralblatt MATH des FIZ Karlsruhe und das Heidelberger Institut für Theoretische Studien (HITS) Projekt diesem Problem annehmen und mit Hilfe des aktuellen Forschungsstandes gemeinsame Lösungsstrategien entwickeln. Die Datensätze von Zentralblatt MATH und dblp teilen dabei die Probleme bei der Identifikation von Autorennamen. Die Kombination beider Datensätze, bestehend aus teils überlappenden und teils disjunkten Einträgen, stellt dabei eine interessante Möglichkeit dar, Fehler in den Datensätzen aufzudecken und von einander zu lernen. Die Natural-Lan-

■ Agenda 2014

In 2014, the following developments are planned for dblp:

- quantitative consolidation of the annual rate of new additions to dblp
- team-oriented reorganization of the internal data acquisition workflows
- establishment of a quality assuring selection process
- addition of monographs and dissertations to the dblp data stock
- semantic enrichment of dblp data, such as adding information on person's affiliations
- investigation of the possibility of additional dblp services to support the Dagstuhl seminars
- networking with other infrastructure providing institutes of the Leibniz Association
- obtaining third-party funding grants and activation of alternative sources of income

■ Cooperation of dblp, Zentralblatt MATH, and HITS

The correct attribution of scholarly material to their unambiguous authors ranks among the most critical challenges for digital libraries. More generally, the problem of determining which records in a database refer to the same entities is known as “entity resolution” or “author name disambiguation” and constitutes an important field of research within the discipline of natural language processing. In a joint project, the dblp computer science bibliography and the Zentralblatt MATH (located at FIZ Karlsruhe) aim to begin partnering with the Heidelberg Institute for Theoretical Studies (HITS) to find and implement new and state-of-the-art strategies to overcome the challenges of author identification and disambiguation. Zentralblatt MATH and dblp share the challenges associated with author name disambiguation. Due to their partially overlapping, but also partially disjointed data, a joint effort to identify authors based on the combination of the two data sets appears to be very promising. The Natural Language Processing (NLP) Group at the HITS, lead by

³⁴ <http://survey.dagstuhl.de>

guage-Processing (NLP) Forschungsgruppe des HITS um Prof. Dr. Michael Strube bringt dabei ihre Erfahrung mit graph- und netzwerkbasierter NLP-Methoden bei der Co-Referenz-Resolution und der Konzept- bzw. Entitäts-Disambiguierung ein.

Im Frühjahr 2013 wurde ein Projektantrag für den Leibniz Wettbewerb 2014 in der Förderlinie „Nationale und internationale Vernetzung“ eingereicht, der jedoch leider auf Grund der hohen Anzahl eingereicherter Anträge nicht berücksichtigt wurde. Eine überarbeitete Version des Antrages wurde in den Leibniz Wettbewerb 2015 eingebracht. Das Projekt wird von Schloss Dagstuhl (dblp-Team) koordiniert und soll von 2015 bis 2018 laufen. Auch ungeachtet des Erfolgs des Förderantrages wollen dblp und HITS bereits in 2014 mit dem Austausch von Daten beginnen.

■ Zusammenarbeit im Leibniz-Forschungsverbund „Science 2.0“

Die Bibliographiedatenbank dblp ist ein wesentlicher Beitrag von Schloss Dagstuhl zum Leibniz-Forschungsverbund „Science 2.0“. Dank Dagstuhls Erfahrung in den Bereichen Publikationswesen und Indexierung, sowie dank des Engagement für Open Access und Open Data, kann Schloss Dagstuhl Projekte des Forschungsverbundes als Use-Case-Partner wesentlich unterstützen.

Eine Reihe von Kooperationen mit Partnern des Forschungsverbundes sind bereits geplant. So soll in Zusammenarbeit mit Prof. Dr. Isabella Peters (Deutsche Zentralbibliothek für Wirtschaftswissenschaften, ZBW) der Zusammenhang zwischen wissenschaftlichem Zitationsverhalten und der Wahrnehmung in sozialen Netzen des Web 2.0, sowie die Anwendung alternativer Szientometrien (so genannte „Altmetrics“) mit Hilfe von dblp-Daten erforscht werden. In einem weiteren Projekt, in Kooperation mit GESIS – Leibniz-Institut für Sozialwissenschaften, soll die Integration von Daten digitaler Bibliotheken in soziale Netzwerkdienste wie etwa Facebook an Hand von aggregierten Datenquellen (unter anderem dblp) untersucht werden.

Die offenen Daten von dblp erlauben es zudem, auch ohne direkte Kooperation Projekte des Forschungsverbundes zu unterstützen. So untersucht etwa derzeit die Forschungsgruppe von Prof. Dr. Robert Jäschke (Forschungszentrum L3S) das Twitter-Verhalten von Informatikforschenden unter zu Hilfe nahme von dblp-Daten.

Prof. Dr. Michael Strube, joins the project by providing its extensive experience with graph-based and network methods for NLP tasks such as co-reference resolution, cross-document co-reference resolution, concept and entity disambiguation.

In early 2013, a project proposal was submitted to the “National and international networking” funding line of the Leibniz Competition 2014. Unfortunately, the proposal was not funded due to the large number of submissions received in that year. A revised proposal was submitted to the Leibniz Competition 2015. The project is coordinated by Schloss Dagstuhl (i.e., the dblp team) and is intended to run from 2015 to 2018. Regardless of the success of the funding application, dblp and HITS will establish a data exchange cooperation in 2014.

■ Cooperation in the Leibniz Research Alliance “Science 2.0”

The dblp computer science bibliography is recognized as a major contribution of Schloss Dagstuhl to the newly founded Leibniz Research Alliance “Science 2.0.” Particularly the experiences from the publishing and indexing departments and the commitment to the Open Access and Open Data movements enable Schloss Dagstuhl to contribute significantly to the Research Alliance’s projects.

A number of collaborations with partners of the research alliance are already planned. In cooperation with Prof. Dr. Isabella Peters (German National Library of Economics, ZBW) the connection between scientific citation behavior and perception within the social networks of the Web 2.0, as well as the use of alternative scientometrics (so-called “altmetrics”) will be studied with the help of dblp data. Another project, in cooperation with GESIS – Leibniz Institute for the Social Sciences, will study the integration of digital libraries in social networking services such as Facebook on the basis of aggregated data sources (including dblp).

The Open Data approach of dblp also allows for Schloss Dagstuhl to support research projects without an immediate collaboration. For instance, the research group of Prof. Dr. Robert Jäschke (L3S Research Center) is studying the Twitter usage of computer scientists by using data from dblp.

8

Einrichtung und Service ***Facilities and Services***

Tagungsräume

8.1

Conference Facilities

Schloss Dagstuhl bietet drei Hörsäle für 25 bis 60 Personen. Alle Hörsäle sind mit einem Beamer, einen MS-Windows-Arbeitsplatz und einer Audioanlage einschließlich Mikrophone ausgestattet. Durch diese Technik werden Vorträge, Präsentationen und Live-Vorfürungen auch verteilter Systeme optimal unterstützt. Mittels einem Presenter können Vortragende ihre vorbereiteten Materialien präsentieren, ohne zum Laptop oder Arbeitsplatz zurückzukehren. In 2013 ist ein Hörsaal zusätzlich mit einem Apple TV ausgestattet, so dass Gäste mit kompatiblen Geräten drahtlos Inhalte an den Beamer übertragen können.

Neben den Hörsälen bietet Dagstuhl sechs Seminarräume. Davon sind zwei mit modernen HDMI-fähigen Beamern ausgestattet, während in einem Hörsaal ein großes Plasmadisplay montiert ist. Fünf Beamer auf Rollwagen stehen zusätzlich zur Benutzung in allen Räumen zur Verfügung.

Die beiden größten Hörsäle sind jeweils mit mehreren Tafeln ausgestattet, während in den anderen Tagungsräumen jeweils große Whiteboards an den Wänden montiert sind. In einem Seminarraum kann sogar eine ganze Wand als Whiteboard (über 12m²) benutzt werden, da diese mit einer speziellen Farbe gestrichen wurde.

Daneben gibt es eine Anzahl weiterer Orte, an denen Gäste sich zur Diskussion in entspannter Atmosphäre treffen können. Am Abend zieht es viele Gäste in den Weinkeller und die Cafeteria, zwei der gemütlichsten Räume im Haus und hervorragend geeignet für die Fortsetzung einer produktiven Diskussion.

Schloss Dagstuhl has three lecture halls with a seating capacity of 25 to 60 each. All lecture halls are equipped with a projector, an MS Windows workplace, and an audio system including a microphone. These facilities not only enable talks and papers to be presented in an optimal manner but also permit online demonstrations of active and distributed systems to be given to large audiences. A presenter for use of those who wish to go through their presentations without physical access to a computer is also available. In 2013 one of the lecture halls was equipped with an Apple TV. Guests can thus access the beamer wireless and presenting content from their compatible computer without connecting physically.

In addition to the lecture halls, the center has six meeting rooms. Two are equipped with up-to-date HDMI projectors and one has a large plasma display at the wall. Five mobile projectors are available for use in all of the rooms.

Whereas the two main lecture halls are equipped with several blackboards, whiteboards are provided in the other rooms. One of the conference rooms features a complete "whiteboard wall" painted with a special paint which allows to use this whole wall (over 12m²) as one large whiteboard.

The center also offers a spectrum of other spaces where guests can sit and work together in a relaxed atmosphere. In the evening, guests gravitate towards our wine cellar and cafe, two of the coziest places in the house and great places for continuing with a productive discussion.

Freizeit und Ambiente

8.2

Leisure Facilities

Die Atmosphäre im Schloss wird von den Teilnehmern als kommunikativ, zur Arbeit anregend und angenehm beschrieben. Immer wieder wird berichtet, dass die schönen Räume für abendliche Treffen und die Möglichkeit, im Barocksaal zu musizieren, gute außerfachliche Erinnerungen entstehen lassen. Die im Zentrum angebotenen Freizeiteinrichtungen wurden so ausgewählt, dass sie die Kommunikation unter den Teilnehmern fördern. Neben dem mit diversen Instrumenten und Notenmaterial ausgestatteten Musikraum gibt es einen Billardraum, eine Sauna, Tischfußball, Darts und einen Freizeitkeller mit einer Tischtennisplatte sowie Fitnessgeräten. Im Sommer können ein Ballplatz mit Netz, Boulespiele im Garten sowie die beliebten Mountainbikes genutzt werden.

The participants typically describe the atmosphere at Schloss Dagstuhl as being surprisingly pleasant and instrumental in promoting valuable work and communication between the guests. Former participants frequently mention fond memories of the pleasant evenings spent in the beautiful rooms of the manor house and making music in the baroque music room. The leisure activities offered in the center have been chosen so as to promote communication among the participants. Apart from the music room which features a grand piano and various other instruments as well as sheet music, the center also has a sauna, a pool table, table football facilities, dartboard, and a recreation room with gym equipment and table tennis table. During the summer guests can use the outdoor sports grounds fitted with a net, play boules in the yard, or ride one of our mountain bikes.

Dagstuhls Küche

8.3

Dagstuhl's Kitchen

Die Mahlzeiten sind ein wichtiger Bestandteil des wissenschaftlichen Programms von Schloss Dagstuhl. Die Sitzordnung wird absichtlich stets zufällig gemischt, um

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are deliberately mixed in order to break up cliques and

eingefahrene Gruppen aufzuteilen und Gäste zu ermuntern, während ihres Aufenthalts möglichst viele verschiedene Kollegen kennen zu lernen. Runde Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten. Während der Pausen können sich Gäste bei heißen Getränken und einer Kleinigkeit in der Cafeteria erholen. Abends gibt es Brot und eine Auswahl an Käsen im Weinkeller — dem Brennpunkt der Dagstuhl Aktivitäten nach dem offiziellen Programm. Jeder Raum hat im Einklang mit den wechselnden Stimmungen eines jeden Tages seine ganz eigene Atmosphäre und unterstützt die vielfältigen Treffen von Teilnehmergruppen.

Dagstuhls Philosophie des Kochens ist einfach: saisonal, gesund, schmackhaft und ausschließlich aus frischen Zutaten bereitet. Alles bis hin zu Gelees und Konfitüren wird jeden Tag frisch von den 12 Mitarbeitern der Küche zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dieses steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber durchaus der Mehrheit der internationalen Gästen des Zentrums entgegen. Hausgemachter Kuchen, Kaffee und Tee unterbrechen angenehm die tägliche Routine und anstrengenden Diskussionen.

Sowohl die Zutaten als auch die Gerichte wechseln entsprechend der Saison. An warmen Sommerabenden wird häufig auf der Terrasse vor dem Speisesaal gegrillt und unter anderem saarländischer *Schwenker* (eine lokale Variante des Grillsteaks, die unter dauerndem Schwenken des Rosts zubereitet wird,) den Gästen angeboten. In den kalten Monaten stehen häufig wärmende Suppen auf dem Speiseplan, Spezialitäten wie Spargelcreme und Kürbis-suppe gibt es den Jahreszeiten entsprechend über das ganze Jahr. Im Allgemeinen werden im Winter schwerere und im Sommer leichtere Speisen von der Küche zubereitet, die über das Jahr eine ausgewogene Mischung an regionalen und internationalen Spezialitäten aus neuen sowie bewährten und geliebten Gerichte bietet. Auf Nachfrage bereiten unsere Mitarbeiter auch abgestimmte Gerichte für Vegetarier, Veganer und Gäste mit Lebensmittelunverträglichkeiten zu.

Um all diesen Anforderungen innerhalb auf eine vernünftige Größe beschränktem Budget gerecht zu werden, bietet die Küche zu festen Zeiten ein Menü am Abend und Morgens und Mittags ein Buffet an. Entsprechend bedienen sich die Gäste bei Frühstück und Mittagessen selbst, während am Abend von unseren Mitarbeitern das Menü serviert wird. Bei allen Mahlzeiten räumen die Gäste selbst das gebrauchte Geschirr auf die bereitstehende Servierwagen. Dies ist verglichen mit den frühen Anfängen, in denen alle Mahlzeiten eher formal in der europäischen Tradition mit Tischdecken und Stoffservietten stattfanden. Dagstuhl bietet daher nun eine viel entspanntere fast familiäre Atmosphäre in dem Speisesaal mit den großen Fenster zum Garten des Hauptgebäudes. Der Speisesaal bietet bis zu 95 Personen Platz wähen im allgemeinen 60 bis 75 Gäste hier speisen.

encourage guests to talk to as many different people as possible during the course of their stay. Round tables in the dining hall promote collaborative interaction at breakfast and lunch. During breaks and in the evening, guests can enjoy a warm drink and a snack in the café or a round of bread and cheese in the wine cellar – the soul of the after-hours scene at Dagstuhl. Each space has its own unique atmosphere, in keeping with the shifting mood of each day and the different kinds of gatherings they support.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy and tasty meals using only whole, unprocessed foods. Everything, down to the jams and jellies, are freshly prepared each day by the kitchen's 12-person staff. The focus is on lighter fare during the day in order to aid scientists' concentration and on a warm meal in the evening, which breaks with the German tradition of a cold evening meal but fits well with internationality of the center's guests. Homemade cake, coffee and tea each afternoon punctuate the daily routine.

Both ingredients and dishes vary with the changing seasons. On warm summer evenings, guests are frequently invited to partake of grilled *Schwenker* (the local variant of barbecued steak) on the outdoor patio adjacent to the dining hall. Warm soups often appear on the menu during the colder winter months, and German culinary favorites such cream of asparagus or butternut squash are served whenever the fresh ingredients are in season. In general, the kitchen tries to keep meals lighter in the summertime and heavier in the winter, offering a blend of regional and international dishes year-round that include some new recipes and many tried-and-true Dagstuhl favorites. Special dishes for those with medical food conditions and vegetarians or vegans are prepared upon request.

To accomplish all of this within a reasonable budget, the center offers a set menu for dinner and a buffet-style breakfast and lunch at fixed times. Guests normally serve themselves during the daytime meals, sitting to a served meal only in the evening. Everybody clears his own tray. This is a great change from the center's earlier years, when guests were served in formal European fashion, with tablecloths and linen napkins. Today's Dagstuhl offers a much more relaxed, family-style atmosphere in the large dining hall that opens onto the garden of the main building. The dining hall seats up to 95 guests, serving typically 60 to 75 persons at each meal.

Die Geschichte der Speisekarte von Dagstuhl

Als wir im Sommer 1990 mit dem Seminarbetrieb anfangen, übernahm die Küche die deutsche Tradition, mittags eine warme Mahlzeit, abends kalte Küche. Bald merkten wir, dass das warme, meist ziemlich schwere Mittagessen unsere Gäste in ein Verdauungstief versetzte und ihre Aufmerksamkeit für die Nachmittagsvorträge reduzierte. Wir änderten sofort die Mahlzeiten; mittags gab es eine Suppe und ein Büffet bestehend aus Rohkost und Salaten, abends eine warme Mahlzeit, zwischendurch Kaffee und Kuchen. Das rettete den Nachmittag. Womit ich nicht gerechnet hatte, was der Ehrgeiz unseres Küchenpersonals. Ein schlichtes Rohkostbüffet jeden Mittag war einfach zu spartanisch. Deshalb wuchs das Angebot auf dem Mittagbüffet langsam aber sicher. Zarte Hinweise, sanfte Ermahnungen und deutliche Worte zu der Ausrichtung des Mittagessens, alles wurde überhört. Die Bemühungen, meine Kritik in freundliche Komplimente zu verpacken, „Das Büffet war heute ja wieder außerordentlich und reichlich!“, stießen auf taube Ohren. Also muss ich in diesem Punkt eine meiner wenigen Niederlagen auf der ganzen Linie in meinem Wirken in Dagstuhl bekennen.

Ich weiß nicht, ob es nur mit dieser eben eingestandenen Niederlage zusammenhing. Aber die Komplimente für die Küche nahmen über die Jahre ständig zu. Das hing auch damit zusammen, dass die Küchenmannschaft immer besser und experimentierfreudiger wurde. Die Bibliothek an Rezeptbüchern wuchs; es wurde immer wieder Neues ausprobiert. Aber erstaunlicherweise kam die regionale Küche immer besonders gut an. Hoorische am Freitag Mittag waren der Renner!

Erwähnt werden muss allerdings auch die Backfisch-Affäre. Die hatte nichts mit einem pubertären Küchenlehrling zu tun, sondern mit dem Speiseplan. Da wir selten dieselben Gäste in aufeinanderfolgenden Wochen hatten, war es für die Küche ohne Risiko, denselben Speiseplan für mehrere aufeinanderfolgende Wochen beizubehalten. Nur, es gab einen, der in vielen aufeinanderfolgenden Wochen Donnerstag Abend in Dagstuhl aß, nämlich mich. Als ich zum dritten Mal und dann zum vierten Mal Donnerstag abends Backfisch serviert bekam, war dies der Küche sehr peinlich.

Die Küche bietet Gästen Diäten bei Unverträglichkeiten an. Es ist erstaunlich, wie Unverträglichkeiten zunehmen! Gut, Vegetarier gab es immer schon, Veganer auch schon einige Zeit. Aber, wenn heute der Bauch zwickt, ist das nicht mehr irgendetwas Unverdauliches wie zu viel gezähnte Erbsen oder zu viel ausgeschiedene Korinthen, nein es muss an einer Unverträglichkeit gegen ein Nahrungsmittel liegen! Es ist erstaunlich, welche Vielfalt an Unverträglichkeiten der moderne Mensch aufweisen kann, und dass diese auch noch manchmal alle gleichzeitig bei Gästen einer Woche in Dagstuhl vorhanden zu sein scheinen. Da kommt die Küchenmannschaft ins Schwitzen. Wenn anschließend moniert wird, dass die Küche nicht innerhalb einer Diät noch Wahlmöglichkeiten anbietet, dann merke ich, wie die Küchenmannschaft anfängt, Unverträglichkeit von gewissen Gästen zu entwickeln.

The Story of Dagstuhl's Menu

When we first launched our seminar program in the summer of 1990, the kitchen planned meals in accordance with the German tradition of a hot meal at noon and a cold one in the evening. However, we soon realized that the warm, usually quite heavy lunch slowed our guests' digestion and reduced their alertness during the afternoon lectures. We immediately changed to serving soup and a buffet consisting of raw vegetables and salads for lunch, and a hot meal in the evening, with coffee and cake in between. That saved the afternoon. What I did not expect was the ambitiousness of our kitchen staff. A simple raw food buffet for every lunch was just too spartan for them. Slowly but surely, the lunch buffet spread. Delicately phrased notes, gentle admonitions and clear words with regard to the orientation of the lunch – nothing worked. My efforts to couch friendly criticism in compliments, such as “The buffet was so plentiful and extraordinarily again today!” fell on deaf ears. In the end, I had to concede defeat – one of my few defeats as far as my work at Dagstuhl is concerned – on this detail.

I do not know if it is related only to this confessed defeat, but the compliments for our kitchen have grown constantly over the years. This also has to do with the fact that the kitchen staff have continually improved and become more adventurous over time. The library of recipe books has also grown, with new recipes being tried out again and again. But surprisingly, the local food dishes have always been particularly well received. Hoorische (a Saarland speciality consisting of potato dumplings with sauerkraut and meat sauce) is always a hit!

However, the Baked Fish Affair must also be mentioned. This had nothing to do with a pubertal kitchen apprentice, but rather with the menu. Since guests rarely visit Dagstuhl for longer than one week, it was in the beginning no risk for the kitchen to maintain the same menu plan for several consecutive weeks. The only problem was that there was one person who dined at Dagstuhl on many consecutive Thursday evenings, namely me. When I was served bake fish for dinner a third, and then a fourth time, the kitchen was very embarrassed.

The kitchen prepares special meals for our guests who have allergies and food intolerances. It's amazing how intolerances grow! Well, we've always had vegetarians, and vegans too for some time now. But today, if something slightly tweaks the stomach... well, this can't just be a passing indigestion, a matter of counting too many beans – no, it must be due to a food intolerance! It's amazing what diversity of intolerances modern man can sustain. Sometimes all of these intolerances even seem to be present among our Dagstuhl guests simultaneously on the same week! Then the kitchen crew really starts to sweat. If it is subsequently remarked that the kitchen doesn't offer a wide enough selection of dishes, suited to every diet, then I notice how the kitchen staff develops an allergy to certain guests.

Fig. 8.1

An anecdotal account by Schoss Dagstuhl's Scientific Director, Professor Reinhard Wilhelm.



Fig. 8.2
The lunch buffet offers an appetizing assortment of salads and other light foods.



Fig. 8.3
Homemade cake and coffee or tea are an afternoon tradition at Schloss Dagstuhl.

Kinderbetreuung

8.4

Childcare

Viele unserer internationalen Gäste möchten ihre Kinder nach Schloss Dagstuhl mitbringen, da sie wegen fehlender Kinderbetreuung zu Hause ansonsten nicht an den Veranstaltungen teilnehmen könnten. Zur Familienförderung bietet Schloss Dagstuhl seinen Gästen seit einigen Jahren während den Vortragszeiten eine Kinderbetreuung an. Bei Bedarf wird dazu eine erfahrene, staatlich-geprüfte Betreuerin verpflichtet.

Ebenso wie für die Kinder übernimmt Schloss Dagstuhl die Aufenthaltskosten einer Begleitperson zur Kinderbetreuung. In 2013 besuchten 21 Kinder Schloss Dagstuhl. Davon wurden 14 Kinder durch einen Tagesmutter und sieben weitere durch Verwandte betreut.

Many of our international guests who would otherwise be unable to take part in the events due to a lack of childcare options at home would like to bring their children with them to Dagstuhl. In order to promote family friendliness, Schloss Dagstuhl offers to arrange qualified child care for participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops who bring young children with them. Children are looked after on-site during the seminar meeting times by a qualified nanny.

Parents also have the option to bring along their own "nanny," usually a spouse or relative, whose room and board costs are gladly absorbed by the center just as they are for children. In 2013, Dagstuhl hosted 21 children, 14 of whom were cared for by a nanny on site and seven by relatives.

Computer und Vernetzung

8.5

Computers and Networks

Schloss Dagstuhl bietet seinen Gästen eine adäquate Anbindung an das Internet. Seit 2013 erfolgte die Anbindung an das Internet über das DFN mit zwei redundanten 100 Mbit/s Leitungen. Fast im ganzen Zentrum können sich Gäste über WLAN (IEEE 802.11 b,g,n) mit dem Internet verbinden. Der Zugriff erfolgt entweder über eduroam oder über eine Dagstuhl-eigene Kennung. Die Seminar-Organisatoren haben im größten Seminarraum „Saarbrücken“ die Möglichkeit, den WLAN Empfang abzuschalten.

Neben dem Zugang über mitgebrachte Laptops, Tablet Computer oder Smartphones stehen den Gästen einige fest installierte Arbeitsplätze zur Verfügung. Schloss Dagstuhl hat einen Rechnerraum mit acht festen Arbeitsplätze. Davon sind zwei Arbeitsplätze mit Apple Macs ausgerüstet, zwei sind dedizierte MS-Windows-Arbeitsplätze. Vier weitere Arbeitsplätze können wechselweise entweder mit Linux oder mit MS Windows gestartet werden. Zusätzlich steht Benutzern eines Laptops ein externer Monitor samt Tastatur und Maus zur Verfügung. In diesem Raum stehen auch direkte Ethernet-Anschlüsse zur Verfügung, um das WLAN wegen Bandbreite oder Kapazitätsgründen zu umgehen. Weiterhin bietet Schloss Dagstuhl seinen Gästen zwei iPads sowie auf Nachfrage einen MacBook Pro und einen Laptop mit MS Windows.

Im Zentrum steht den Gästen ein Multifunktions-Farbdrucker mit Scanner und Kopierer, ein weitere Farbdrucker und ein S/W-Drucker zur Verfügung. Der Zugriff erfolgt vorzugsweise über eine Weboberfläche, die das direkte Drucken zahlreicher Dokumentenformate erlaubt. Alternativ können die Drucker mittels entsprechender Treiber auch direkt aus dem lokalen Netzwerk angesteuert werden.

Zu der IT-Ausstattung gehören weiterhin fünf Recherche-Arbeitsplätze in der Bibliothek sowie drei fest installierte Rechner in den Seminarräumen.

Schloss Dagstuhl offers its guests an adequate connection to the Internet. Since 2013 the center is connected to the Internet by the DFN (German Research Network) using two redundant 100 Mbit/s lines. Throughout the grounds guests have Internet access by Wi-Fi (IEEE 802.11 b,g,n). Access is either via eduroam or a Dagstuhl-hosted private account. In our largest seminar room, "Saarbrücken," organizers may choose to disable the main Wi-Fi connection during meeting times.

Most of our guests prefer to access the Internet via their laptops, tablet computers and smartphones, but they are also free to use the workstations in our computer room. Schloss Dagstuhl offers one computer room including eight workstations. Among them there are two Apple Macs, two dedicated MS Windows workstations and four workstations providing either Linux or MS Windows by a dual boot method. There is also a large display together with an external keyboard and mouse for users with their own laptop. Several Ethernet cables with Internet connection are also provided to bypass the rate-restricted Wi-Fi connection. Two iPads, and upon request a MacBook Pro and a laptop with Windows are also available for use throughout the grounds.

Schloss Dagstuhl provides a multifunction color printer with scanner and copier, a color printer, and a black and white printer. The preferable access method is to use a dedicated web front end which allows to upload and print the most used document formats without converting them. Alternatively, guests can use the appropriate printer drivers on their computers to directly access the printers via the network.

The center's IT equipment also includes five workstations in the library for literature research, as well as three fixed computers in the lecture halls.

Dagstuhl's Web-basierte Dienste

8.6

Dagstuhl's Web-based Services

8

Schloss Dagstuhl bietet allen Organisatoren und Gästen eine wachsende Anzahl Web-basierter Dienste. Während der Vorbereitungsphase können alle Organisatoren tagesaktuell überprüfen, welche eingeladen Gäste bereits zu- oder abgesagt haben. Sie können ebenfalls einen (vorläufigen) Zeitplan auf der seminarspezifischen Webseite hochladen. Alle Teilnehmer können Dokumente zu ihrem Vortrag oder dem Seminar hochladen, die für alle anderen zugreifbar sind. Weiterhin werden jedem Seminar ein MediaWiki und ein WebDAV-Repository angeboten. Die Erstellung der Seminardokumentation innerhalb der Reihe Dagstuhl Reports wird ebenfalls durch ein Web-Frontend unterstützt.

Schloss Dagstuhls Internetauftritt³⁵ bietet nicht nur seinen Gästen sondern allen Nutzern weltweit Informationen über die folgenden Themen:

- Verbreitung allgemeiner Informationen über das Zentrum, wie Konzept, Programm, Antragsmodalitäten, Stiftung
- Informationen zur Anreise der Teilnehmer, wie Lageplan, Fahrpläne, Taxidienste
- Die Bibliothek mit der Möglichkeit zur Recherche im Dagstuhl-Bibliothekskatalog
- Informationen zu Seminaren und Veranstaltungen, wie Seminarziele, angemeldete Wissenschaftler und Publikationen
- Angebot einer Plattform zum Austausch von Material unter den Seminarteilnehmern

Der Webserver verwaltet die Inhalte mit dem Content-Management-System Typo3. Außer statischen Seiten – fast alle in deutschen und in englischen Versionen – werden auch dynamische Seiten angeboten, die über eigene Software generiert werden. So gibt es zu jedem Seminar eine dynamisch generierte Seite, die zu Motivationstext, Teilnehmerliste, Publikationen, etc. führt.

Schloss Dagstuhl offers an increasing number of web-based services to seminar organizers and participants. During the preparation phase, the seminar organizers can check how invited participants are responding to the invitation and which of them have committed to attending. They can also upload a (preliminary) schedule to the seminar web page. All participants can upload seminar- or presentation-related documents to the page, which are then accessible to everyone else. A MediaWiki and WebDAV-related repository are also offered. The making of the seminar documentation inside our Dagstuhl Reports periodical is also supported by a Web-based service.

In keeping with the center's philosophy, its Internet³⁵ offerings are not only available to the guests at Dagstuhl but to netizens throughout the world. Objectives and content:

- Dissemination of general information on the center, e.g. concept, program, particulars pertaining to proposal submission, the Foundation
- Offering participants travel information on how to get to the center (site plan, train and bus schedules, taxi services, etc.)
- Presenting the Dagstuhl Informatics Research Library along with its offerings and resources and enabling research in the Dagstuhl library catalogue
- Provision of information about seminars and events (e.g. seminar objectives, scientists from whom proposals have been accepted, publications)
- Providing a platform for exchanging materials among seminar participants

The web server administers the content using the Typo3 content management system. Apart from static pages, almost all of which are in German and English, dynamic pages are also offered which are generated by the center's proprietary software. Each seminar has a dynamically generated page of its own featuring links to a seminar description, list of participants, publications, etc.

³⁵ <http://www.dagstuhl.de/>

9

Bibliothek

Research Library

Bestand und Angebot

9.1

Inventory and Offering

Die Forschungsbibliothek bildet eines der wichtigsten Angebote. Sie hat sich Dank der Startfinanzierung der Volkswagen-Stiftung und durch zahlreiche Buchspenden von Verlagen und Seminarteilnehmern zu einer der bedeutendsten Informatik-Forschungsbibliotheken in Deutschland entwickelt.

Die Bibliothek erwirbt aktuelle Informatik-Forschungsliteratur thematisch zu den jeweiligen Seminaren, überwiegend in englischer Sprache. Am 31. Dezember 2013 umfasste der Bibliotheksbestand 60 394 bibliographische Einheiten, die vollständig im Online-Katalog verzeichnet sind. Eine Besonderheit ist der umfangreiche Zeitschriftenbestand, der fast komplett elektronisch bezogen wird. Neben den abonnierten Zeitschriftentiteln, ermöglicht die Bibliothek Zugriff auf mehrere Tausend weitere elektronische Zeitschriftentitel und Zeitschriftenarchive über die DFG-geförderten National- und Allianzlizenzen.

Die Literatur wird in einem attraktiven Bibliotheksturm auf vier Ebenen präsentiert, der auch zahlreiche Leseplätze zum Studium anbietet. Als Präsenzbibliothek steht sie den Dagstuhl-Seminarteilnehmern für ihre Forschungsarbeit vor Ort rund um die Uhr offen. Auch externe Wissenschaftler können die Bibliothek nach Voranmeldung nutzen.

Durch die Teilnahme an der Online-Fernleihe steht der komplette Zeitschriftenbestand im Rahmen des internationalen Leihverkehrs Bibliotheken aus der ganzen Welt zur Verfügung. Dazu ist der komplette Zeitschriftenbestand auch in der Zeitschriftendatenbank (ZDB) sowie in der Elektronischen Zeitschriftenbibliothek (EZB) nachgewiesen. Zusätzlich ist die Bibliothek Teilnehmer an LITexpress, einem Lieferdienst rückgabepflichtiger Medien für Bürgerinnen und Bürger in Rheinland-Pfalz, dem Saarland und der deutschsprachigen Gemeinschaft Belgiens. Speziell die Archivtitel der Bibliothek sollen dadurch zur Ausleihe bereitgestellt werden.

Die Bibliothek präsentiert regelmäßig umfangreiche Buchausstellungen. Jede Woche wird im 1. Obergeschoss eine Ausstellung aller vorhandenen Bücher der Autoren präsentiert, die an den aktuellen Dagstuhl-Seminaren teilnehmen. An die Autoren ergeht gleichzeitig die Bitte, ihre Bücher zu signieren. Andere Buchausstellungen werden auf Wunsch von Organisatoren zu einem speziellen Thema zusammengestellt. Weiter werden alle Buchspenden von Verlagen separat ausgestellt und regelmäßig aktualisiert. Dieser Service wird von Gästen und Verlagen sehr geschätzt.

Über die Internetseite der Bibliothek³⁶ sind u.a. der Online-Bibliothekskatalog, die Zeitschriftenbestandsliste mit Zugang zu den abonnierten online verfügbaren Zeitschriften sowie weitere Informationsangebote der Bibliothek zu erreichen.

The Dagstuhl Informatics Research Library is one of the center's most impressive offerings. Thanks to the startup financing by the Volkswagen Foundation and numerous book donations of publishing houses and seminar participants, it numbers among Germany's key informatics research libraries.

The library collects current research literature on informatics topics for the respective seminars, primarily in English. As of December 31, 2013, the library's collection totaled 60 394 bibliographic units, all of which are contained in the online catalog. One distinguishing feature is the center's impressive holdings of journals and periodicals, almost all of which are in electronic form. Apart from subscribed journals, the library also provides access to several thousand other electronic journals and journal archives via the DFG-funded national and alliance licenses.

The literature is arranged on four levels in an attractive library tower, which also offers a large number of recesses for quiet study and research. Being a reference library, it is at the disposal of the Dagstuhl Seminar participants 24/7 for their research work on site. External scholars can also use the library provided they register beforehand.

In order to support informatics research in Germany and throughout the world, the center's entire holdings of periodicals are also made available to other libraries, particularly by way of inter-library loan. The library's entire holdings of journals and periodicals are additionally listed in the Zeitschriftendatenbank (ZDB), the world's largest specialized database for serial titles, and in the Electronic Journals Library (EZB). The library is a member of LITexpress, the Virtual Library of Rhineland-Palatinate, Saarland and the German-speaking community of Belgium, a media loan service for the citizens of these three areas. The library's archive items in particular are designed to be made available for loan.

The library regularly arranges comprehensive book exhibits. Each week all the books authored by the participants in the current Dagstuhl Seminars which are available in the library are displayed on the first floor. The authors are kindly asked to sign their books. If desired, book exhibits on a particular topic are also put together by the organizers. In addition, all book donations received from publishers are exhibited separately and the exhibits are regularly updated. This service is highly appreciated by the center's guests and publishers alike.

The online catalogue and a comprehensive journal list with access to the subscribed journals as well as other information offerings can be accessed via the library's webpage.³⁶

³⁶ <http://www.dagstuhl.de/de/library/>

Spenden an die Bibliothek

9.2

Die Bibliothek von Schloss Dagstuhl profitiert durch zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2013 Buchspenden von den Verlagen, die in Fig. 9.1 aufgeführt sind. Auch viele Seminarteilnehmer spenden der Bibliothek ihre Bücher. Autorenexemplare, insbesondere von wichtigen, bereits vergriffenen Büchern, werden ebenso dankbar entgegengenommen. Insgesamt erhielt das Zentrum im Berichtszeitraum 951 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

9

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. During 2013 the Informatics Research Library received book donations from the publishers listed in Fig. 9.1. The center is also grateful for donations of author's copies, particularly those of major works that are out of print. The center received a total of 951 volumes during the year 2013 as donations from publishing houses and seminar participants.

Birkhäuser Verlag

<http://www.birkhaeuser-science.com>

Eurographics – European Association for Computer Graphics

<https://www.eg.org>

SIAM – Society for Industrial and Applied Mathematics

<http://www.siam.org>

Springer-Verlag GmbH | Springer Science+Business Media

<http://www.springer.com>

Fig. 9.1

Donations from publishers to the Dagstuhl library.

10 Kunst *Art*

Dagstuhl als Galerie

10.1

Dagstuhl as Art Gallery

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des schmalen Flurs sowie die hervorragende Ausleuchtung mit starken Kontrasten zwischen Tag und Nacht bieten den Künstlern sehr gute Möglichkeiten, ihre Werke darzustellen. Die Kunstwerke an den Wänden des schmalen Gangs durchbrechen die Nüchternheit des Neubaus in anregender und angenehmer Weise. Die wechselnden Ausstellungen bieten einen erfrischenden und dynamischen Kontrast zu der ständigen Kunstsammlung von Schloss Dagstuhl. Seitens den Gästen wird oft die Ausstrahlung, die von dem Kunstangebot ausgeht, gelobt und als Teil der „Magie“ von Dagstuhl wahrgenommen.

Der Wissenschaftliche Direktor Prof. Reinhard Wilhelm fungiert als Kurator und legt die Ausstellungen fest. Oft leitet er auch die Vernissagen. Jedes Jahr veranstaltet Schloss Dagstuhl etwa fünf Ausstellungen für jeweils zwei bis drei Monate.

Insgesamt fünf Ausstellungen fanden 2013 in Schloss Dagstuhl statt, von denen zwei unter Beteiligung von Seminarteilnehmern entstanden sind.

■ Forschungsaufenthalt von HBK Saar Kunststudenten

Diese ungewöhnliche Ausstellung war der Arbeit von zehn Studenten für Fotografie und Bildende Kunst sowie ihrer Professorin Gabriele Langendorf von der Hochschule für Bildende Künste Saarbrücken gewidmet. Die Erfahrungen der Seminarteilnehmer der Dagstuhl-Seminare 13071 und 13072 war das zentrale Thema der Skizzen, Zeichnungen und Fotografien. Die Studenten, die bereits an ähnlichen Projekten mit der Deutschen Radio Philharmonie Saarbrücken Kaiserslautern und dem Landtag teilgenommen hatten, besuchten Schloss Dagstuhl vom 10. bis zum 14. Februar 2014. Sie hatten den Auftrag, ihre Eindrücke der Dagstuhl-Seminare 13071 und 13072 darzustellen, ohne die Seminarteilnehmer zu stören. Das Ergebnis war eine höchst kreative Kombination aus Fotos, Collagen, Skizzen, Gemälden und mit Mischtechniken erstellten Werken, die auf kraftvolle Weise die Energie und Dynamik des „Dagstuhl-Erlebnisses“ vermitteln.

Reporter des SR Fernsehen dokumentierten die Veranstaltung im Nachrichtenmagazin „Aktueller Bericht“ vom 14. Februar 2013. Die Fernsehreportage zeigt Nahaufnahmen der von den Studenten angefertigten Zeichnungen und Fotos, Interviews mit den Künstlern und ein Gespräch mit Reinhard Wilhelm. Er gab einen kurzen Überblick zu dem Kunstprogramm von Schloss Dagstuhl und die Gemeinsamkeit des „abstrakten Denkens“, die Informatiker und Künstler verbindet.

Die Studenten stifteten zwei gerahmte Zusammenstellungen von Fotos und Skizzen, die während ihres Aufenthaltes entstanden waren. Das Geschenk ist im Gästehaus von Schloss Dagstuhl ausgestellt.

Exhibitions of artists are regularly organized in the so-called cloister of the new building. The spacious surroundings, excellent lighting, and dramatic day-to-night contrast offer artists a very special and unique exhibition space. Artworks are arranged along the walls of the narrow gallery, providing an intriguing juxtaposition to the otherwise ascetic nature of the new building. The temporary exhibits offer a fresh and dynamic counterpoint to center's permanent collection, which can be found scattered throughout the three buildings. The center's scientific guests often praise the special atmosphere created by the art offerings as contributing to the overall “magic of Dagstuhl.”

Schloss Dagstuhl Scientific Director Reinhard Wilhelm acts as curator and sets the artistic program of the center, often also personally hosting exhibit openings. The center holds approximately five art exhibits per year, with each exhibit generally running for two to three months.

In 2013, the center organized a total of five exhibits, including two that drew on the active participation of Dagstuhl seminar participants.

■ HBK Saar Student Research Visit

This unusual art exhibit featured the work of ten photography and fine art students from the Saarbrücken University of Art and Design (HBK Saar) whose photographs, sketches and drawings captured the experience of participants in Dagstuhl Seminars 13071 and 13072. The students, who had already participated in similar projects with the Saarland radio symphony orchestra and state parliament, stayed at Schloss Dagstuhl together with their instructor, Professor Gabriele Langendorf, on February 10–14, 2013. They were instructed to hang around the Schloss and portray their impressions of Dagstuhl Seminars 13071 and 13072 without disturbing seminar guests. The result was a highly creative mix of photos, collages, sketches, paintings and mixed-media compositions that vibrantly transmitted the energy and dynamism of the Dagstuhl experience.

The “Aktueller Bericht” news program of SR TV featured the event on February 14, 2013, giving close-ups of the students' drawings and photos, artist interviews, and a discussion with Reinhard Wilhelm, who gave a brief overview of Dagstuhl's art program and the “abstract thinking” connection between computer scientists and artists.

The students donated two framed collections of photos and sketches produced during their stay. The gift is now on display in Dagstuhl's new guest house.



Fig. 10.1
The Dagstuhl Seminar experience. Portrayed by HBKsaar student artist Joanna Crittendon.



Fig. 10.2
Another seminar perspective. Drawn by Joanna Crittendon.

■ **Gekrümmte Wirklichkeit: Bögen zwischen Kunst und Wissenschaft**

Zusammen mit den Teilnehmern des Seminars „Drawing Graphs and Maps with Curves“ (13151) organisierten Maxwell J. Roberts, Stephen Kobourov und Martin Nöllenburg diese Ausstellung. Verbindendes Thema war die Benutzung von gekrümmten Linien in der Darstellung von Routen, Strukturen, Bewegungen, Zeit und Verbindungen in zahlreichen Zusammenhängen. Werke dieser Ausstellung werden nun auf einer eigenen [Webseite](#)³⁷ von Schloss Dagstuhl präsentiert.

■ **Mechthild Schneider Fotografie**

Die saarländische Fotografin und Künstlerin Mechthild Schneider beschäftigt sich bereits seit Jahren mit ihrer Fotografie auch in künstlerischer Weise mit der Natur. Nach der Darstellung der Weite des Raumes in großen Querformaten und der vergänglichen Schönheit in ihren Blütenfotografien geht es ihr jetzt um das Überzeitliche, das Ewige in der Natur. Zentrale Rolle in den gezeigten Werken spielt die Verwendung von oftmals alten oder fast pompösen Bilderrahmen, die den Betrachter die Bilder auf eine ganz neue Art erleben lassen.

■ **Peter Amici, „Bilder“**

Peter Amici, in der Schweiz geboren, ist ein freischaffender Künstler, der in Deutschland lebt. Seine Werke umfassen Zeichnungen, Bilder, Bücher und Objekte. Amici's Bilder leben aus einem reichen Fundus von Gesehenem und Erfundenem. Sie zeugen von lebendiger Beziehung zur Natur wie von einer reichen Kenntnis der verwendeten Materialien.

■ **Sabine Eisenbrand, „Sprünge“**

Die im Saarland geborne Sabine Eisenbrand arbeitet und lehrt als freischaffende Künstlerin in Varel. Sowohl in ihren Werken als auch in den von ihr verwendeten Materialien zeigt sie eine ungewöhnlich große Bandbreite. Ihre Werke zeugen jedoch immer wieder von einer Gemeinsamkeit: Die lebhafteste Freude am Tun.

■ **Bending Reality: where arc and science meet**

Psychologist Maxwell J. Roberts teamed up with computer scientists Stephen Kobourov and Martin Nöllenburg to organize this exhibit in connection with Dagstuhl Seminar 13151, “Drawing Graphs and Maps with Curves,” focusing on the use of curves in routes, structures, movement, time and connections, in a variety of contexts.

A permanent digital exhibit of the collection is available on the Dagstuhl [webpage](#)³⁷.

■ **Photography of Mechthild Schneider**

Saarland-based photographer and artist Mechthild Schneider has been focusing on artful nature photography for a number of years. Following her portrayal of the vastness of space in large landscape format and of ephemeral beauty in her flower photographs, she is now focusing on the theme of the eternal in nature. Old, sometimes almost pretentious framing plays a central role in Ms. Schneider's work, showing it from a completely new perspective.

■ **Peter Amici, “Pictures”**

The Swiss-born artist Peter Amici is a freelance artist based in Germany whose working media include drawing, painting, objects, and books. Amici's images are evidence of a living relationship to nature and a rich knowledge of the materials used.

■ **Sabine Eisenbrand, “Sprünge”**

Sabine Eisenbrand, born in Saarland, is a freelance artist and art teacher based in Varel, Germany. Her art works and the materials she uses cover an extraordinarily broad spectrum. One central theme runs through all of her work: the sheer fun of doing.

Kunstankauf durch Spenden

Das Internetangebot von Dagstuhl enthält eine Seite, die es Teilnehmern, Einzelpersonen und Gruppen ermöglicht, Kunst für Dagstuhl zu stiften. Die Kunstobjekte werden über das Internet angeboten, dabei wird der Preis in kostengünstige Anteile aufgeteilt. Sobald alle Anteile eines Bildes gezeichnet sind, werden die Teilnehmer aufgefordert, den Gegenwert der bestellten Anteile als Spende einzuzahlen, wodurch dann das Objekt angekauft werden kann. Die Stifter werden sowohl in der virtuellen Internet-Galerie von Schloss Dagstuhl als auch an dem realen Objekt

10.2

Art Sponsorship and Donations

Dagstuhl's website contains a page featuring an Internet gallery enabling participants, individuals, and groups to make contributions to Dagstuhl for art donations. The works of art are featured on the Internet and donations are made by acquiring shares at affordable prices. Donors pay the value of their pledged shares as soon as a piece is fully subscribed for, thus allowing it to be purchased. Donors' names appear in Dagstuhl's online art gallery and also on the art items themselves. The art donation program also benefits the center, enabling Schloss Dagstuhl to purchase

³⁷ <http://www.dagstuhl.de/en/ueber-dagstuhl/kunst/bending-reality/>

genannt. Gleichzeitig ist es fruchtbar für das Zentrum, da die Möglichkeit besteht, Werke von Künstlern, die auf Schloss Dagstuhl ausgestellt haben, anzukaufen.

Auch 2013 erhielt Schloss Dagstuhl dankenswerterweise mehrere außergewöhnliche Spenden von Privatpersonen. Unter den gespendeten Werken sind zwei gerahmte Zusammenstellungen mit Skizzen und Fotos von Kunststudenten der HBK Saar, die während ihres „Forschungsaufenthalts“ in unserem Zentrum entstanden waren (siehe oben) und das Werk „Blooming Businesses“, das Professor Ben Schneiderman von der Universität Maryland mit Hilfe des Treemap-Algorithmus zur Datenvisualisierung geschaffen hat. Für seine historische Kunstaustellung erhielt Schloss Dagstuhl als Spende von Manfred Stein eine fotografische Reproduktion eines Portraits von Octavie de Lasalle von Louisenthal (1811–1890).

Weitere Spenden und Details werden in Kapitel 1 diskutiert. Nähere Informationen und aktuelle Neuigkeiten finden sich auf der [Kunst-Webseite](http://www.dagstuhl.de/art/)³⁸ von Dagstuhl. Wir möchten diese Stelle nutzen, allen Spendern, die 2013 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

works of art from those who exhibit at the center. Three works were purchased this way in 2013.

Schloss Dagstuhl also received several extraordinary private donations in 2013. The gifts included two framed collections of sketches and photos done by HBK Saar fine arts students during their visit to the center in February 2013 (see above), and the digital image “Blooming Businesses,” created by University of Maryland Professor Ben Schneiderman using the tree map algorithm for data visualization. For its historical art collection, the center also received as a gift from Manfred Stein a framed photographic reproduction of a portrait of Octavie de Lasalle von Louisenthal (1811–1890).

More details about these and other donations can be found in Chapter 1. For further information about Dagstuhl’s art program in general, please visit Dagstuhl’s [art webpage](http://www.dagstuhl.de/art/)³⁸. We would like to take this opportunity to thank all those who have made art donations in 2013.

³⁸ <http://www.dagstuhl.de/art/>

Forschungsaufenthalt
Works by students Joanna Crittendon, Pascal Elsen, Charlotte Geisler, Malika Hagemann, Tobias O. Heitz, Tanja Huberti, Chris Kolondra, Jennifer Lubahn, Lucie Sahner, Anna Katharina Schäfer, Anne-Luise Rieche, Margareta Marx, Stephanie Stieren February 13 to April 4, 2013
Gekrümmte Wirklichkeit: Bögen zwischen Kunst und Wissenschaft
Organized by Maxwell Roberts, Stephen Kobourov and Martin Nöllenburg as part of Dagstuhl Seminar 13151 “Drawing Graphs and Maps with Curves” April 8 to 21, 2013
Mechthild Schneider Fotografie
Works by artist Mechthild Schneider April 22 to July 12, 2013
Peter Amici
Works by Peter Amici July 30 to October 2, 2013
Sabine Eisenbrand, «Sprünge»
Works by artist Sabine Eisenbrand October 8 to December 19, 2013

Fig. 10.3
Art exhibitions in 2013.

11

**Stiftung „Informatikzentrum
Schloss Dagstuhl“
*The Dagstuhl Foundation***

Zielsetzung

11.1

Aims

Die Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH ist als Förderer von Wissenschaft und Forschung als gemeinnützig anerkannt. Für den Betrieb ist Schloss Dagstuhl neben seinen eigenen Einnahmen von den Teilnehmern seiner Veranstaltungen von einer öffentlichen Förderung abhängig, die 2013 über 75 % der Ausgaben trug. Daher ist Schloss Dagstuhl zu Sparsamkeit verpflichtet und unterliegt bei seinen Ausgaben einer öffentlichen Kontrolle. Schloss Dagstuhl ist deshalb dankbar für Spenden, die die Gesellschaft freier, flexibler und kurzfristiger verwenden kann.

Die Gesellschaft verwaltet zusätzlich seit 1995 ein Sondervermögen mit dem Ziel, damit eine rechtsfähige Stiftung „Informatikzentrum Schloss Dagstuhl“ zu gründen, die die Arbeit von Schloss Dagstuhl unterstützt. Die zu gründende Stiftung wird sich an den Gesellschaftszielen von Schloss Dagstuhl orientieren, die im Gesellschaftsvertrag der GmbH festgeschrieben sind. Sie fördert

- die Informatikforschung auf international anerkanntem Niveau
- die interdisziplinäre Forschungsdiskussion und Forschungskoooperation
- den Forschungsnachwuchs durch dessen Einbeziehung in die Forschungsdiskussion und durch intensive Fortbildung
- das Wirksamwerden neuer Informatikentwicklungen durch wissenschaftliche Weiterbildung auf hohem fachlichen Niveau
- die Erschließung neuer Anwendungsfelder der Informatik
- den Wissenstransfer zwischen Forschung und Wirtschaft.

Die Förderung des Nachwuchses ist dabei ein besonderes Anliegen. Das Sondervermögen speist sich aus jährlichen Zuwendungen von Privatpersonen und Institutionen, sowie Einzelspenden. Im Jahr 2013 wuchs das Sondervermögen aus diesen Zuwendungen um 15 979 € an.

Schloss Dagstuhl is recognized in Germany as a scientific non-profit organization. For its operation it depends, in addition to income from fees, on public funding which for 2013 provided over 75 % of the necessary funds. As such Schloss Dagstuhl is closely monitored and therefore grateful for additional donations that it can manage more independently.

For collecting donations Schloss Dagstuhl established in 1995 the “Informatikzentrum Schloss Dagstuhl” fund which is wholly part of the Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH. Using this fund, Schloss Dagstuhl plans to establish a legally independent foundation that supports the aims of Schloss Dagstuhl as:

- computer science on an international level
- interdisciplinary cooperation and discussion in science
- support of junior researchers
- establishing of new trends in computer science by providing training
- unlocking new applications in computer science
- exchange between research and industry

Supporting young researchers is a special priority of the Dagstuhl Foundation, which it maintains through annual contributions from individuals and institutions, as well as one-time donations. In 2013 the endowment grew by 15,979 €.

12 **Organe und Gremien** *Organs and Bodies*

Struktur der Gesellschaft

12.1

Schloss Dagstuhl wird als eine gemeinnützige GmbH betrieben, deren Gesellschafter die Gesellschaft für Informatik e.V. (GI), die Universität des Saarlandes, die Technische Universität Darmstadt, die Technische Universität Kaiserslautern, das Karlsruher Institut für Technologie, die Johann Wolfgang Goethe-Universität Frankfurt am Main, die Universität Stuttgart und die Universität Trier sind. Weitere Gesellschafter sind drei international renommierte Forschungsgesellschaften: Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), Centrum Wiskunde & Informatica (CWI, Niederlande), und die Max-Planck-Gesellschaft (MPG, Deutschland).

Schloss Dagstuhl wurde durch Beschluss der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (BLK) 2005 als Serviceeinrichtung für die Forschung in die gemeinsame Forschungsförderung von Bund und Ländern aufgenommen. Es ist Mitglied der Leibniz-Gemeinschaft. Entsprechend wurde 2008 der Name des Zentrums von vormals „Internationales Begegnungs- und Forschungszentrum für Informatik“ in „Schloss Dagstuhl – Leibniz-Zentrum für Informatik“ geändert.

Im Juli 2009 wurde Dagstuhl erstmals durch die Leibniz-Gemeinschaft evaluiert. Die Stellungnahme der Evaluierungskommission vom März 2010 ergab ein positives Bild: Schloss Dagstuhl widmet sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen.

Organe und Gremien der Gesellschaft

12.2

Folgende sechs Organe und Gremien sind für die Aktivitäten von Schloss Dagstuhl verantwortlich.

■ Die Gesellschafterversammlung

Die Vertreter der Gesellschafter berufen die Mitglieder des Aufsichtsrates und sind zuständig für Änderungen im Gesellschaftsvertrag und die Aufnahme weiterer Gesellschafter, siehe Fig. 12.1.

■ Der Aufsichtsrat

Der Aufsichtsrat ist verantwortlich dafür, dass die Geschäftsführung die Ziele der Gesellschaft rechtmäßig, zweckmäßig und wirtschaftlich sinnvoll erfüllt. Er wirkt in allen wesentlichen Angelegenheiten der Gesellschaft betreffend Forschung und Finanzplanung mit. Die 12 Mitglieder des Aufsichtsrats setzen sich zusammen aus vier Repräsentanten der Gesellschaft für Informatik, je einem

Structure of the Center

Schloss Dagstuhl is operated as a non-profit organization whose associates include the Gesellschaft für Informatik e.V.³⁹ (GI), the Universität des Saarlandes, the Technischen Universität Darmstadt, the Technische Universität Kaiserslautern, the Karlsruher Institut für Technologie, the Johann Wolfgang Goethe-Universität Frankfurt am Main, the Universität Stuttgart and the Universität Trier. Other associates of the center are three research institutes of international renown: the Institut National de Recherche en Informatique et en Automatique (INRIA, France), the Centrum Wiskunde & Informatica (CWI, The Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the “Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung” (BLK)⁴⁰ in 2005, the center was included as a research service institution in the joint funding of the German federal and state governments. The center is a member of the Leibniz Association. Accordingly its name was changed from “Internationales Begegnungs- und Forschungszentrum für Informatik”⁴¹ to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik.”⁴²

Dagstuhl was evaluated for the first time in July of 2009 by the Leibniz Association. The findings of the Evaluation Commission of March 2010 showed a positive image and established that the center has shown outstanding commitment to its designated task of supporting the international informatics research community by providing a seminar center for academic events.

Dagstuhl Organs and Bodies

The following six organs and bodies are in charge of the activities offered by Schloss Dagstuhl.

■ Associates’ Meeting

The representatives of the Associates’ Meeting convene meetings of the Supervisory Board and are responsible for amendments to the articles of incorporation and the admission of other associates; see Fig. 12.1.

■ Supervisory Board

The Supervisory Board is responsible for ensuring that management complies with the center’s objectives in a meaningful legal and economic manner. It is involved in all essential matters regarding research and financial planning. The board with its 12 members is composed of four representatives of the German Informatics Society, one representative each of the three founding universities (Uni-

³⁹ engl.: German Informatics Society

⁴⁰ engl.: Federal Government–State Commission for Educational Planning and Research Promotion

⁴¹ engl.: International Conference and Research Center for Computer Science

⁴² engl.: Schloss Dagstuhl – Leibniz Center for Informatics.

Vertreter der drei Gründungsuniversitäten (Universität des Saarlandes, Karlsruher Institut für Technologie, Technische Universität Kaiserslautern), zwei Vertretern der später hinzugekommenen Universitäten (Technische Universität Darmstadt, Johann Wolfgang Goethe-Universität Frankfurt am Main, Universität Stuttgart, Universität Trier) und je einem Vertreter des Bundes und der beiden Sitzländer (Saarland und Rheinland-Pfalz). Die Amtszeit der Mitglieder des Aufsichtsrates beträgt vier volle abgeschlossene Geschäftsjahre. Der Aufsichtsrat beruft das Wissenschaftliche Direktorium sowie die Mitglieder des Wissenschaftlichen Beirates und des Industriellen Kuratoriums. Siehe Fig. 12.2.

■ Die Geschäftsführung

Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH hat zwei Geschäftsführer, die gemeinsam die Gesellschaft vertreten. Die Geschäftsführer sind der Wissenschaftliche Direktor Professor Dr. Dr. h.c. Reinhard Wilhelm und der Technisch-administrative Geschäftsführer Dr. Christian Lindig.

■ Der Wissenschaftliche Beirat

Der Wissenschaftliche Beirat ist international besetzt und soll die Leitung des Zentrums hinsichtlich der wissenschaftlichen Ausrichtung sowie der Nutzerorientierung des Serviceangebotes kritisch begleiten und in grundlegenden Entscheidungen zur Weiterentwicklung unterstützen. Aufsichtsrat und Direktorium soll er in fachlich-wissenschaftlicher Hinsicht beraten. Zudem soll er die Leistungen des Zentrums bewerten und zwischen zwei Evaluierungen durch die Leibniz-Gemeinschaft (alle sieben Jahre) eine sogenannte interne Evaluierung (Audit) von Schloss Dagstuhl durchführen. Das Ergebnis dieser Evaluierung wird dem Senatsausschluss Evaluierung der Leibniz-Gemeinschaft übermittelt.

■ Das Industrielle Kuratorium

Das Industrielle Kuratorium (siehe Fig. 12.4) erfüllt eine Transmissionsfunktion zwischen Schloss Dagstuhl und den Forschungsabteilungen und Entwicklungslaboren der Industrie. Zudem hat es die Aufgabe, die Akzeptanz des Zentrums in Verwaltung, Industrie und Wirtschaft abzusichern und als Förderungsorganisation die wirtschaftliche Basis des Zentrums zu verbreitern. Die Mitglieder des Kuratoriums unterstützen das Zentrum dabei, aktuelle Themen zu identifizieren und dazu passende zugkräftige Organisatoren aus der Industrie zu gewinnen. Das Kuratorium wird regelmäßig aufgefordert, aus seinem Wirkungskreis passende Teilnehmer zu den Seminaren vorzuschlagen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

■ Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 12.5) ist für die Gestaltung des Seminarprogramms verantwortlich, begutachtet die Anträge auf Dagstuhl-Seminare und Dag-

stuhl des Saarlandes, Karlsruher Institut für Technologie, Technische Universität Kaiserslautern), two representatives of the universities that subsequently joined (Technische Universität Darmstadt, Johann Wolfgang Goethe-Universität Frankfurt am Main, Universität Stuttgart, Universität Trier), and one representative each of the federal government and the two host state governments (Saarland and Rhineland-Palatinate). The members of the Supervisory Board hold office for four full fiscal years. The Supervisory Board convenes meetings of the Scientific Directorate and of members of the Scientific Advisory Board and the Industrial Curatory Board. See Fig. 12.2.

■ The Management

Scientific Director Professor Dr. Dr. h.c. Reinhard Wilhelm and Technical Administrative Director Dr. Christian Lindig serve as the joint representatives of Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH.

■ Scientific Advisory Board

The Scientific Advisory Board is an internationally diverse body. The purpose of the board is to lend critical support to the management of the center with regard to its scientific orientation and the user orientation of its service offerings, and in policy decisions pertaining to the center's continued development, by advising the Supervisory Board and Scientific Directorate in a scientific or subject-matter capacity. Another task is to evaluate the center's performance and achievements (the so-called internal audit) between two evaluations, carried out by the Leibniz Association every seven years. The result of the internal audit is sent to the Senate Evaluation Committee of the Leibniz Association.

■ Industrial Curatory Board

The Industrial Curatory Board (see Fig. 12.4) performs a transmission function between the center and the R&D departments and industry laboratories. It also helps to secure the center's acceptance by government authorities and industry and, being a promotional organization, works to expand Schloss Dagstuhl's economic base. The members of the Curatory Board help the center to identify current R&D topics for seminars and locate attractive organizers in industry. The Curatory Board is regularly called upon to propose suitable participants for seminars known to it from its activities. The Industrial Curatory Board convenes once a year together with the Scientific Advisory Board.

■ Scientific Directorate

The Scientific Directorate (see Fig. 12.5) is responsible for the center's seminar program. It reviews the proposals for Dagstuhl Seminars and the Dagstuhl Perspec-

stuhl-Perspektiven-Workshops und entscheidet über ihre Annahme. Es behält sich vor, sowohl auf die Fokussierung des Themas als auch auf die Zusammensetzung des Teilnehmerkreises Einfluss zu nehmen. Außerdem gibt das Direktorium Anregungen zu Seminarthemen, wenn einzelne Gebiete der Informatik nicht gut vertreten sind, und plant neue Veranstaltungskonzepte. Das Direktorium setzt sich zusammen aus jeweils einem oder einer von den Gesellschafteruniversitäten und -forschungsinstituten entsandten Professor oder Professorin für Informatik, sowie vier Delegierten der GI. Von diesen werden zwei vom GI-Präsidium und zwei von dem vom Präsidium unabhängigen GI-Beirat der Universitätsprofessoren (GIBU) nominiert. Das Direktorium hat insgesamt 14 Mitglieder. Die Amtszeiten der Mitglieder und des Direktors betragen drei Jahre.

Der Aufsichtsrat hat in seiner Sitzung im Oktober 2013 beschlossen, dass die dreijährigen Amtszeiten der Direktoren am 1. November des Jahres ihrer Ernennung beginnen und im allgemeinen am 30. Oktober drei Jahre später enden. Um einen nahtlosen Übergang von den alten zu den neuen Amtszeiten zu gewährleisten, wurde die Amtszeit der Mitglieder des Direktoriums einmalig um fünf Monate verlängert.

Die Mitglieder wählen aus ihrer Mitte den wissenschaftlichen Direktor. Das Amt wird seit Bestehen des Zentrums von Professor Dr. Reinhard Wilhelm wahr genommen.

tives Workshops and decides whether they merit approval, reserving the right to approve the focus of topics and the individuals included in the participant group. It also makes recommendations to the Scientific Directorate concerning seminar topics when individual informatics fields are not well represented, and develops new event concepts. The Scientific Directorate is composed of one informatics professor delegated from each of the university and research center associate members, and four GI delegates. Of these individuals, two are nominated by the GI Executive Board and two by the GI Advisory Board of University Professors (GIBU), which is independent of the Executive Board. The Scientific Directorate is composed of 14 members in total. Each member holds office for three years, as does the Scientific Director.

In its October 2013 meeting, the Supervisory Board decided to implement the following change with respect to the Scientific Directorate appointments: The term of office begins on November 1 of the year of election, and generally ends on October 31 three years later. To guarantee a smooth transition between the old and the new regulations, the term of office of all members is subject to a one-time extension of five months.

The members elect a Scientific Director from their midst. Professor Dr. Reinhard Wilhelm has been the center's Scientific Director since its founding.

Gesellschafter Associates
Gesellschaft für Informatik e.V., Germany
Universität des Saarlandes, Germany
Technische Universität Kaiserslautern, Germany
Karlsruher Institut für Technologie (KIT), Germany
Technische Universität Darmstadt, Germany
Universität Stuttgart, Germany
Universität Trier, Germany
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Centrum Wiskunde & Informatica (CWI), Netherlands
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V., Berlin

Fig. 12.1

Associates

Vertreter der Gesellschafter Representatives of the Associates
Prof. Alejandro P. Buchmann, Ph.D. Technische Universität Darmstadt, Germany Representative of Technische Universität Darmstadt
Dr. Peter Federer Gesellschaft für Informatik e.V., Bonn, Germany Representative of Gesellschaft für Informatik e.V.
Prof. Oliver Günther, Ph.D. Universität Potsdam, Germany Representative of Gesellschaft für Informatik e.V.
Prof. Dr.-Ing. Dr. h.c. Theo Härder Technische Universität Kaiserslautern, Germany Representative of Technische Universität Kaiserslautern
Prof. Dr.-Ing. Dr. h.c. Stefan Jähnichen Technische Universität Berlin, Germany Representative of Gesellschaft für Informatik e.V. Chairman of the Supervisory Board and the Associates' Meeting
Prof. Dr. Volker Linneweber Universität des Saarlandes, Germany Representative of Universität des Saarlandes
Prof. Dr. Erhard Plödereder Universität Stuttgart, Germany Representative of Universität Stuttgart
Prof. Dr. Peter H. Schmitt Karlsruher Institut für Technologie, Germany Representative of Karlsruher Institut für Technologie
Prof. em. Dr.-Ing. Dr.-Ing. h.c. Roland Vollmar Karlsruher Institut für Technologie, Germany Representative of Gesellschaft für Informatik e.V.

Vertreter des Bundes und der Länder Representatives of the German federal government and states
Dr. Doreen Becker Bundesministerium für Bildung und Forschung, Bonn, Germany Representative of the German federal government
Wolfgang Habelitz Ministeriums für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany Representative of the Rhineland-Palatinate
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany Representative of the Saarland

Fig. 12.2

Supervisory Board members

Wissenschaftlicher Beirat Scientific Advisory Board
Prof. Dr. Manuel V. Hermenegildo IMDEA Software Institute, Madrid, and Technical University of Madrid, Spain
Prof. Dr. Claude Kirchner Institut National de Recherche en Informatique et en Automatique (INRIA), France
Prof. Dr. Mila E. Majster-Cederbaum Ludwig-Maximilians-Universität München, Germany Chairwoman of the Scientific Advisory Board <i>since May, 2013</i>
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institut and Computer Science Departement, Universität Paderborn, Germany <i>tenure started on May, 2013</i>
Prof. Dr.-Ing. Dr. h.c. Andreas Reuter HITS GmbH, Heidelberg, Germany
Prof. em. Dr. Dr. h.c. Otto Spaniol RWTH Aachen, Germany
Prof. Dr. Dorothea Wagner Karlsruher Institut für Technologie, Germany Chairwoman of the Scientific Advisory Board <i>until May 2013</i> <i>tenure ended on May, 2013</i>
Dr. Anne Norekian Staatskanzlei des Saarlandes, Saarbrücken, Germany <i>Guest</i>

Fig. 12.3

Scientific Advisory Board

Industrielles Kuratorium Industrial Curatory Board
Dr. Udo Bub EIT ICT Labs, Berlin, Germany
Dr. Jorge R. Cuéllar Siemens AG, München, Germany
Dr.-Ing. Elmar Dörner SAP Research, Karlsruhe, Germany
Dr. Jo Ebergen Oracle Labs, United States
Dr. Goetz Graefe HP Labs, United States
Prof. Dr. Ralf Guido Herrtwich Daimler AG, Böblingen, Germany
Prof. Dr. Thomas Hofmann ETH Zürich, Switzerland (<i>since 07/2013 at ETH Zürich; previously at Google Research, Zürich, Switzerland</i>)
Prof. Dr. Ulrich Lauther Siemens AG, München, Germany
Prof. Dr. Prabhakar Raghavan Google Inc. and Consulting Professor at Stanford University, United States
Prof. Dr.-Ing. Dr. h.c. Andreas Reuter HITS GmbH, Heidelberg, Germany
Prof. Dr. Frank Tip David R. Cheriton School of Computer Science, University of Waterloo, Ontario, Canada (<i>since 09/2012 at University of Waterloo; before at IBM T.J. Watson Research Center, Hawthorne, United States</i>) <i>tenure ended on May 2013</i>
Prof. Dr. Volker Tresp Siemens AG, München, Germany and Ludwig-Maximilians-Universität München, Germany

Fig. 12.4

Industrial Curatory Board

Wissenschaftliches Direktorium Scientific Directorate
Prof. Dr. Susanne Albers Humboldt-Universität, Berlin, Germany Delegate of Gesellschaft für Informatik e.V. (GIBU)
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany Delegate of Gesellschaft für Informatik e.V. (GIBU)
Prof. Dr. Karsten Berns Technische Universität Kaiserslautern, Germany Delegate of Technische Universität Kaiserslautern
Prof. Dr. Stefan Diehl Universität Trier, Germany Delegate of Universität Trier
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany Delegate of Karlsruher Institut für Technologie
Prof. Dr. Han La Poutré Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands Delegate of Centrum voor Wiskunde en Informatica (CWI)
Dr. Stephan Merz Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy, France Delegate of INRIA
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany Delegate of Universität Stuttgart
Prof. Dr. Bernhard Nebel Albert-Ludwigs-Universität Freiburg, Germany Delegate of Gesellschaft für Informatik e.V. (GI-Präsidium)
Prof. Dr. Bernt Schiele Max-Planck-Institut für Informatik, Saarbrücken, Germany Delegate of Max-Planck-Gesellschaft
Prof. Dr. Nicole Schweikardt Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany Delegate of Johann Wolfgang Goethe-Universität Frankfurt am Main
Prof. Dr. Raimund Seidel Universität des Saarlandes, Germany Delegate of Gesellschaft für Informatik e.V. (GI-Präsidium)
Prof. Dr. Michael Waidner Technische Universität Darmstadt, Germany Delegate of Technische Universität Darmstadt
Prof. Dr. Dr. h.c. Dr. h.c. Reinhard Wilhelm Universität des Saarlandes, Germany Delegate of Universität des Saarlandes Scientific Director of Schloss Dagstuhl
Members-at-Large
Prof. Dr. Friedemann Mattern ETH Zürich, Switzerland
Prof. Dr. Luca Benini ETH Zürich, Switzerland and University of Bologna, Italy
Prof. em. Dr. Jan-Olof Eklundh Royal Institute of Technology, Stockholm, Sweden
Prof. Dr. David Notkin University of Washington, United States † April 22, 2013

Fig. 12.5
Scientific Directorate

13 Statistik 2013 *Statistics 2013*

In diesem Kapitel werden statistische Daten zum wissenschaftlichen Programm und der Zusammenstellung der Teilnehmer aufgeführt.

Teilnehmer-bezogene Daten: Fig. 13.1 zeigt die Verteilung der Herkunftsländer unserer Gäste. Die Anzahl von früheren Seminarbesuchen kann man Fig. 13.2 entnehmen. Fig. 13.3 gibt Auskunft über die Altersstruktur der Teilnehmer.

Veranstaltungs-bezogene Daten: Daten zu der Anzahl unserer Veranstaltungen sind in Fig. 13.5 angegeben. Fig. 13.4 zeigt die Verteilung in Bezug auf kleine und große Seminare. Teilnehmerzahlen abhängig vom Veranstaltungstyp gibt Fig. 13.6 an. Schlussendlich findet man in Fig. 13.7 Zahlen zu den Gasttagen.

Antrags-bezogene Daten: Die Akzeptanzraten für eingereichte Anträge sind in Fig. 13.8 dargestellt.

Gender-bezogene Daten: Fig. 13.9 enthält Daten zur Gender-Komposition der Seminarleitung. Die Abbildungen Fig. 13.10 und Fig. 13.11 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler.

This chapter provides statistical data about the scientific program and its composition with regard to participants.

Participant-related data: Fig. 13.1 shows the distribution of country affiliations. Fig. 13.2 displays how often participants have attended a seminar before. Fig. 13.3 gives data about the seniority of participants.

Event-related data: Fig. 13.5 provides data about the number of events and Fig. 13.4 shows the distribution with regard to large and small seminars. Fig. 13.6 shows the number of participants according to the event type. Finally, Fig. 13.7 states the number of guest days.

Proposal-related data: Fig. 13.8 shows acceptance rates for the recent years.

Gender-related data: Fig. 13.9 shows mixed-gender data. In Fig. 13.10 and Fig. 13.11 data is given with regard to female participants and invitees, respectively.


Country	A	B	Total
Germany	653	646	1299
United States	488	30	518
United Kingdom	240	13	253
France	234	3	237
Netherlands	110	20	130
Switzerland	77	26	103
Canada	88	2	90
Italy	86	3	89
Austria	72	13	85
Belgium	41	25	66
Israel	49	4	53
Australia	51	1	52
Sweden	35	15	50
Denmark	38	9	47
Japan	41	5	46
Spain	39	6	45
Finland	36	0	36
Norway	32	3	35
Poland	21	5	26
India	21	2	23
Portugal	23	0	23
Ireland	19	1	20
China	17	0	17
Greece	11	4	15
Brazil	14	0	14
Hungary	12	1	13
Singapore	11	0	11
Russian Federation	10	0	10
Czech Republic	6	3	9
Pakistan	0	9	9
Estonia	7	1	8
Republic of Korea	8	0	8
Turkey	6	2	8
Luxembourg	6	1	7
Lithuania	0	6	6
South Africa	5	1	6
Chile	5	0	5
Slovenia	4	0	4
Argentina	3	0	3
Hong Kong	3	0	3
New Zealand	3	0	3
Serbia	0	3	3
United Arab Emirates	3	0	3
Colombia	2	0	2
Slovak Republic	2	0	2
Bulgaria	1	0	1
Cyprus	1	0	1
Latvia	1	0	1
Morocco	1	0	1
Oman	1	0	1
Saudi Arabia	1	0	1
Sudan	0	1	1
Thailand	1	0	1
Total	2639	864	3503

(a)Details by country


Fig. 13.1
Number of Dagstuhl guests by country of origin in 2013. A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and other events).

Region	A		B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	1170	44.3	163	18.9	1333	38.1
Germany	653	24.7	646	74.8	1299	37.1
North America	576	21.8	32	3.7	608	17.4
Asia	156	5.9	20	2.3	176	5
Australia	54	2	1	0.1	55	1.6
South America	24	0.9	0	0	24	0.7
Africa	6	0.2	2	0.2	8	0.2
Total	2639	100	864	100	3503	100

(b)Details by region



(c)Graphical distribution by region




Previous visits	Participants	
	#	%
2008–2013		
0	1444	54.7
1	572	21.7
2	316	12
3	140	5.3
4	68	2.6
5	57	2.2
≥ 6	42	1.6

(a) Graphical distribution

(b) Detailed numbers

Fig. 13.2

Dagstuhl participants in 2013 and their previous attendances in a Dagstuhl Seminar or Dagstuhl Perspectives Workshop from 2008 to 2013.




Year	Junior (%)	Senior (%)	Neither (%)
2008	23.8	55.5	20.7
2009	25.2	52.9	21.9
2010	28.9	49.7	21.4
2011	27.9	51.2	20.9
2012	27.6	52.1	20.3
2013	28.9	52.8	18.3

(a) Graphical distribution

(b) Detailed numbers

Fig. 13.3

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants



Year	Small	Large	Total
2008	14	38	52
2009	22	38	60
2010	19	40	59
2011	21	34	55
2012	22	42	64
2013	34	41	75

(a) Graphical distribution

(b) Detailed numbers

Fig. 13.4

Small vs. large Dagstuhl Seminars and Dagstuhl Perspectives Workshops. Small = 30-person seminar, Large = 45-person seminar.

Year	DS	PW	GI	EDU	OE	Total
2008	45	7	1	6	50	109
2009	53	7	1	4	36	101
2010	55	4	1	6	39	105
2011	53	2	0	3	35	93
2012	59	5	2	4	52	122
2013	74	1	0	5	33	113

Fig. 13.5 **Types of events at Dagstuhl.** DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl-Seminar, EDU = educational event, OE = other event.

Year	DS		PW		GI		EDU		OE		Total #
	#	%	#	%	#	%	#	%	#	%	
2008	1622	55.7	179	6.1	32	1.1	166	5.7	912	31.3	2911
2009	1983	65.9	185	6.1	26	0.9	131	4.4	686	22.8	3011
2010	1950	64.7	103	3.4	25	0.8	192	6.4	743	24.7	3013
2011	1894	70.2	64	2.4	0	0.0	103	3.8	637	23.6	2698
2012	2226	64.4	120	3.5	48	1.4	144	4.2	916	26.5	3454
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503


Fig. 13.6 **Number of participants by event type and year.** DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl-Seminar, EDU = educational event, OE = other event.

Year	DS	PW	GI	EDU	OE	Total
2008	7309	525	109	379	2206	10528
2009	8717	657	77	378	1776	11605
2010	8572	381	125	722	2002	11802
2011	8415	228	0	266	1604	10513
2012	9798	458	190	393	2031	12870
2013	11612	130	0	741	1614	14097

Fig. 13.7 **Number of guest days at Dagstuhl.** DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl-Seminar, EDU = educational event, OE = other event.

Year	Proposals		Accepted		Rejected	
	#	%	#	%	#	%
2008	83	60	72.3	23	27.7	
2009	95	68	71.6	27	28.4	
2010	94	65	69.1	29	30.9	
2011	80	54	67.5	26	32.5	
2012	90	69	76.7	21	23.3	
2013	107	72	67.3	35	32.7	

Fig. 13.8 **Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals and acceptance rates**



(a) Graphical distribution

Year	Teams	Organizers	Mixed Teams		Women	
	#	#	#	%	#	%
2008	52	200	27	51.9	31	15.5
2009	60	228	20	33.3	20	8.8
2010	59	233	32	54.2	34	14.6
2011	55	213	27	49.1	31	14.6
2012	64	256	32	50.0	39	15.2
2013	75	282	36	48.0	43	15.2

(b) Detailed numbers

Fig. 13.9

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

Year	Participants	Female Participants	
	#	#	%
2008	1801	244	13.5
2009	2168	296	13.7
2010	2053	293	14.3
2011	1958	294	15.0
2012	2346	378	16.1
2013	2639	401	15.2

Fig. 13.10

Female participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops by year

Year	Invitees	Female Invitees		Decliners		Female Decliners	
	#	#	%	#	%	#	%
2008	4268	594	13.9	2467	57.8	350	58.9
2009	4671	648	13.9	2503	53.6	352	54.3
2010	4499	630	14.0	2446	54.4	337	53.5
2011	4223	600	14.2	2265	53.6	306	51.0
2012	5033	822	16.3	2687	53.4	444	54.0
2013	5591	889	15.9	2952	52.8	488	54.9

Fig. 13.11

Gender of Dagstuhl Seminar and Dagstuhl Perspectives Workshop invitees and decliners including their rates

14 **Veranstaltungen 2013** *Schedule of Events 2013*

Dagstuhl-Seminare**14.1****Dagstuhl Seminars****13021 – Symbolic Methods in Testing**

Thierry Jéron (INRIA Rennes – Bretagne Atlantique, FR), Margus Veanes (Microsoft Research – Redmond, US), Burkhart Wolff (University of Paris South XI, FR)

January 6–11, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13021>

13022 – Engineering Resilient Systems: Models, Methods and Tools

Nicolas Guelfi (University of Luxembourg, LU), Maritta Heisel (Universität Duisburg-Essen, DE), Mohamed Kaaniche (LAAS – Toulouse, FR), Alexander Romanovsky (Newcastle University, GB), Elena Troubitsyna (Abo Akademi University, FI)

January 6–11, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13022>

13031 – Computational Counting

Peter Bürgisser (Universität Paderborn, DE), Leslie Ann Goldberg (University of Liverpool, GB), Mark Jerrum (Queen Mary University of London, GB), Pascal Koiran (ENS – Lyon, FR)

January 13–18, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13031>

13041 – Civilian Crisis Response Models

Ozlem Ergun (Georgia Institute of Technology, US), Bernhard Katzy (Leiden University, NL), Ulrike Lechner (Universität der Bundeswehr – München, DE), Luk van Wassenhove (INSEAD – Fontainebleau, FR)

January 20–25, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13041>

13042 – Epidemic Algorithms and Processes: From Theory to Applications

Benjamin Doerr (MPI für Informatik – Saarbrücken, DE), Robert Elsässer (Universität Salzburg, AT), Pierre Fraigniaud (University Paris-Diderot, FR), Rachid Guerraoui (EPFL – Lausanne, CH)

January 20–25, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13042>

13051 – Software Certification: Methods and Tools

Darren Cofer (Rockwell Collins – Cedar Rapids, US), John Hatchiff (Kansas State University, US), Michaela Huhn (TU Clausthal, DE), Mark Lawford (McMaster University – Hamilton, CA)

January 27 to February 1, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13051>

13052 – Multicore Enablement for Embedded and Cyber Physical Systems

Andreas Herkersdorf (TU München, DE), Michael G. Hinchey (University of Limerick, IE), Michael Paulitsch (EADS Deutschland – München, DE)

January 27 to February 1, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13052>

13061 – Fault Prediction, Localization, and Repair

Mary Jean Harrold (Georgia Institute of Technology, US), Friedrich Steimann (Fernuniversität in Hagen, DE), Frank Tip (University of Waterloo, CA), Andreas Zeller (Universität des Saarlandes, DE)

February 3–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13061>

13062 – Decentralized Systems for Privacy Preservation

Sonja Buchegger (KTH Royal Institute of Technology, SE), Jon Crowcroft (University of Cambridge, GB), Balachander Krishnamurthy (AT&T Labs Research – Florham Park, US), Thorsten Strufe (TU Darmstadt, DE)

February 3–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13062>

13071 – Dependence Logic: Theory and Applications

Samson Abramsky (University of Oxford, GB), Juha Kontinen (University of Helsinki, FI), Jouko Väänänen (University of Helsinki, FI & University of Amsterdam, NL), Heribert Vollmer (Leibniz Universität Hannover, DE)

February 10–15, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13071>

13072 – Mechanisms of Ongoing Development in Cognitive Robotics

Jacqueline Fagard (Paris Descartes University, FR), Roderic A. Grupen (University of Massachusetts – Amherst, US), Frank Guerin (University of Aberdeen, GB), Norbert Krüger (University of Southern Denmark – Odense, DK)

February 10–15, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13072>

13081 – Consistency In Distributed Systems

Bettina Kemme (McGill University – Montreal, CA), Ganesan Ramalingam (Microsoft Research India – Bangalore, IN), André Schiper (EPFL – Lausanne, CH), Marc Shapiro (INRIA & LIP6 – Paris, FR), Kapil Vaswani (Microsoft Research India – Bangalore, IN)

February 17–22, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13081>

13082 – Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices

LeRoy B. Beasley (Utah State University, US), Hartmut Klauck (Nanyang TU – Singapore, SG), Troy Lee (National University of Singapore, SG), Dirk Oliver Theis (University of Tartu, EE)

February 17–22, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13082>

13091 – Analysis, Test and Verification in The Presence of Variability

Paulo Borba (University of Pernambuco – Recife, BR), Myra B. Cohen (University of Nebraska – Lincoln, US), Axel Legay (INRIA Rennes – Bretagne Atlantique, FR), Andrzej Wasowski (IT University of Copenhagen, DK)

February 24 to March 1, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13091>

13101 – Computational Geometry

Otfried Cheong (KAIST – Daejeon, KR), Kurt Mehlhorn (MPI für Informatik – Saarbrücken, DE), Monique Teillaud (INRIA Sophia Antipolis – Méditerranée, FR)

March 3–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13101>

13111 – Scheduling

Susanne Albers (HU Berlin, DE), Onno J. Boxma (TU Eindhoven, NL), Kirk Pruhs (University of Pittsburgh, US)

March 10–15, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13111>

13121 – Bidimensional Structures: Algorithms, Combinatorics and Logic

Erik D. Demaine (MIT – Cambridge, US), Fedor V. Fomin (University of Bergen, NO), MohammadTaghi Hajiaghayi (University of Maryland, US), Dimitrios M. Thilikos (National and Kapodistrian University of Athens, GR)

March 17–22, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13121>

13131 – Future Internet

Jon Crowcroft (University of Cambridge, GB), Markus Fidler (Leibniz Universität Hannover, DE), Klara Nahrstedt (University of Illinois – Urbana Champaign, US), Ralf Steinmetz (TU Darmstadt, DE)

March 24–27, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13131>

13141 – Formal Verification of Distributed Algorithms

Bernadette Charron-Bost (Ecole Polytechnique – Palaiseau, FR), Stephan Merz (LORIA – Nancy, FR), Andrey Rybalchenko (TU München, DE), Josef Widder (TU Wien, AT)

April 1–5, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13141>

13142 – Correct and Efficient Accelerator Programming

Albert Cohen (ENS – Paris, FR), Alastair F. Donaldson (Imperial College London, GB), Marieke Huisman (University of Twente, NL), Joost-Pieter Katoen (RWTH Aachen, DE)

April 1–4, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13142>

13151 – Drawing Graphs and Maps with Curves

Sara Fabrikant (Universität Zürich, CH), Stephen G. Kobourov (University of Arizona – Tucson, US), Martin Nöllenburg (KIT – Karlsruhe Institute of Technology, DE), Monique Teillaud (INRIA Sophia Antipolis – Méditerranée, FR)

April 7–12, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13151>

13161 – Interface of Computation, Game Theory, and Economics

Sergiu Hart (The Hebrew Univ. of Jerusalem, IL), Éva Tardos (Cornell University, US), Bernhard von Stengel (London School of Economics, GB)

April 14–19, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13161>

13162 – Pointer Analysis

Ondrej Lhotak (University of Waterloo, CA), Yannis Smaragdakis (University of Athens, GR), Manu Sridharan (IBM TJ Watson Research Center – Yorktown Heights, US)

April 14–19, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13162>

13171 – Customizing Service Platforms

Luciano Baresi (Technical University of Milan, IT), Andreas Rummeler (SAP Research Center – Dresden, DE), Klaus Schmid (Universität Hildesheim, DE)

April 21–26, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13171>

13181 – VaToMAS – Verification and Testing of Multi-Agent Systems

Alessio R. Lomuscio (Imperial College London, GB), Sophie Pinchinat (University of Rennes, FR), Holger Schlingloff (HU Berlin, DE)

April 28 to May 3, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13181>

13182 – Meta-Modeling Model-Based Engineering Tools

Tony Clark (Middlesex University, GB), Robert B. France (Colorado State University, US), Martin Gogolla (Universität Bremen, DE), Bran V. Selic (Malina Software Corp. – Nepean, CA)
 April 28 to May 3, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13182>

13192 – Tree Transducers and Formal Methods

Sebastian Maneth (NICTA & University of New South Wales, Sydney, AU), Helmut Seidl (TU München, DE)
 May 5–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13192>

13201 – Information Visualization – Towards Multivariate Network Visualization

Andreas Kerren (Linnaeus University – Växjö, SE), Helen C. Purchase (University of Glasgow, GB), Matthew O. Ward (Worcester Polytechnic Institute, US)
 May 12–17, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13201>

13211 – Automated Reasoning on Conceptual Schemas

Diego Calvanese (Free University of Bozen-Bolzano, IT), Sven Hartmann (TU Clausthal, DE), Ernest Teniente (UPC – Barcelona, ES)
 May 19–24, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13211>

13212 – Computational Methods Aiding Early-Stage Drug Design

Andreas Bender (University of Cambridge, GB), Hinrich Göhlmann (Janssen Pharmaceutica – Beerse, BE), Sepp Hochreiter (University of Linz, AT), Ziv Shkedy (Hasselt University – Diepenbeek, BE)
 May 19–24, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13212>

13231 – Belief Change and Argumentation in Multi-Agent Scenarios

Jürgen Dix (TU Clausthal, DE), Sven Ove Hansson (KTH Royal Institute of Technology, SE), Gabriele Kern-Isberner (TU Dortmund, DE), Guillermo R. Simari (National University of the South – Bahia Blanca, AR)
 June 2–7, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13231>

13232 – Indexes and Computation over Compressed Structured Data

Sebastian Maneth (University of Oxford, GB), Gonzalo Navarro (University of Chile, CL)
 June 2–7, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13232>

13241 – Virtual Realities

Guido Brunnett (TU Chemnitz, DE), Sabine Coquillart (INRIA Rhône-Alpes, FR), Robert van Liere (CWI – Amsterdam, NL), Gregory F. Welch (University of Central Florida – Orlando, US)
 June 9–14, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13241>

13251 – Parallel Data Analysis

Artur Andrzejak (Universität Heidelberg, DE), Joachim Giesen (Universität Jena, DE), Raghu Ramakrishnan (Microsoft Research – Redmond, US), Ion Stoica (University of California – Berkeley, US)
 June 16–21, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13251>

13252 – Interoperation in Complex Information Ecosystems

Andreas Harth (KIT – Karlsruhe Institute of Technology, DE), Craig A. Knoblock (University of Southern California – Marina del Rey, US), Kai-Uwe Sattler (TU Ilmenau, DE), Rudi Studer (KIT – Karlsruhe Institute of Technology, DE)
 June 16–19, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13252>

13271 – Theory of Evolutionary Algorithms

Benjamin Doerr (MPI für Informatik – Saarbrücken, DE), Nikolaus Hansen (INRIA Saclay – Île-de-France – Orsay, FR), Jonathan L. Shapiro (University of Manchester, GB), L. Darrell Whitley (Colorado State University, US)
 June 30 to July 5, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13271>

13272 – Computer Science in High Performance Sport – Applications and Implications for Professional Coaching

Koen A.P.M. Lemmink (University of Groningen, NL), Stuart Morgan (Australian Institute of Sport – Bruce, AU), Jaime Sampaio (Universidade de Trás-os-Montes – Vila Real, PT), Dietmar Saupe (Universität Konstanz, DE)
 June 30 to July 3, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13272>

13311 – Duality in Computer Science

Mai Gehrke (University Paris-Diderot, FR), Jean-Eric Pin (University Paris-Diderot, FR), Victor Selivanov (A. P. Ershov Institute – Novosibirsk, RU), Dieter Spreen (Universität Siegen, DE)
 July 28 to August 2, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13311>

13312 – “My Life, Shared” – Trust and Privacy in the Age of Ubiquitous Experience Sharing

Alessandro Acquisti (Carnegie Mellon University, US), Ioannis Krontiris (Goethe-Universität Frankfurt am Main, DE), Marc Langheinrich (University of Lugano, CH), Martina Angela Sasse (University College London, GB)

July 28 to August 2, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13312>

13321 – Reinforcement Learning

Peter Auer (Montan-Universität Leoben, AT), Marcus Hutter (Australian National University, AU), Laurent Orseau (AgroParisTech – Paris, FR)

August 4–9, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13321>

13322 – The Critical Internet Infrastructure

Georg Carle (TU München, DE), Jochen Schiller (FU Berlin, DE), Steve Uhlig (Queen Mary University of London, GB), Matthias Wählisch (FU Berlin, DE), Walter Willinger (AT&T Labs Research – Florham Park, US)

August 4–9, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13322>

13331 – Exponential Algorithms: Algorithms and Complexity Beyond Polynomial Time

Thore Husfeldt (IT University of Copenhagen, DK), Ramamohan Paturi (University of California – San Diego, US), Gregory B. Sorkin (London School of Economics, GB), Ryan Williams (Stanford University, US)

August 11–16, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13331>

13341 – Verifiably Secure Process-Aware Information Systems

Rafael Accorsi (Universität Freiburg, DE), Jason Crampton (Royal Holloway University of London, GB), Michael Huth (Imperial College London, GB), Stefanie Rinderle-Ma (Universität Wien, AT)

August 18–23, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13341>

13351 – Coding Theory

Hans-Andrea Loeliger (ETH Zentrum – Zürich, CH), Emina Soljanin (Bell Labs – Murray Hill, US), Judy L. Walker (University of Nebraska – Lincoln, US)

August 25–30, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13351>

13352 – Interaction with Information for Visual Reasoning

David S. Ebert (Purdue University, US), Brian D. Fisher (Simon Fraser University – Surrey, CA), Petra Isenberg (INRIA Saclay – Île-de-France – Orsay, FR), Shixia Liu (Microsoft Research – Beijing, CN)

August 25–30, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13352>

13361 – Crowdsourcing: From Theory to Practice and Long-Term Perspectives

Claudio Bartolini (HP Labs – Palo Alto, US), Tobias Hoßfeld (Universität Würzburg, DE), Phuoc Tran-Gia (Universität Würzburg, DE), Maja Vukovic (IBM TJ Watson Research Center – Yorktown Heights, US)

September 1–4, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13361>

13362 – Cloud-based Software Crowdsourcing

Michael N. Huhns (University of South Carolina – Columbia, US), Wei Li (Beihang University – Beijing, CN), Wei-Tek Tsai (ASU – Tempe, US), Wenjun Wu (Beihang University – Beijing, CN)

September 1–4, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13362>

13371 – Quantum Cryptanalysis

Serge Fehr (CWI – Amsterdam, NL), Michele Mosca (University of Waterloo, CA), Martin Roetteler (Microsoft Research – Redmond, US), Rainer Steinwandt (Florida Atlantic University – Boca Raton, US)

September 8–13, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13371>

13372 – Integration of Tools for Rigorous Software Construction and Analysis

Uwe Glässer (Simon Fraser University – Burnaby, CA), Stefan Hallerstede (Aarhus University, DK), Michael Leuschel (Heinrich-Heine-Universität Düsseldorf, DE), Elvinia Riccobene (University of Milan, IT)

September 8–13, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13372>

13381 – Algorithms and Scheduling Techniques for Exascale Systems

Henri Casanova (University of Hawaii at Manoa, US), Yves Robert (ENS – Lyon, FR), Uwe Schwiegelshohn (TU Dortmund, DE)

September 15–20, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13381>

13382 – Collaboration and learning through live coding

Alan Blackwell (University of Cambridge, GB), Alex McLean (University of Leeds, GB), James Noble (Victoria University – Wellington, NZ), Julian Rohrer (Robert Schumann Hochschule für Musik, DE)

September 15–20, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13382>

13391 – Algorithm Engineering

Andrew V. Goldberg (Microsoft Research – Mountain View, US), Giuseppe F. Italiano (University of Rome “Tor Vergata”, IT), David S. Johnson (Madison, US), Dorothea Wagner (KIT – Karlsruhe Institute of Technology, DE)

September 22–27, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13391>

13392 – Inter-Vehicular Communication – Quo Vadis

Onur Altintas (TOYOTA InfoTechnology Center – Tokyo, JP), Falko Dressler (Universität Innsbruck, AT), Hannes Hartenstein (KIT – Karlsruhe Institute of Technology, DE), Ozan K. Tonguz (Carnegie Mellon University, US)

September 22–25, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13392>

13401 – Automatic Application Tuning for HPC Architectures

Siegfried Benkner (Universität Wien, AT), Franz Franchetti (Carnegie Mellon University, US), Hans Michael Gerndt (TU München, DE), Jeffrey K. Hollingsworth (University of Maryland – College Park, US)

September 29 to October 4, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13401>

13402 – Physical-Cyber-Social Computing

Payam M. Barnaghi (University of Surrey, GB), Ramesh Jain (University of California – Irvine, US), Amit P. Sheth (Wright State University – Dayton, US), Steffen Staab (Universität Koblenz-Landau, DE), Markus Strohmaier (Universität Koblenz-Landau, DE)

September 29 to October 4, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13402>

13411 – Deduction and Arithmetic

Nikolaj Björner (Microsoft Research – Redmond, US), Reiner Hähnle (TU Darmstadt, DE), Tobias Nipkow (TU München, DE), Christoph Weidenbach (MPI für Informatik – Saarbrücken, DE)

October 6–11, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13411>

13412 – Genomic Privacy

Kay Hamacher (TU Darmstadt, DE), Jean Pierre Hubaux (EPFL – Lausanne, CH), Gene Tsudik (University of California – Irvine, US)

October 6–9, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13412>

13421 – Algorithms for Optimization Problems in Planar Graphs

Glencora Borradaile (Oregon State University, US), Philip N. Klein (Brown University, US), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU), Claire Mathieu (Brown University, US)

October 13–18, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13421>

13422 – Nominal Computation Theory

Mikolaj Bojanczyk (University of Warsaw, PL), Bartek Klin (University of Warsaw, PL), Alexander Kurz (University of Leicester, GB), Andrew M. Pitts (University of Cambridge, GB)

October 13–16, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13422>

13431 – Real-World Visual Computing

Oliver Grau (Intel Visual Computing Institute – Saarbrücken, DE), Marcus A. Magnor (TU Braunschweig, DE), Olga Sorkine-Hornung (ETH Zürich, CH), Christian Theobalt (MPI für Informatik – Saarbrücken, DE)

October 20–25, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13431>

13441 – Evaluation Methodologies in Information Retrieval

Maristella Agosti (University of Padova, IT), Norbert Fuhr (Universität Duisburg-Essen, DE), Elaine Toms (Sheffield University, GB), Pertti Vakkari (University of Tampere, FI)

October 27 to November 1, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13441>

13451 – Computational Audio Analysis

Meinard Müller (Universität Erlangen-Nürnberg, DE), Shrikanth S. Narayanan (University of Southern California, US), Björn Schuller (TU München, DE)

November 3–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13451>

13452 – Proxemics in Human-Computer Interaction

Saul Greenberg (University of Calgary, CA), Kasper Hornbæk (University of Copenhagen, DK), Aaron Quigley (University of St. Andrews, GB), Roman Rädle (Universität Konstanz, DE), Harald Reiterer (Universität Konstanz, DE)

November 3–8, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13452>

13461 – Electronic Markets and Auctions

Yishay Mansour (Tel Aviv University, IL), Benny Moldovanu (Universität Bonn, DE), Noam Nisan (The Hebrew University of Jerusalem, IL), Berthold Vöcking (RWTH Aachen, DE)

November 10–15, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13461>

13462 – Computational Models of Language Meaning in Context

Hans Kamp (Universität Stuttgart, DE), Alessandro Lenci (University of Pisa, IT), James Pustejovsky (Brandeis University – Waltham, US)

November 10–15, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13462>

13471 – Synchronous Programming

Albert Benveniste (INRIA Rennes – Bretagne Atlantique, FR), Stephen A. Edwards (Columbia University – New York, US), Alain Girault (INRIA Grenoble – Rhône-Alpes, FR), Klaus Schneider (TU Kaiserslautern, DE)

November 17–22, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13471>

13472 – Global Measurement Framework

Philip Eardley (British Telecom R&D – Ipswich, GB), Marco Mellia (Polytechnic University of Torino, IT), Jörg Ott (Aalto University, FI), Jürgen Schönwälder (Jacobs Universität – Bremen, DE), Henning Schulzrinne (Columbia University – New York, US)

November 17–20, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13472>

13481 – Unleashing Operational Process Mining

Rafael Accorsi (Universität Freiburg, DE), Ernesto Damiani (Università degli Studi di Milano – Crema, IT), Wil van der Aalst (TU Eindhoven, NL)

November 24–29, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13481>

13482 – Forensic Computing

Felix C. Freiling (Universität Erlangen-Nürnberg, DE), Gerrit Hornung (Universität Passau, DE), Radim Polcák (Masaryk University, CZ)

November 24–29, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13482>

13491 – Computational Mass Spectrometry

Rudolf Aebersold (ETH Zürich, CH), Oliver Kohlbacher (Universität Tübingen, DE), Olga Vitek (Purdue University, US)

December 1–6, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13491>

13492 – Geosensor Networks: Bridging Algorithms and Applications

Matt Duckham (The University of Melbourne, AU), Stefan Dulman (TU Delft, NL), Jörg-Rüdiger Sack (Carleton University – Ottawa, CA), Monika Sester (Leibniz Universität Hannover, DE)

December 1–6, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13492>

13502 – Approaches and Applications of Inductive Programming

Sumit Gulwani (Microsoft Research – Redmond, US), Emanuel Kitzelmann (Universität Duisburg – Essen, DE), Ute Schmid (Universität Bamberg, DE)

December 8–11, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13502>

13511 – Software Engineering for Self-Adaptive Systems: Assurances

Rogério de Lemos (University of Kent, GB), David Garlan (Carnegie Mellon University, US), Carlo Ghezzi (Technical University of Milan, IT), Holger Giese (Hasso-Plattner-Institut – Potsdam, DE)

December 15–19, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13511>

13512 – Social Issues in Computational Transportation Science

Glenn Geers (NICTA – Kensington, AU), Monika Sester (Leibniz Universität Hannover, DE), Stephan Winter (The University of Melbourne, AU), Ouri E. Wolfson (University of Chicago, US)

December 15–19, 2013 | Dagstuhl Seminar | <http://www.dagstuhl.de/13512>

Dagstuhl-Perspektiven- Workshops

14.2

Dagstuhl Perspectives Workshops

13342 – ICT Strategies for Bridging Biology and Precision Medicine

Jonas Almeida (University of Alabama – Birmingham, US), Andreas Dress (Shanghai Institutes for Biological Sciences, CN & infinity3, DE), Titus Kühne (Deutsches Herzzentrum, DE), Laxmi Parida (IBM TJ Watson Research Center – Yorktown Heights, US)

August 18–23, 2013 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/13342>

GI-Dagstuhl-Seminare**14.3****GI-Dagstuhl Seminars**

Schloss Dagstuhl did not host any GI-Dagstuhl Seminars in 2013.

Lehrveranstaltungen**14.4****Educational Events****13152 – Summer School “Implementation Techniques for Data Management Software”**

Goetz Graefe (HP Labs – Madison, US), Wey Guy (Redmond, US), Harumi Anne Kuno (HP Labs – Palo Alto, US), Thomas Neumann (TU München, DE)

April 7–12, 2013 | Meeting | <http://www.dagstuhl.de/13152>

13242 – Workshop Wissenschaftsjournalismus

Roswitha Bardohl (Schloss Dagstuhl, DE), Gordon Bolduan (Universität des Saarlandes, DE), Tim Schröder (Oldenburg, DE)

June 9–12, 2013 | Meeting | <http://www.dagstuhl.de/13242>

13262 – 2nd Summer School in Computational Oncology

Norbert Graf (Universitätsklinikum des Saarlandes, DE)

June 23–28, 2013 | Meeting | <http://www.dagstuhl.de/13262>

13282 – Third European Business Intelligence Summer School (eBISS 2013)

Esteban Zimanyi (Université Libre de Bruxelles, BE)

July 7–12, 2013 | Meeting | <http://www.dagstuhl.de/13282>

13503 – Lehrerfortbildung in Informatik

Roswitha Bardohl (Schloss Dagstuhl, DE), Manuel Garcia Mateos (LPM Saarbrücken, DE), Martin Zimmol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)

December 11–13, 2013 | Meeting | <http://www.dagstuhl.de/13503>

Sonstige Veranstaltungen**14.5****Other Events****13032 – Erneuerbare Mobilität**

Karl-Heinz Krempels (RWTH Aachen, DE), Christoph Terwelp (RWTH Aachen, DE)

January 13–16, 2013 | Meeting | <http://www.dagstuhl.de/13032>

13034 – Retreat Forschungsbereich Agenten und Simulierte Realität: “Besprechungen produktiver gestalten”

Philipp Slusallek (DFKI – Saarbrücken, DE)

January 18, 2013 | Meeting | <http://www.dagstuhl.de/13034>

13043 – Workshop “Annotation and Alignment of Parallel Corpora for Linguistic Research”

Gintare Grigonyte (Universität Zürich, CH & Vytautas Magnus University – Kaunas, LT), Ruta Marcinkeviciene (Vytautas Magnus University – Kaunas, LT), Andrius Utka (Vytautas Magnus University – Kaunas, LT), Martin Volk (Universität Zürich, CH)

January 22–25, 2013 | Meeting | <http://www.dagstuhl.de/13043>

13044 – Offsite Meeting Commercial Performance Management

Thomas In der Rieden (T-Systems International GmbH, DE)

January 20–22, 2013 | Meeting | <http://www.dagstuhl.de/13044>

13073 – Kunstprojekt “HBK Saarbrücken”

Ingeborg Knigge (HBKS – Saarbrücken, DE), Gabriele Langendorf (HBKS – Saarbrücken, DE)

February 10–14, 2013 | Meeting | <http://www.dagstuhl.de/13073>

13083 – Lehrstuhltreffen Hanebeck

Uwe D. Hanebeck (KIT – Karlsruhe Institute of Technology, DE)

February 20–22, 2013 | Meeting | <http://www.dagstuhl.de/13083>

13092 – FOSD-Treffen

Sven Apel (Universität Passau, DE), Christian Kästner (Carnegie Mellon University, US), Christian Lengauer (Universität Passau, DE), Janet Siegmund (Universität Magdeburg, DE)

February 26 to March 1, 2013 | Meeting | <http://www.dagstuhl.de/13092>

13093 – Quantum Information: Theory & Implementation

Jürgen Eschner (Universität des Saarlandes, DE), Jörg Hettel (FH Kaiserslautern-Zweibrücken, DE),
Hans-Jürgen Steffens (FH Kaiserslautern-Zweibrücken, DE)

February 24–26, 2013 | Meeting | <http://www.dagstuhl.de/13093>

13102 – Clusterseminar “Intelligente Systeme zur Entscheidungsunterstützung”

Lars Mönch (FernUniversität in Hagen, DE)

March 4–6, 2013 | Meeting | <http://www.dagstuhl.de/13102>

13112 – Counting and Enumerating of Plane Graphs

Kevin Buchin (TU Eindhoven, NL), André Schulz (Universität Münster, DE), Csaba D. Tóth (University
of Calgary, CA)

March 10–13, 2013 | Meeting | <http://www.dagstuhl.de/13112>

13113 – Lehrstuhltreffen “Embedded Intelligence”

Bernhard Sick (Universität Kassel, DE)

March 13–15, 2013 | Meeting | <http://www.dagstuhl.de/13113>

13122 – Klausurtagung “LST Rannenberg”

Kai Rannenberg (Goethe-Universität Frankfurt am Main, DE)

March 17–20, 2013 | Meeting | <http://www.dagstuhl.de/13122>

13124 – Lehrstuhltreffen AG Zeller

Andreas Zeller (Universität des Saarlandes, DE)

March 20–22, 2013 | Meeting | <http://www.dagstuhl.de/13124>

13132 – GIBU 2013: GI-Beirat der Universitätsprofessoren

Gregor Snelting (KIT – Karlsruhe Institute of Technology, DE)

March 24–27, 2013 | Meeting | <http://www.dagstuhl.de/13132>

13172 – NSF/SRC/DFG Joint Workshop on “Bugs and Defects in Electronic Systems: the Next Frontier”

Wolfgang Kunz (TU Kaiserslautern, DE), Subhasish Mitra (Stanford University, US)

April 21–24, 2013 | Meeting | <http://www.dagstuhl.de/13172>

13173 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)

Bernhard Schätz (fortiss GmbH – München, DE)

April 24–25, 2013 | Meeting | <http://www.dagstuhl.de/13173>

13193 – Kolloquium zum GI Dissertationspreis 2012

Steffen Hölldobler (TU Dresden, DE)

May 5–8, 2013 | Meeting | <http://www.dagstuhl.de/13193>

13202 – Klausurtagung CCS@BTH

Markus Fiedler (Blekinge Institute of Technology – Karlskrona, SE)

May 12–15, 2013 | Meeting | <http://www.dagstuhl.de/13202>

13203 – CELSTEC Retreat Meeting

Rob Koper (Open University – Heerlen, NL)

May 12–17, 2013 | Meeting | <http://www.dagstuhl.de/13203>

13204 – Lehrstuhltreffen AG Schneider / Sturm

Peter Sturm (Universität Trier, DE)

May 15–17, 2013 | Meeting | <http://www.dagstuhl.de/13204>

13222 – Gemeinsamer Workshop der Graduiertenkollegs: GRK 1362 und GRK 1564

Julian Bader (Universität Siegen, DE), Rodrigo do Carmo (TU Darmstadt, DE), Christian Feinen
(Universität Siegen, DE), Jens Hedrich (Universität Koblenz-Landau, DE), Andreas Kolb (Universität
Siegen, DE), Philipp M. Scholl (TU Darmstadt, DE), Oskar von Stryk (TU Darmstadt, DE)

May 26–29, 2013 | Meeting | <http://www.dagstuhl.de/13222>

13243 – Deutsch-Pakistanischer Workshop

Karsten Berns (TU Kaiserslautern, DE)

June 13–14, 2013 | Meeting | <http://www.dagstuhl.de/13243>

13329 – Forschungsaufenthalt

Eike Best (Universität Oldenburg, DE)

August 9–17, 2013 | Meeting | <http://www.dagstuhl.de/13329>

13332 – Klausurtagung “LST Freiling”

Felix C. Freiling (Universität Erlangen-Nürnberg, DE)

August 11–15, 2013 | Meeting | <http://www.dagstuhl.de/13332>**13333 – Koordinationstreffen des Lehrstuhls Organic Computing der Universität Augsburg und des Fachgebiets Angewandte Informationssicherheit der Universität Kassel**

Jörg Hähner (Universität Augsburg, DE), Sven Tomforde (Universität Augsburg, DE), Arno Wacker (Universität Kassel, DE)

August 12–16, 2013 | Meeting | <http://www.dagstuhl.de/13333>**13363 – Optimierung von Diskriminierungsnetzwerken**

Karl-Heinz Krempels (RWTH Aachen, DE)

September 1–4, 2013 | Meeting | <http://www.dagstuhl.de/13363>**13365 – Offsite Meeting Commercial Performance Management**

Thomas In der Rieden (T-Systems International GmbH, DE)

September 1–3, 2013 | Meeting | <http://www.dagstuhl.de/13365>**13393 – Klausurtagung “AG Goesele”**

Michael Goesele (TU Darmstadt, DE)

September 25–27, 2013 | Meeting | <http://www.dagstuhl.de/13393>**13414 – Workshop**

Wolfgang Thomas (RWTH Aachen, DE)

October 9–11, 2013 | Meeting | <http://www.dagstuhl.de/13414>**13423 – Klausurtagung Graduierten-Kolleg 1194**

Uwe D. Hanebeck (KIT – Karlsruhe Institute of Technology, DE)

October 16–18, 2013 | Meeting | <http://www.dagstuhl.de/13423>**13432 – Facilitating Process and Metadata-Driven Automation in the Social, Economic, and Behavioural Sciences with the Data Documentation Initiative (DDI)**

Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (Univ. of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)

October 20–25, 2013 | Meeting | <http://www.dagstuhl.de/13432>**13442 – DDI Lifecycle: Moving Forward (Part 2)**

Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (Univ. of Minnesota – Minneapolis, US), Mary Vardigan (University of Michigan – Ann Arbor, US), Joachim Wackerow (GESIS – Mannheim, DE)

October 27 to November 1, 2013 | Meeting | <http://www.dagstuhl.de/13442>**13473 – Klausurtagung “LST Schmeck”**

Florian Allerding (KIT – Karlsruhe Institute of Technology, DE), Hartmut Schmeck (KIT – Karlsruhe Institute of Technology, DE)

November 20–22, 2013 | Meeting | <http://www.dagstuhl.de/13473>



Mitglied der

Leibniz
Leibniz-Gemeinschaft



© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2013 | ISSN 2199-1995
<http://www.dagstuhl.de>