

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report

2012



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2012

Herausgeber	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Deutschland
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. 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 des Kompetenzzentrums Informatik der Universität des Saarlandes
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
Veröffentlichungsdatum	Oktober 2013
Druckauflage	200
Druck	Digitaldruck Pirrot GmbH, Saarbrücken-Dudweiler, Deutschland
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
Online-Version	http://drops.dagstuhl.de/portals/dagstuhl_annual_reports
Redaktion	Laura Cunniff, Dr. Andreas Dolzmann, Dr. Marc Herbstritt, Christina Schwarz (Übersetzung)
Fotos	Gruppenbilder der Seminare: Heike Clemens, Jutka Gasiorowski Kunst-Bilder: Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Statistiken	Dr. Andreas Dolzmann, Dr. Christian Lindig

Vorwort

Foreword

Forschung in der Informatik vollzieht sich heute mit hohem Tempo und einer Spezialisierung, die durch die Kommunikation und Kollaboration mit Hilfe des Internets weiter zunimmt. Für internationalen Forschungsgruppen ist eine Kommunikation über E-Mail, Video-Konferenzen und Datenbanken für Arbeitsergebnisse der Normalfall. Gleichzeitig zwingt das Auftreten neuer Probleme und Fragestellungen die Informatik, sich neuen Themen in hochspezialisierten Forschungsgemeinschaften anzunehmen. Dies führt zu einer paradoxen Situation: Forscher sind so gut vernetzt wie nie zuvor, und dennoch fällt es zunehmend schwerer, den Austausch von Ideen zwischen den verschiedenen Informatik-Teildisziplinen aufrecht zu erhalten.

Vor diesem Hintergrund einer zunehmenden Beschleunigung und Virtualisierung der Spitzenforschung ist der persönliche Austausch wichtiger denn je und Schloss Dagstuhl ist stolz, dabei eine besondere Rolle zu spielen. Zahlreiche wissenschaftliche Gemeinden in der Informatik betrachten Schloss Dagstuhl als ihre wissenschaftliche Heimat und drücken dies bei ihrer Ankunft auch aus: "I love this place – it's so good to be back!"

Daneben war Schloss Dagstuhl auch 2012 der Ort für eine erstmalige Begegnung mit neuen Themen und zwischen Forschern, die noch nicht zusammen gearbeitet haben. Gerade für junge Themen an den Grenzen existierender Forschungsgebiete hat es sich Schloss Dagstuhl zur Aufgabe gemacht, ein Forum zu bieten. Das wissenschaftliche Programm 2012 war deswegen gleichermaßen von spezialisierten und interdisziplinären Themen wie Linguistik, Recht, Paläographie und Medizin geprägt. Gesellschaftlich aktuelle Fragen zu Computer-Sicherheit und Datenschutz wurden in zahlreichen Dagstuhl-Seminaren diskutiert.

Auch die Publikation wissenschaftlicher Ergebnisse in der Informatik befindet sich durch das Internet in einem Umbruch. Die Öffentlichkeit als Geldgeber fordert unter dem Schlagwort "Open Access" den kostenfreien Zugang zu Forschungsergebnissen, die traditionell in Zeitschriften veröffentlicht werden. So entstehen neue Publikationsmodelle,

Research in computer science is advancing rapidly as virtual communication tools make collaboration and information exchange via email, video chat, and shared repositories a matter of daily routine. At the same time, traditional research areas are branching out into highly specialized subcommunities as new questions and issues arise. The paradoxical result is that while computer scientists today are more connected than ever, communication across research communities is becoming harder and harder.

Against this backdrop, having a stable point of encounter for face-to-face meetings is even more important now than it was in the past and Schloss Dagstuhl is proud of its role in this regard. Ask a computer scientist "Where do you feel at home in the academic world?" and he or she may well say "Dagstuhl." "I love this place – it's so good to be back!" is a comment we frequently hear from our guests. For some core communities, Schloss Dagstuhl has indeed become an academic home. For others, especially highly interdisciplinary ones, Dagstuhl represents the only place in the world where that exact constellation of scientists can hope to meet in person.

Our 2012 program struck a balance between specialized and interdisciplinary topics covering a spectrum of informatics areas and others in linguistics, law, paleontology, and medicine, to name a few. Schloss Dagstuhl in 2012 was also very much the place where scientists met to discuss new and emerging research areas, which is a Dagstuhl priority. Security and privacy proved to be particularly hot topics, inspiring several seminars.

Bringing the focus back to research was the goal of the high-profile Dagstuhl Perspectives Workshop "Publication Culture in Computing Research," led by Professor Moshe Y. Vardi and Professor Kurt Mehlhorn. The workshop addressed the fundamental shift in publication practices that is currently taking place in our community, in part due to the increasing predominance of electronic publications. The public that funds research demands access to research results free of charge in an "Open Access" regime, opening up new

die sich im Spannungsfeld von Wissenschaftlern, Verlagen und Bibliotheken bewegen. Schloss Dagstuhl beteiligte sich 2012 intensiv an dieser Diskussion durch die Ausrichtung des Perspektiven-Workshops “Publication Culture in Computing Reseach” zusammen mit Professor Moshe Y. Vardi und Professor Kurt Mehlhorn. Der Workshop war ein Erfolg, nicht zuletzt weil Schloss Dagstuhl einen neutralen Ort für ein sehr politisches Thema bot.

Schloss Dagstuhl wird weiterhin die weltweit aufregendsten und besten Arbeiten und Wissenschaftler in der Informatik unterstützen. Wir freuen uns weiter auf Ihren Beitrag und Ihre Unterstützung.

Saarbrücken/Wadern, Oktober 2013

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

Dr. Christian Lindig
Technisch-administrativer Geschäftsführer

publication models within the highly contested space of researchers, publishers and libraries. The workshop turned out to be an instant success, in no small part because Schloss Dagstuhl provided a neutral ground that participants valued for an open and honest discussion of a highly charged topic.

Schloss Dagstuhl will continue to support exciting and excellent research in computer science for the years to come. We are looking forward to your contribution and support.

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 2012	
	<i>News from 2012</i>	4
2	Das wissenschaftliche Programm 2012	
	<i>Scientific Program 2012</i>	11
2.1	Dagstuhl-Seminare	
	<i>Dagstuhl Seminars</i>	12
2.2	Dagstuhl-Perspektiven-Workshops	
	<i>Dagstuhl Perspectives Workshops</i>	13
2.3	Einreichung der Anträge und Begutachtungsverfahren	
	<i>Proposal Submission and Review Process</i>	14
2.4	Nähere Betrachtung des Dagstuhl-Seminarprogramms 2012	
	<i>A Closer Look at the Dagstuhl Seminar Program in 2012</i>	15
2.5	Angaben zu Teilnehmern und Organisatoren	
	<i>Participant and Organizer Data</i>	16
2.6	Themen und Forschungsgebiete	
	<i>Topics and Research Areas</i>	19
2.7	Weitere Veranstaltungstypen	
	<i>Further Event Types</i>	20
2.8	Qualitätssicherung	
	<i>Quality Assurance</i>	20
2.9	Auslastung des Zentrums	
	<i>Utilization of the Center</i>	20
3	Resonanz	
	<i>Feedback</i>	23
4	Die Seminare in 2012	
	<i>The 2012 Seminars</i>	29
4.1	Foundations for Scripting Languages	31
4.2	Computability, Complexity and Randomness	32
4.3	Symmetric Cryptography	34
4.4	Learning in Multiobjective Optimization	36
4.5	Analysis of Executables: Benefits and Challenges	38
4.6	Network Attack Detection and Defense Early Warning Systems – Challenges and Perspectives	40
4.7	Software Clone Management Towards Industrial Application	42
4.8	Information Visualization, Visual Data Mining and Machine Learning	44
4.9	Principles of Provenance	46

4.10	Computation and Incentives in Social Choice	48
4.11	Normative Multi-Agent Systems	50
4.12	Applications of Combinatorial Topology to Computer Science	52
4.13	Open Models as a Foundation of Future Enterprise Systems	54
4.14	Touching the 3rd Dimension	56
4.15	Software Synthesis	58
4.16	Abstractions for scalable multi-core computing	59
4.17	Semantic Data Management	60
4.18	Quality of Experience: From User Perception to Instrumental Metrics	64
4.19	Social, Supply-Chain, Administrative, Business, Commerce, Political Networks: a Multi-Discipline Perspective	65
4.20	Artificial and Computational Intelligence in Games	66
4.21	Co-Design of Systems and Applications for Exascale	68
4.22	Cognitive Approaches for the Semantic Web	70
4.23	Future Internet for eHealth	72
4.24	Data Reduction and Problem Kernels	74
4.25	Putting Data on the Map	76
4.26	AI meets Formal Software Development	78
4.27	Architecture-Driven Semantic Analysis of Embedded Systems	80
4.28	Security and Dependability for Federated Cloud Platforms	81
4.29	Database Workload Management	82
4.30	Structure Discovery in Biology: Motifs, Networks & Phylogenies	84
4.31	Robust Query Processing	86
4.32	Mobility Data Mining and Privacy	90
4.33	Verifying Reliability	92
4.34	Engineering Multiagent Systems	94
4.35	Interaction Beyond the Desktop	96
4.36	Information Flow and Its Applications	97
4.37	Information-centric networking – Ready for the real world?	98
4.38	The Multilingual Semantic Web	100
4.39	Software Defined Networking	102
4.40	Machine Learning Methods for Computer Security	104
4.41	Biological Data Visualization	106
4.42	Privacy-Oriented Cryptography	110
4.43	Computation and Palaeography: Potentials and Limits	111
4.44	Algorithms and Complexity for Continuous Problems	112
4.45	Web Application Security	114
4.46	Coalgebraic Logics	118
4.47	Algebraic and Combinatorial Methods in Computational Complexity	120
4.48	Time-of-Flight Imaging: Algorithms, Sensors and Applications	122
4.49	Foundations and Challenges of Change and Evolution in Ontologies	126
4.50	Requirements Management – Novel Perspectives and Challenges	128
4.51	The Constraint Satisfaction Problem: Complexity and Approximability	130
4.52	Publication Culture in Computing Research	134
4.53	Games and Decisions for Rigorous Systems Engineering	136
4.54	Symbolic Methods for Chemical Reaction Networks	137
4.55	SAT Interactions	138
4.56	Is the Future of Preservation Cloudy?	140
4.57	Quantitative Security Analysis	141
4.58	Analysis of Security APIs	142
4.59	Interpreting Observed Action	143
4.60	Human Activity Recognition in Smart Environments	144
4.61	Organizational Processes for Supporting Sustainable Security	146
4.62	Securing Critical Infrastructures from Targeted Attacks	148
4.63	Divide and Conquer: the Quest for Compositional Design and Analysis	152
4.64	Representation, Analysis and Visualization of Moving Objects	153

5	Öffentlichkeitsarbeit	155
	<i>Public Relations and Outreach</i>	
5.1	Pressemitteilungen und Medienarbeit	
	<i>Press Releases and Media Work</i>	156
5.2	Fortbildung	
	<i>Educational Training</i>	157
6	Dagstuhl Publishing	159
	<i>Dagstuhl Publishing</i>	
6.1	Portfolio	
	<i>Portfolio</i>	160
6.2	Infrastruktur	
	<i>Infrastructure</i>	165
7	LZI+DBLP: Konsolidierung der bibliometrischen Datenbasis in der Informatik	167
	<i>LZI+DBLP: Consolidation of the Bibliometric Data Base in Informatics</i>	
7.1	Projektbeschreibung	
	<i>Project Description</i>	168
7.2	Ergebnisse	
	<i>Results</i>	168
7.3	Ausblick	
	<i>Outlook</i>	173
8	Einrichtung und Service	175
	<i>Facilities and Services</i>	
8.1	Tagungsräume	
	<i>Conference Facilities</i>	176
8.2	Computer und Vernetzung	
	<i>Computers and Networks</i>	176
8.3	Dagstuhl's Web-basierte Dienste	
	<i>Dagstuhl's Web-based Services</i>	177
8.4	Dagstuhl-Website	
	<i>The Dagstuhl Website</i>	177
8.5	Freizeit und Ambiente	
	<i>Leisure Facilities</i>	177
8.6	Kinderbetreuung	
	<i>Childcare</i>	178
9	Bibliothek	179
	<i>Research Library</i>	
9.1	Bestand und Angebot	
	<i>Inventory and Offering</i>	180
9.2	Spenden an die Bibliothek	
	<i>Library Donations</i>	181
10	Kunst	183
	<i>Art</i>	
10.1	Dagstuhl als Galerie	
	<i>Dagstuhl as Art Gallery</i>	184
10.2	Kunstankauf durch Spenden	
	<i>Purchasing of Art through Donations</i>	184
11	Stiftung "Informatikzentrum Schloss Dagstuhl"	187
	<i>The Dagstuhl Foundation</i>	
11.1	Zielsetzung	
	<i>Aims</i>	188
11.2	Förderung	
	<i>Contributions</i>	188
11.3	Ausblick	
	<i>Outlook</i>	189

12	Organe und Gremien	
	<i>Dagstuhl Bodies</i>	191
12.1	Struktur der Gesellschaft	
	<i>Structure of the Center</i>	192
12.2	Organe und Gremien der Gesellschaft	
	<i>Dagstuhl Bodies</i>	192
13	Statistik 2012	
	<i>Statistics 2012</i>	199
14	Veranstaltungen 2012	
	<i>Schedule of Events 2012</i>	205
14.1	Dagstuhl-Seminare	
	<i>Dagstuhl Seminars</i>	206
14.2	Dagstuhl-Perspektiven-Workshops	
	<i>Dagstuhl Perspectives Workshops</i>	210
14.3	GI-Dagstuhl-Seminare	
	<i>GI-Dagstuhl Seminars</i>	210
14.4	Lehrveranstaltungen	
	<i>Educational Events</i>	211
14.5	Sonstige Veranstaltungen	
	<i>Other Events</i>	211

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

Wissenschaftlicher Austausch in anregender Umgebung

1.1

Schloss Dagstuhl, Leibniz-Zentrum für Informatik (LZI), 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 achtziger 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 richtungsweisende 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 bietet den Gästen eine einzigartige Arbeitsumgebung, die den Austausch

Stimulating Exchanges in Relaxed Surroundings

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 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 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 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

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.

Das Leibniz-Zentrum für Informatik in Schloss Dagstuhl wird durch eine Bund-Länder-Förderung finanziert und ist seit 2006 Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 86 Forschungsinstituten, Bibliotheken und Museen. Schloss Dagstuhl beherbergt mehr als 3 000 internationale Gäste pro Jahr.

and state governments and hosts over 3 000 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 86 research institutes, libraries and museums throughout Germany.

1



Fig. 1.1
The “castle” part of Schloss Dagstuhl

Neuigkeiten in 2012

1.2

News from 2012

■ Neue Einrichtungen und Programmänderungen

Mit der Einweihung des neuen Gästehauses (siehe Fig. 1.2 and Fig. 1.2), das von dem Berliner Architekturbüro *Veauthier Meyer Architekten* entworfen wurde, startete Dagstuhl besonders gut in das Jahr 2012. Es bietet 7 Gästezimmer, eine offene Küche und einen Konferenzraum für 14 Personen, der mit Whiteboard und Beamer ausgestattet ist. Ein überdachter Außengang verbindet das Gästehaus mit dem Hauptgebäude, das einen wundervollen Blick auf die Südseite des Geländes und der umliegenden Natur bietet. Es ist somit eine gute Alternative für Einzelpersonen und kleine Gruppen, die eine gewisse Abgeschiedenheit vorziehen. Die offizielle Einweihung des Gästehauses fand am Montag, 21. Mai 2012, in einer kleinen Zeremonie statt, an der die Geschäftsführer des Zentrums, Mitglieder des Aufsichtsrats, Vertreter der Stadt Wadern und eingeladene Gäste teilnahmen. Das Zentrum verfügt in seinen drei Gebäuden nunmehr über drei Vortragsräume, sechs Konferenzräume, eine umfassende Forschungsbibliothek und 71 Gästezimmer für bis zu 89 Gäste. Weitere Details zur Ausstattung von Schloss Dagstuhl können Kapitel 8 entnommen werden.

Das Gästehaus wurde sorgfältig so geplant, dass Schloss Dagstuhl nun zwei Seminare pro Woche veranstalten kann, wozu bislang noch einige Gästezimmer fehlten. Beginnend mit der Antragsrunde Januar 2011 wurde nun begonnen, systematisch zwei Seminare pro Woche einzuplanen. 2012 bestanden 35% des Jahresprogramms des Zentrums (17 von 48 Wochen; zum Vergleich: 10 von 48 Wochen 2011) aus parallelen Seminaren. Die zusätzlichen Seminare wurden auch für neue Themen und Gemeinden von Wissenschaftlern genutzt und führten dazu, dass das Zentrum eine besonders hohe Anzahl von an Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops teilnehmenden Wissenschaftlern hatte. Eine detaillierte Analyse des wissenschaftlichen Programms findet sich in Kapitel 2.

Die wachsende Anzahl von Seminaren trug auch dazu bei, dass Schloss Dagstuhl mehr Übernachtungen (12 870) allgemein und mehr Übernachtungen (10 256) von Teilnehmern an Seminaren und Workshops hatte als jemals zuvor. Im neuen Gästehaus gab es dabei 1 014 Übernachtungen.

Ab dem 7. Oktober 2012 wurden die Preise für Kost und Logis erstmals seit 2006 in allen Kategorien dauerhaft um 10 € pro Tag angehoben, wobei jedoch gleichzeitig die Differenzierung zwischen akademischen Gästen und Gästen aus der Industrie abgeschafft wurde. Die Preisänderung wurde seitens des Aufsichtsrats von Schloss Dagstuhl zur Aufrechterhaltung der Servicequalität angesichts steigender Grundkosten als notwendig erachtet. Die neuen Preise wurden von unseren Gästen mit überwältigender Zustimmung aufgenommen. Nach dieser Umstrukturierung decken Subventionen von Bund und Ländern 75% bis 85% der Unterbringungskosten, während die Gäste je nach Veranstaltungsart für die restlichen 15% bis 25% aufkommen.

■ New Facilities and Program Changes

The year started on an excellent note with the inauguration in January of a modern guest house at Schloss Dagstuhl (see Fig. 1.2 and Fig. 1.2). Designed by the Berlin-based architectural firm *Veauthier Meyer Architekten*, the house features seven individual guest rooms, an open kitchen, and a 14-person meeting room equipped with a whiteboard wall and beamer. It is connected to the main building via a covered outdoor walkway and overlooks the quiet southern border of the conference center grounds, providing an accommodation option for those who prefer secluded quarters. The official inauguration of the building took place on Monday, May 21, 2012 in a small ceremony attended by the directors of Schloss Dagstuhl, members of the Dagstuhl Supervisory Board, representatives of the city of Wadern, and invited guests. The center now includes three lecture halls, six conference rooms, a comprehensive research library, and 71 guest rooms for up to 89 guests distributed over three buildings. See Chapter 8 for further details about the facilities at Schloss Dagstuhl.

The guest house had been carefully planned to support an expansion of the Dagstuhl Seminar program from one to two seminars per week, for which a critical number of guest rooms was needed. In keeping with this plan we began systematically booking two seminars per week after the January 2011 proposal submission round, with the result that 35% of the center's annual program featured parallel seminars (17 of 48 weeks, as compared to 10 of 48 weeks in 2011). The expanded seminar offerings embraced many new topics and communities, with a record number of scientists attending Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2012. A detailed review of our 2012 scientific program is provided in Chapter 2.

More Dagstuhl Seminars meant that Schloss Dagstuhl also hosted more overnight stays (12 870) and more stays involving Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests (10 256) in 2012 than ever before, including 1 014 overnight stays in the new guest house alone.

Schloss Dagstuhl further strengthened its financial base in 2012 by raising guest fees for room and board by 10 € per day in all categories and eliminated the distinction between industry-affiliated and academic seminar participants. The price update, the center's first since 2006, was approved by the Schloss Dagstuhl Supervisory Board as necessary to maintain the center's high quality of service in the face of rising basic costs. The new fee became effective on October 7, 2012 and met with overwhelming positive acceptance among our guests. Under the new pricing structure, German federal and state subsidies account for 75% to 85% of room and board costs while guest fees account for the remaining 15% and 25%, depending on program area.



Fig. 1.2
The new guesthouse as seen from the garden



Fig. 1.3
Balcony view

■ Das Team

Im Dagstuhl-Team fanden 2012 einige Veränderungen statt, sowohl in Bezug auf die Leitung als auch die Größe. Die Zahl der Vollzeitstellen stieg von 30,5 auf 33,6, was zum Teil der Förderung des LZI+DBLP-Projekts durch die Leibniz-Gemeinschaft und die Klaus Tschira Stiftung zu verdanken ist. Insgesamt acht neue Mitarbeiter ersetzten vier Mitarbeiter, die Schloss Dagstuhl verließen oder beurlaubt wurden, und um das wachsende Programm des Zentrums zu unterstützen, das nunmehr über drei Niederlassungen verfügt: Saarbrücken (Dagstuhl-Geschäftsstelle), Trier (dblp) und Wadern (Schloss Dagstuhl Konferenzzentrum). Die neuen Mitarbeiter besetzen eine neugeschaffene Stelle für eine studentische Hilfskraft sowie sieben Stellen in der Stammbesetzung für akademische, administrative und organisatorische Aufgaben.

Wolfgang Lorenz, der seit Mai 2000 als technisch-administrativer Direktor und Geschäftsführer tätig war – eine Position, die er seit 2009 gemeinsam mit Dr. Christian Lindig inne hatte – ging im Mai 2012 nach vielen Jahren wertvoller Zusammenarbeit in den Ruhestand. Seitdem wird Schloss Dagstuhl vom wissenschaftlichen Direktor Professor Dr. Reinhard Wilhelm zusammen mit dem technisch-administrativen Geschäftsführer Dr. Christian Lindig gemeinsam als Geschäftsführer geführt.

Außerdem verabschiedete die Dagstuhl-Geschäftsstelle mit Angelika Mueller-von Brochowski eine weitere leitende Angestellte, die seit der Gründung des Zentrums 1990 als Leiterin der Geschäftsstelle tätig gewesen war und am 1. Mai 2012 in den Ruhestand ging. Laura Cunniff, die ihren Master an der University of Michigan at Ann Arbor erhielt und zuvor im Bereich der Forschungsförderung für das spanische Wissenschaftsministerium tätig war, übernahm die Stelle im Januar 2012. In den ersten Monaten führte sie die Geschäftsstelle übergangsweise gemeinsam mit Angelika Mueller-von Brochowski.

Ebenfalls im Januar 2012 begann Dr. Michael Wagner im Rahmen eines Zweijahresvertrags mit dem Aufbau einer auf dblp basierten bibliometrischen Infrastruktur als Grundlage für szientometrische Analysen innerhalb der Informatik. Dr. Wagner hat Informatik an der Universität Trier studiert und im Anschluss als wissenschaftlicher Angestellter an der Universität Kassel gearbeitet. Seine Arbeitsstelle bei Schloss Dagstuhl wurde durch die Zuwendungen der Klaus Tschira Stiftung finanziert und erweitert die seit Mitte 2011 vom Senatsausschuss Wettbewerb (SAW) geförderte Kooperation mit der Universität Trier (Dr. Michael Ley) zur Unterstützung des dblp-Projekts (SAW-2011-LZ1-3, "LZI+DBLP"). Das Fördervolumen des SAW-Projekts, das von Mitte 2011 zwei Jahre läuft, beträgt 336 827 €. Die Klaus Tschira Stiftung stellte seit 2011 zusätzlich 60 000 € jährlich zur Verfügung. Mit Dr. Wagner stieg die Zahl von Dagstuhl-Mitarbeitern in der dblp-Gruppe auf 3,5 Vollzeitstellen.

Alle Stellen werden aus dem Kernhaushalt des Zentrums finanziert, bis auf drei Wissenschaftler, die in 2012 über Drittmittel bezahlt wurden.

■ The Dagstuhl Team

The Dagstuhl team underwent some changes in 2012 with regard to its leadership and size, which grew from 30.5 to 33.6 in full-time equivalent positions thanks in part to LZI+DBLP project funding received from the Leibniz Association and the Klaus Tschira Foundation. A total of eight new staff members were hired to replace four outgoing or on-leave members and to support the center's expanding program, which now extends over three locations in Saarbrücken (Dagstuhl Office), Trier (dblp) and Wadern (the Schloss Dagstuhl conference center itself). The incoming staff filled one newly created student position and seven core positions charged with academic, administrative or housekeeping responsibilities.

Throughout 2012 the executive management of Schloss Dagstuhl remained in the hands of the center's Scientific Director, Professor Dr. Reinhard Wilhelm, and the Technical and Administrative Director Dr. Christian Lindig. Wolfgang Lorenz, who had served as Technical and Administrative Director of Schloss Dagstuhl since 1990 – a position he held jointly with Dr. Lindig since 2009 – retired in May of 2012 following many years of valuable service and collaboration.

The Dagstuhl Office said goodbye to another key staff member when Angelika Mueller-von Brochowski, head of the Dagstuhl Office since Schloss Dagstuhl's earliest days, retired on May 1, 2012. Laura Cunniff, who earned her Master's degree at the University of Michigan at Ann Arbor and served as an executive program administrator under the Spanish Ministry of Science and Innovation prior to coming to Dagstuhl, took over the position in January of 2012. She and Ms. Mueller-von Brochowski enjoyed a successful transitional period of joint collaboration during the first four months of 2012.

Also in January of 2012, Dr. Michael Wagner joined the Dagstuhl team to set up a dblp-based bibliometric infrastructure as a basis for scientometric analyses in computer science. Dr. Wagner has studied computer science at the University of Trier and worked as research assistant at the University of Kassel. His position, which is funded by a grant from the Klaus Schira Foundation, extends the cooperation between Schloss Dagstuhl and the University of Trier (Dr. Michael Ley) to support the dblp project. This cooperation started in mid-2011 and is financed by a special grant from the Senatsausschuss Wettbewerb (SAW-2011-LZ1-3, "LZI+DBLP"), which provides a total of 336 827 € for a 2-year period beginning in mid-2011. Additionally, the Klaus Schira Foundation has provided 60 000 € annually starting in 2011. With Dr. Wagner's addition to the team, the number of Dagstuhl staff working out of the dblp group rose to 3.5 full-time equivalent positions.

All staff at Schloss Dagstuhl are funded under the center's core budget, with the exception of three academic positions that were supported by third-party funding in 2012.

■ Dagstuhl Publishing

Schloss Dagstuhl hat auch in 2012 erfolgreich seine Arbeit zur Förderung von Open Access, d.h. den kostenfreien Zugang zu wissenschaftlicher (Online-)Literatur, fortgesetzt, insbesondere durch die Verlagsangebote, die sich an die weltweite Informatik-Community richten. Die Diskussion über Open Access bekam in den Medien ziemlichen Aufwind durch die von Timothy Gowers ins Leben gerufene Boykott-Initiative “The cost of knowledge”, die sich insbesondere gegen den Wissenschaftsverlag Elsevier und dessen exorbitante Zeitschriftenpreise gerichtet hat. Aber auch innerhalb der Informatik wird Open Access bei den Entscheidungsträgern der Fachgesellschaften zunehmend zu einem wichtigen Thema. Bestehende Verträge mit herkömmlichen Verlagen verzögern allerdings die Umstellung von Publikationsserien auf Open Access. Die von Schloss Dagstuhl angebotenen “goldenen Open Access”-Angebote, d.h., die Artikel sind direkt mit ihrer Veröffentlichung online frei zugänglich, sind im Informatik-Umfeld einzigartig was Qualität, Kostenmodell und Service angeht. Es ist selbstverständlich, dass Schloss Dagstuhl sich auch in den kommenden Jahren seinem Engagement beim Thema Open Access treu bleibt.

Ein anderes interessantes Projekt in 2012 war die Retrodigitalisierung von Seminarberichten aus den Anfängen von Schloss Dagstuhl (1990–1993), die bisher nur in Papierform vorlagen. Diese Dokumente haben auch einen historischen Wert und können nun einfach in elektronischer Form über die jeweilige Seminar-Webseite abgerufen werden.

Weitere Details zu Schloss Dagstuhl’s Publikationsaktivitäten finden sich in Kapitel 6.

■ Wissenschaftliche und institutionelle Öffentlichkeitsarbeit

Nachdem bereits 2011 die ACM (Association for Computing Machinery) den begeisterten Leitartikel “Where Have All the Workshops Gone?” (Communications of the ACM, Vol. 54, No. 1) des amerikanischen Informatikers Moshe Y. Vardi veröffentlicht hatte, publizierte sie bereits 2012 mit “A Workshop Revival” (Communications of the ACM, Vol. 55, No. 5) einen zweiten Artikel über Schloss Dagstuhl. Grundlage dieses Artikels war ein Interview über den internationalen Erfolg des Konzepts von Schloss Dagstuhl, das der amerikanische Wissenschaftsjournalist Paul Hyman mit Dagstuhls Wissenschaftlichem Direktor Professor Dr. Reinhard Wilhelm Anfang 2012 geführt hatte.

Außerdem wurde Professor Wilhelm im Februar 2012 von Cris Calude interviewt (EATCS Bulletin Nr. 106, News from New Zealand, “A Dialogue with Reinhard Wilhelm about Compiler Construction and Dagstuhl”). Diese Gelegenheit ermöglichte es Professor Wilhelm offen über die Gründung von Schloss Dagstuhl und die Gemeinsamkeiten und Unterschiede zum Mathematisches Forschungsinstitut Oberwolfach zu sprechen sowie über den Einsatz des Zentrums für die bildenden Künste und seine eigene Rolle als Wissenschaftlicher Direktor von Schloss Dagstuhl.

Als eingeladener Redner besuchte Professor Wilhelm anlässlich des “NII Shonan Meeting Memorial Symposium – For Making Future Value From Asia” am 1. November 2012 das Shonan Meeting Center in Japan. Er berichtete, dass das

■ Dagstuhl Publishing

In 2012, Schloss Dagstuhl successfully continued its work on promoting open access to scientific literature by offering publishing services to the informatics community on a global scale. The open access debate gained media prominence early in the year thanks to Timothy Gower’s “The cost of knowledge” initiative, which aimed at a collective boycott of Elsevier due to their exorbitant journal prices. Also within the informatics community, the policymakers of learned societies are becoming increasingly aware of open access. However, existing contracts with legacy publishers lead to a delay in transforming publishing venues to open access. The so-called “gold open access” services offered by Schloss Dagstuhl (which means that articles are accessible online and free of charge immediately upon publishing) are unique in the computing research communities with regard to quality, pricing, and service. Schloss Dagstuhl will stick to its commitment to open access during the coming years.

Another effort in 2012 was the digitization of documentary reports of the very first Dagstuhl Seminars from 1990–1993, which were previously available only in print versions. These documents have also a historical value and can now be retrieved easily in electronic form from the respective seminar webpages.

For more details on Schloss Dagstuhl’s publishing activities, please see Chapter 6.

■ Scientific and Institutional Outreach

Following the glowing review of Schloss Dagstuhl given by prominent computer scientist Moshe Y. Vardi in his 2011 editorial, “Where Have All the Workshops Gone?” (Communications of the ACM, Vol. 54, No. 1), the ACM published a second article about our center in 2012 entitled “A Workshop Revival” (Communications of the ACM, Vol. 55, No. 5). The article came out of science journalist Paul Hyman’s interview of Dagstuhl Scientific Director Professor Dr. Reinhard Wilhelm early in 2012 about the success of the Dagstuhl concept internationally.

Professor Wilhelm was also interviewed by Cris Calude in February of 2012 (EATCS Bulletin No. 106, *News from New Zealand*, “A Dialogue with Reinhard Wilhelm about Compiler Construction and Dagstuhl”). The opportunity gave him a chance to speak candidly about the founding of Schloss Dagstuhl, its similarities and differences to/from the Mathematisches Forschungsinstitut Oberwolfach, Dagstuhl’s dedication to the fine arts, and his own role as Scientific Director of Schloss Dagstuhl.

On November 1, Professor Wilhelm traveled to Japan to visit the Shonan Meeting Center as an invited speaker at the “NII Shonan Meeting Memorial Symposium – For Making Future Value From Asia.” He observed that Shonan operates under less favorable conditions with respect to facilities and financing than those enjoyed by Schloss Dagstuhl.

Closer to home, in October of 2012 Schloss Dagstuhl was

japanische Gegenstück zu Schloss Dagstuhl bezüglich finanzieller und operativer Aspekte unter weniger günstigen Bedingungen wie Dagstuhl selbst geführt wird.

Im Oktober 2012 besuchte der Präsident der Leibniz-Gemeinschaft, Professor Dr. Karl Ulrich Mayer, Schloss Dagstuhl. Zusätzlich zu einem Rundgang mit den Mitarbeiterinnen und Mitarbeitern und einem Treffen mit Seminarorganisatoren sprach Professor Mayer mit den beiden Geschäftsführern über die Geschichte, Funktionsweise, Probleme, Gremien und das wissenschaftliche Programm des Zentrums.

Dr. Marc Herbstritt, Mitglied des wissenschaftlichen Stabs und verantwortlich für die Open Access-Verlagsangebote von Schloss Dagstuhl, hielt am 16. Juni 2012 einen eingeladenen Vortrag mit dem Titel "Pawlow, Turing und Open Access – Gedanken zur Konditionierung in der Informatik" in der neu aufgelegten Kolloquiumsreihe des ZPID (Leibniz-Zentrum für Psychologische Information und Dokumentation, Trier). Der Vortrag diskutierte klassische Begriffe der Konditionierung im Kontext heutiger (missbräuchlicher) Praktiken beim wissenschaftlichen Publizieren.

In Zusammenarbeit mit der Universität Trier/dblp und der Universität des Saarlandes hat Schloss Dagstuhl während des Sommersemesters 2012 eine Vortragsreihe zum Thema Bibliometrie veranstaltet, da dieses Thema im Rahmen von dblp zunehmend an Bedeutung gewinnt. Mit Stefanie Hausteine (FZ Jülich), Peter van den Besselaar (VU Amsterdam) und Debora Weber-Wulff (HTW Berlin) konnten kompetente Redner gewonnen werden.

■ Öffentlichkeitsarbeit und Lehrerfortbildung

Um junge Journalisten und Volontäre zu ermutigen, auch über komplexe Themen der Informatik zu berichten, bietet Schloss Dagstuhl alljährlich einen Workshop über Wissenschaftsjournalismus an. Der Workshop fand 2012 vom 3. bis 6. Juni parallel zum Dagstuhl-Seminar 12231, "Future Internet for eHealth" statt. Die Trainer waren Tim Schröder aus Oldenburg (wissenschaftlicher Autor und Mediencoach) und Gordon Bolduan von der Universität des Saarlandes (M2CI).

Die Europäische Union der Gesellschaften für Wissenschaftsjournalisten (European Union of Science Journalists' Associations, EUSJA) besuchte Schloss Dagstuhl vom 29. bis 30. März 2012 im Rahmen ihrer Saarlandtour. Bereits 2008 war das Zentrum während der EUSJA-Tour Gastgeber, was einige Berichte über Informatik im Saarland mit sich brachte. Die EUSJA-Journalisten übernachteten im Schloss und erfuhren Genaueres über Schloss Dagstuhl und die Wissenschaftler, die an dem Dagstuhl-Seminar 12331 "Open Models as a Foundation of Future Enterprise Systems" teilnahmen, das zu diesem Zeitpunkt gerade stattfand.

Der Lehrerfortbildungsworkshop, der seit 1991 einmal jährlich stattfindet, wurde vom 12. bis 14. Dezember 2012 zum 22. Mal auf Schloss Dagstuhl abgehalten. Durch Unterstützung der beteiligten pädagogischen Institute (LPM Saarland, PL Rheinland-Pfalz) konnten dieses Jahr erstmalig fünf Lehrer von außerhalb des Saarlandes bzw. Rheinland-Pfalz eingeladen werden.

Weitere Details finden sich in Kapitel 5.

visited by the president of the Leibniz Association, Professor Dr. Karl Ulrich Mayer. In addition to touring the grounds with our staff and meeting two Dagstuhl Seminar organizers, Professor Mayer discussed with our directors the history, functioning, current problems, bodies and boards, and scientific program of the center.

In June of 2012 Dagstuhl scientific staff member Dr. Marc Herbstritt, who currently leads the center's effort to promote open-access publishing at Schloss Dagstuhl and in the computer science community as a whole, gave the invited talk "Pawlow, Turing und Open Access – Gedanken zur Konditionierung in der Informatik" within the context of the freshly re-activated colloquium series of the ZPID center (Leibniz-Zentrum für Psychologische Information und Dokumentation, Trier). The talk discussed the relationship between Pavlovian conditioning and today's (mal)practices in scholarly publishing.

In cooperation with the University of Trier/dblp and Saarland University, Schloss Dagstuhl also organized a colloquium series on bibliometrics during the 2012 summer semester at Saarland University as part of its growing involvement in the dblp database initiative. Stefanie Hausteine (FZ Jülich), Peter van den Besselaar (VU Amsterdam) and Debora Weber-Wulff (HTW Berlin) contributed featured talks.

■ Public Relations and Educational Training

In order to encourage young journalists and trainees to report on computer science topics, Schloss Dagstuhl offers an annual workshop on science journalism. In 2012, the workshop took place from June 3 to 6 in parallel with Dagstuhl Seminar 12231 "Future Internet for eHealth." Trainers included scientific writer and media trainer Tim Schröder from Oldenburg and Gordon Bolduan from Saarland University (M2CI).

The European Union of Science Journalists' Associations (EUSJA) visited Schloss Dagstuhl during their tour to Saarland on June 29–30, 2012. The center had already hosted an EUSJA tour in 2008, which resulted in reports on computer science in Saarland. The EUSJA journalists stayed overnight at Dagstuhl and learned more about Dagstuhl and the scientists participating in Dagstuhl Seminar 12331 "Open Models as a Foundation of Future Enterprise Systems" taking place on that week.

The center's 22nd annual Teacher Training Workshop was held on December 12–14, 2012 in Schloss Dagstuhl. With support from both educational institutes that co-organize the workshop (LPM Saarland, PL Rheinland-Pfalz), it was possible for the first time to invite five teachers from outside of Saarland and Rhineland Palatinate respectively.

See Chapter 5 for an overview of public relations and educational outreach at Schloss Dagstuhl.

■ Spender und Förderer von Schloss Dagstuhl

Die Dagstuhl-Stiftung und die Bibliothek wurden dankenswerterweise auch 2012 durch Spenden von Gästen, Instituten und Firmen großzügig unterstützt. Nähere Informationen finden sich in den Kapiteln 9 und 11.

Die Klaus Tschira Stiftung fördert für 2011 und 2012 das LZI+DBLP-Projekt mit einer Spende von 120 000 € (60 000 € jeweils für 2011 und 2012). Dieses Projekt wurde 2012 durch eine Förderung des Senatsausschuss Wettbewerb (SAW) der Leibniz-Gemeinschaft in Höhe von 167 094 € finanziert (SAW-2011-LZI-3). Der Zuschuss belief sich auf insgesamt 336 824 € und erstreckte sich über einen Zeitraum von zwei Jahren, beginnend Mitte 2011.

Die Kunstsammlung von Schloss Dagstuhl wurde seit dem Bestehen des Zentrums von Kunstmäzenen kontinuierlich durch Sach- und Geldspenden gefördert. Jill Knuth, die Ehefrau von Professor Donald Knuth, spendete dem Zentrum zwei von ihr handgefertigte Quilts. Anlässlich der Teilnahme ihres Mannes am Dagstuhl-Seminar 12471 “SAT Interactions” präsentierte sie im November 2012 persönlich die beiden Geschenke (siehe Fig. 4.18). Das ungewöhnliche Muster einer der Decken mit dem Titel “The Hexagonal Prime Quilt”, wurde unter den Seminargästen eifrig diskutiert. Bei ihrem Aufenthalt klärte Jill Knuth, dass die auf der Decke abgebildeten Zahlen tatsächlich zufällig gewählt worden seien.

Schloss Dagstuhl erhielt ebenfalls 2012 das frisch restaurierte Selbstportrait der Gräfin Octavie de Lasalle von Louisenthal (1811–1890), deren künstlerisches und leidenschaftliches Leben sie zu einer der faszinierendsten Persönlichkeiten der historischen Vergangenheit von Schloss Dagstuhl macht. Das Gemälde, das zuletzt einem privaten Sammler in Münster gehörte, wurde dem Zentrum 2011 großzügigerweise von Wolfgang Lorenz anlässlich der Ausstellung “Octavie de Lasalle von Louisenthal und Schloss Dagstuhl” gespendet. Nach der Ausstellung wurde das Bild von Axel Wieland restauriert und mit einem eigens von Angelika Mueller-von Brochowski gespendeten Rahmen versehen. Am 13. Dezember 2012 wurde das Kunstwerk in einer kleinen Feier offiziell an Schloss Dagstuhl übergeben.

Am gleichen Tag erhielt Dagstuhl eine fotografische Reproduktion eines bislang unbekanntem Selbstportraits der Gräfin. Obwohl es keine schriftlichen Überlieferungen zum Verbleib des Gemäldes gibt, ist bekannt, dass es noch vor 1945 auf Schloss Dagstuhl ausgestellt war. Die von Manfred Stein gespendete verkleinerte Reproduktion basiert auf einem wiederentdeckten Foto des Kunstwerks.

Schloss Dagstuhl erhielt 2012 weiterhin Spenden von insgesamt über 3 077 € für den Ankauf von Werken der Künstler Birgit Ginkel, Mane Hellenthal, Josef Linschinger, Uwe Loebens, Marlene Reucher, Maxwell J. Roberts und Ila Wingen, die in Dagstuhl ausgestellt hatten. Weitere Details finden sich in Kapitel 10.

■ Änderungen im Gästeservice

Die Kinderbetreuung von Schloss Dagstuhl wurde 2012 für 14 Kinder in Anspruch genommen, mehr als jemals zu-

■ Sponsors and Donors of the Center

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generously donating funds to support the Dagstuhl Foundation and books for the Dagstuhl library. More details can be found in chapters 9 and 11 of this report.

With respect to project grants, Dagstuhl obtained a two-year donation of 120 000 € from Klaus Tschira Stiftung (60 000 € for 2011 and 2012, respectively) in support of the LZI+DBLP project. The project was further financed in 2012 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 is also 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. Jill Knuth, wife of Professor Donald Knuth, donated to the center two beautiful quilts made by her, which she presented in person together with her husband on the occasion of Professor Knuth’s participation in Dagstuhl Seminar 12471, “SAT Interactions” (November 18–23, 2012), see Fig. 4.18. The unusual pattern on one of the quilts, “The Hexagonal Prime Quilt,” had inspired quite a bit of commentary among our scientific guests. During the visit, Jill Knuth was able to clarify that the numbers on the quilt are, in fact, random.

Schloss Dagstuhl was also pleased to receive in 2012 the freshly restored self-portrait of Countess Octavie de Lasalle von Louisenthal (1811–1890), whose artistic and impassioned life make her one of Dagstuhl’s most fascinating historical personages. The painting, originally in the hands of a private owner from Münster, was acquired by the center thanks to a generous donation from Wolfgang Lorenz following its inclusion in the 2011 exhibit “Octavie de Lasalle von Louisenthal und Schloss Dagstuhl.” Shortly thereafter it was carefully restored by Axel Wieland and given a custom framing as a donation from Angelika Mueller-von Brochowski. The finished piece was formally handed over in a small ceremony at Schloss Dagstuhl on December 13, 2012.

On the same occasion, the center also received a photographic reproduction of a hitherto unknown second self-portrait of the countess. There is no surviving documentation as to the painting, although it is known that prior to 1945 the portrait was in Schloss Dagstuhl. The photograph was a gift of Manfred Stein.

Schloss Dagstuhl also received art donations exceeding 3 077 € in 2012 for a variety of works by artists Birgit Ginkel, Mane Hellenthal, Josef Linschinger, Uwe Loebens, Marlene Reucher, Maxwell J. Roberts, and Ila Wingen. For further details about our art program, see Chapter 10.

■ Changes in Dagstuhl Services

In 2012 Schloss Dagstuhl received child care requests for a total of 14 children, more than in any previous year. To

vor. Um der steigenden Nachfrage nach Kinderbetreuung gerecht zu werden, hat Schloss Dagstuhl dieses Angebot verbessert und weitere Betreuerinnen gewonnen, die bei Nachfrage die Betreuung der Kinder übernehmen. Seit Oktober 2012 betreut bei Nachfrage vor allem Elke Hubertus, eine zertifizierte Erzieherin aus der Umgebung von Wadern, die kleinsten Gäste des Zentrums. Eine konkrete Neuerung war die Einführung eines Fragebogens, der vor dem Aufenthalt auf Schloss Dagstuhl ausgefüllt wird, sodass sichergestellt ist, dass die Kinder unserer Gäste einen schönen Aufenthalt in unserem Zentrum haben, während gleichzeitig die Sicherheitsanforderungen erfüllt werden.

meet the growing demand for child care at our center, the center expanded its pool of on-call nannies and improved its service on several counts in 2012. In October 2012, Elke Hubertus, a professionally trained and certified nanny based in the Wadern area, began supporting our center's smallest guests on an as-needed basis. A specific novelty in our childcare program for 2012 was the introduction of a pre-visit questionnaire to ensure that the children of our scientific guests have a comfortable and fun stay at our center that also meets safety requirements.



Fig. 1.4

Dagstuhl's longstanding Technical Administrative Director Wolfgang Lorenz (left), with Scientific Director Reinhard Wilhelm

2 **Das wissenschaftliche Programm 2012** *Scientific Program 2012*

Dagstuhl-Seminare

2.1

Dagstuhl Seminars

Die Dagstuhl-Seminare haben als wesentliches Instrument der Forschungsförderung Priorität bei der Gestaltung des Jahresprogramms. Das 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 21 Preisträger des Turing-Awards, die höchste 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-Website heruntergeladen werden. Weitere Informationen finden sich unter <http://www.dagstuhl.de/dagrep/>.

In den Kapiteln 14 und 4 finden sich Listen aller Veranstaltungen, die 2012 auf Schloss Dagstuhl stattfanden, sowie Zusammenfassungen der Seminare und Perspektiven-Workshops. Auf der Dagstuhl-Website ist ein tagesaktuelles Programm für die kommenden 18 Monate verfügbar.

Dagstuhl Seminars, our key instrument for promoting research, are accorded top priority in the center's 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 21 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 12 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. See <http://www.dagstuhl.de/dagrep/> for further information.

Chapters 14 and 4 provide a comprehensive list of all events that took place at Schloss Dagstuhl during the year under review and summaries of the 2012 Seminars and Perspectives Workshops. An up-to-the-minute program covering the coming 18 months is available on the Dagstuhl website.

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, welche den aktuellen Stand des Gebietes, offene Probleme, Defizite und vielversprechende Richtungen beschreiben, vorgetragen. 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 2012 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. Weitere Informationen zu den Dagstuhl-Perspektiven-Workshops und ihrer Ergebnisse finden sich auf

- <http://www.dagstuhl.de/pw-list>
- <http://drops.dagstuhl.de/dagman>

Dagstuhl Perspectives Workshops

2

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

The Dagstuhl Perspectives Workshops that took place in 2012 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*. For more information on Dagstuhl Perspectives Workshops and their results, see:

- <http://www.dagstuhl.de/pw-list>
- <http://drops.dagstuhl.de/dagman>

Social, Supply-Chain, Administrative, Business, Commerce, Political networks: a multi-discipline perspective

<http://www.dagstuhl.de/12182>

Co-Design of Systems and Applications for Exascale

<http://www.dagstuhl.de/12212>

Machine Learning Methods for Computer Security

<http://www.dagstuhl.de/12371>

Computation and Palaeography: Potentials and Limits

<http://www.dagstuhl.de/12382>

Publication Culture in Computing Research

<http://www.dagstuhl.de/12452>

Fig. 2.1
Dagstuhl Perspectives Workshops hosted in 2012

Einreichung der Anträge und Begutachtungsverfahren

2.3

Proposal Submission and Review Process

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 vorläufige Gästeliste 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 Organisations- und Unterstützerteam 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 von wissenschaftlichen Gemeinden, geographischer Regionen und besonders auf der Miteinbeziehung von jungen und weiblichen Wissenschaftlern. 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.

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.

Nähere Betrachtung des Dagstuhl-Seminarprogramms 2012

2.4

A Closer Look at the Dagstuhl Seminar Program in 2012

In 2012 wurden 90 Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops eingereicht. Über 76% der eingereichten Anträge wurden in den Direktoriumssitzungen im Januar und Juni 2012 genehmigt, was die ausgezeichnete Qualität der Anträge widerspiegelt (siehe Fig. 2.2).

Durch die Eröffnung des neuen Gästehauses im Januar 2012 wurde es ab Sommer 2012 möglich, in jeder Woche zwei Seminare (ein großes und ein kleines) zu veranstalten. Bei der Ankündigung der beiden Antragsrunden 2012 wurde explizit darauf hingewiesen, dass kleine Seminare schneller als große terminiert würden. Dadurch stieg 2012 die Anzahl der Anträge für kleine Seminare im Vergleich zu den Vorjahren erheblich an. Von den 69 vom Direktorium akzeptierten Seminaren und Workshops waren 32 klein (vgl. Fig. 2.3). Die meisten dieser kleinen Seminare fanden bereits Ende 2012 bzw. Anfang 2013 statt, sodass die durchschnittliche Zeit zwischen Genehmigung und Seminar verkürzt wurde.

In 2012 wurden daher bereits 35% der verfügbaren Kapazitäten von Dagstuhl mit parallelen Seminaren belegt, während es 2011 nur 12% waren. Voraussichtlich wird sich der Anteil 2013 nahezu verdoppeln, da mehr und mehr der nicht durch regulär große Seminare ausgefüllten Kapazitäten nun mit kleinen Seminaren belegt werden. Dies führt ebenfalls zu einer Erhöhung der Gesamtanzahl von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops: Über die Hälfte (64 von 122) aller Veranstaltungen waren 2012 Seminare und Workshops an denen mehr als 70% aller Gäste (2 346 von 3 482) teilgenommen haben (vergleiche Fig. 2.4).

We received 90 proposals for new Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2012. Over 76% of proposals were accepted by Dagstuhl’s Scientific Directorate during the January 2012 and June 2012 submission rounds, reflecting the excellent quality of the proposals received (see Fig. 2.2).

Thanks to the opening of the new guest house in January of 2012, Dagstuhl was able to begin systematically booking two parallel seminars per week (one small one and one large seminar) in the late summer of 2012. The center had prepared for this by actively promoting the fast scheduling of small seminars prior to the January and June 2012 proposal submission rounds. As a result, the number of seminar proposals that applied for a small seminar rose in 2012 in comparison with previous years. Of 69 new Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted by the Directorate in 2012, 32 were small seminars (see Fig. 2.3). Many of these small seminars were scheduled in late 2012 and early 2013, thereby reducing the overall lead time of the accepted seminars and workshops.

In total, 35% of Dagstuhl’s program space featured parallel seminars in 2012, as opposed to only 12% in 2011. That proportion is expected to roughly double in 2013 as more and more of the available program slots are filled with small seminars. More parallel seminars also meant more seminars and workshops overall: over half of the center’s 2012 scientific program (64 out of 122 events) was devoted to Dagstuhl Seminars and Dagstuhl Perspectives Workshop and nearly 70% of those who visited our center in 2012 (2 346 guests out of 3 482) did so in order to participate in one of our seminars or workshops (see Fig. 2.4).

Year	Proposals		Accepted		Rejected	
	#	%	#	%	#	%
2008	83	72.3	60	23	27.7	
2009	95	71.6	68	27	28.4	
2010	94	69.2	65	29	30.9	
2011	80	67.5	54	26	32.5	
2012	90	76.7	69	21	23.3	

Fig. 2.2 Dagstuhl Seminar proposals and acceptance rates

	small	large
short	12	–
long	20	37

Fig. 2.3 Small vs. large and short vs. long Dagstuhl Seminars and Dagstuhl Perspectives Workshops approved in 2012, among them 32 small and 37 large seminars/workshops. 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 Dagstuhl-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. Die Mehrheit (1 293) der Teilnehmer von Dagstuhl-Seminaren 2012 war zum ersten Mal Gast in unserem Zentrum, während 1 053 Teilnehmer bereits mindestens eine Veranstaltung auf Schloss Dagstuhl besucht hatten (siehe Fig. 2.5).

Ein beträchtlicher Anteil der Gäste besteht aus jungen Wissenschaftlern am Anfang ihrer Karriere, die unter Umständen ein Leben lang von dem Dagstuhl-Erlebnis zehren. Etwa 28% der Gäste der Seminare und Workshops in 2012, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein, 52% als erfahrene Forscher (siehe Fig. 2.6). Dieses 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 72% war der Anteil an Gästen aus dem Ausland 2012 wieder sehr hoch. Das Diagramm in Abbildung 2.7 zeigt die regionale Verteilung der Gäste bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops 2012. Eine detaillierte Aufstellung der Herkunftsländer aller Teilnehmer bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops und anderer Veranstaltungen kann Kapitel 13 entnommen werden.

Das Ziel des Seminarprogramms ist es, ein ausgewogenes wissenschaftliches Programm anzubieten. Antragsteller werden angehalten, qualifizierte weibliche Kollegen in das Organisatorenteam und die Gästeliste aufzunehmen. 2012 waren 50% aller Organisatorenteams des wissenschaftlichen Programms hinsichtlich des Geschlechts gemischt, wobei dieses Verhältnis im Vergleich zu den meisten Vorjahren seit 2008 relativ konstant geblieben ist (siehe Fig. 2.8). Im Gegensatz dazu ist der prozentuale Anteil an weiblichen Seminarteilnehmern mit mehr als 16% bei einem Allzeithoch angelangt (siehe Fig. 2.9).

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. The majority (1 293) of Dagstuhl Seminar participants in 2012 were first-time visitors to Dagstuhl, followed by 1 053 participants who had already attended at least one previous event 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 28% of our 2012 seminar and workshop survey respondents self-classified as junior and 52% as senior (see Figure 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.

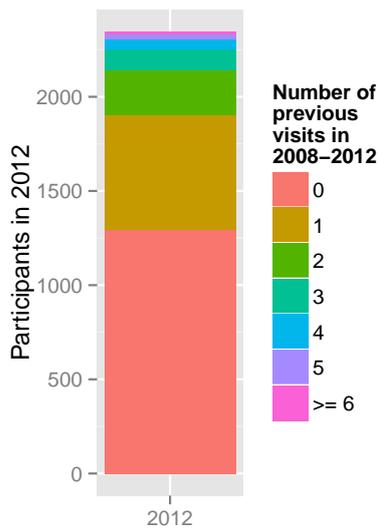
At 72%, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2012. The pie chart in Figure 2.7 shows the regional distribution our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2012. For a detailed breakdown of the countries of origin for all participants in Dagstuhl Seminars, Dagstuhl Perspectives Workshops, and other events at our center, please refer to 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 2012, 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). By contrast, the percentage of female seminar participants rose to an all-time high both in total and relative terms, peaking at over 16% (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

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.

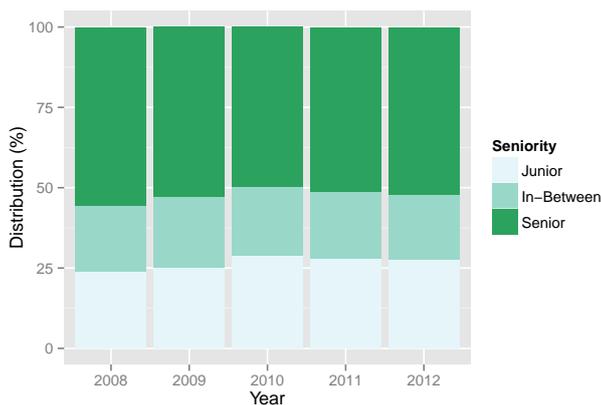


Previous visits 2008–2012	Participants	
	#	%
0	1293	55.1
1	611	26.0
2	237	10.1
3	111	4.7
4	57	2.4
5	22	0.9
≥ 6	15	0.8

(a) Graphical distribution

(b) Distribution of previous visits

Fig. 2.5
Dagstuhl participants in 2012 and the number of Dagstuhl Seminars or Dagstuhl Perspectives Workshops they attended from 2008 to 2012. In 2012, 55.1% of the participants took part in a seminar for the first time.



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

(a) Graphical distribution

(b) Detailed numbers

Fig. 2.6
Self-assigned seniority of Dagstuhl Seminar participants

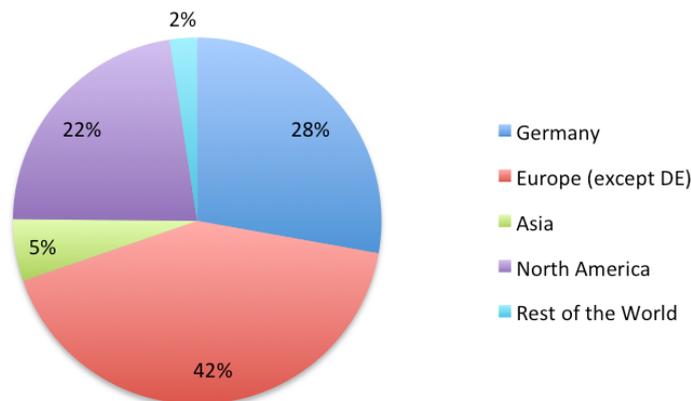
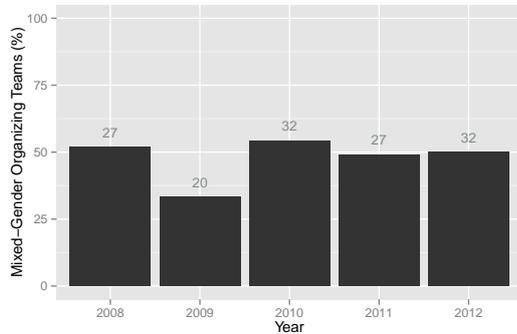


Fig. 2.7
Origin of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2012



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

(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	295	13.6
2010	2053	293	14.3
2011	1958	294	15.0
2012	2346	377	16.1

Fig. 2.9

Female participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops by year

Themen und Forschungsgebiete

2.6

Topics and Research Areas

2

Sicherheit und Datenschutz sind weiterhin wichtige Themen in unserer Gesellschaft. Beide Gebiete waren auch stark im Dagstuhl-Programm 2012 vertreten. Es umfasste nicht nur Dauerbrenner wie *Symmetric Cryptography* (Dagstuhl-Seminar 12031) und *Privacy-Oriented Cryptography* (Dagstuhl-Seminar 12381), sondern auch neu aufgekommene Themen wie *Machine Learning Methods for Computer Security* (Dagstuhl-Perspektiven-Workshop 12371), *Quantitative Security Analysis* (Dagstuhl-Seminar 12481), *Analysis of Security APIs* (Dagstuhl-Seminar 12482), *Web Application Security* (Dagstuhl-Seminar 12401) und *Organizational Processes for Supporting Sustainable Security* (Dagstuhl-Seminar 12501).

Viele Seminare zu diesem Themengebiet waren eng mit Anwendungsfragen verknüpft. Besonders die Sicherheit der netzgebundenen Infrastruktur war ein beliebtes Teilthema, was in Seminaren zu *Securing Critical Infrastructures from Targeted Attacks* (Dagstuhl-Seminar 12502), *Network Attack Detection and Defense Early Warning Systems* (Dagstuhl-Seminar 12061) sowie *Security and Dependability for Federated Cloud Platforms* (Dagstuhl-Seminar 12281) resultierte. Drei Seminare waren allein Cloud Computing gewidmet.

Große Teile der netzgebundenen Infrastruktur und viele Webdienste basieren auf schlecht konstruierten Skriptsprachen. Das erfolgreiche Seminar *Foundations for Scripting Languages* (Dagstuhl-Seminar 12011) sprach die Notwendigkeit weiterer Forschungsarbeit zur Verbesserung dieser Situation an.

Außerdem gab es auch 2012 einige stark interdisziplinär geprägte Seminare wie z.B. *Computation and Incentives in Social Choice* (Dagstuhl-Seminar 12101), das Informatiker, Mathematiker, Sozialwahltheoretiker und Politologen zusammenbrachte. Darüber hinaus konnte ein Trend zu Informatikforschung in Bereichen außerhalb der Informatik wie z.B. Linguistik, Paläographie und Bioinformatik beobachtet werden, wie zum Beispiel im Workshop *Computation and Paleography: Potentials and Limits* (Dagstuhl-Perspektiven-Workshop 12382). Mindestens eine Wissenschaftsgemeinschaft kam 2012 zum ersten Mal nach Dagstuhl, nämlich zum Thema *Artificial and Computational Intelligence in Games* (Dagstuhl-Seminar 12191).

Das Jahr endete mit dem nachdenklich stimmenden Workshop *Publication Culture in Computing Research* (Dagstuhl Perspektiven-Workshop 12452). In diesem Workshop wurde die Notwendigkeit der Neuausrichtung von Zielen und Praktiken rund um die Veröffentlichungskultur in der Informatik adressiert.

Generell umfasste das Seminar-Programm 2012 wie üblich eine breite Palette von Forschungsgebieten, wobei eine Vielzahl von Wissenschaftsgemeinschaften und Disziplinen involviert war.

Security and Privacy remain hot topics for our society, and Schloss Dagstuhl's 2012 program was strong in both areas. The program included some evergreen topics such as *Symmetric Cryptography* (Dagstuhl Seminar 12031) and *Privacy-Oriented Cryptography* (Dagstuhl Seminar 12381), alongside new and emergent ones such as *Machine Learning Methods for Computer Security* (Dagstuhl Perspectives Workshop 12371), *Quantitative Security Analysis* (Dagstuhl Seminar 12481), *Analysis of Security APIs* (Dagstuhl Seminar 12482), *Web Application Security* (Dagstuhl Seminar 12401), and *Organizational Processes for Supporting Sustainable Security* (Dagstuhl Seminar 12501).

Many seminars in this topic area had strong ties to application. The security of our networked infrastructure proved to be a particularly hot sub-topic, giving rise to numerous seminars, including *Securing Critical Infrastructures from Targeted Attacks* (Dagstuhl Seminar 12502), *Network Attack Detection and Defense Early Warning Systems* (Dagstuhl Seminar 12061), and *Security and Dependability for Federated Cloud Platforms* (Dagstuhl Seminar 12281). Cloud computing itself inspired no less than three seminars.

Badly designed scripting languages are the basis for much of our web infrastructure and many web services. The successful seminar on *Foundations for Scripting Languages* (Dagstuhl Seminar 12011) addressed the need for further research to improve this situation.

The 2012 program also included a number of highly interdisciplinary seminars such as *Computation and Incentives in Social Choice* (Dagstuhl Seminar 12101), which brought together computer scientists, mathematicians, social choice theorists, and political scientists. A general tendency towards informatics research in non-informatics areas such as linguistics, paleography, biology and medicine could also be detected, for example in the workshop on *Computation and Paleography: Potentials and Limits* (Dagstuhl Perspectives Workshop 12382). At least one community, that on *Artificial and Computational Intelligence in Games* (Dagstuhl Seminar 12191), came to Dagstuhl for the first time in 2012.

The year ended with the thought-provoking workshop on *Publication Culture in Computing Research* (Dagstuhl Perspectives Workshop 12452), which addressed the need to re-focus the purposes and practices surrounding publication culture in today's computer science community.

In general, Dagstuhl's 2012 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

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. Im Jahr 2012 nahmen 7 Forschungsgäste diese Gelegenheit wahr.

Qualitätssicherung

2.8

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 per Email wöchentlich allen Abteilungen zugänglich gemacht und so eine schnelle Reaktion auf Probleme und Wünsche erreicht. Gleichzeitig werden die anonymisierten Ergebnisse von inhaltlichen Fragen den Teilnehmern eines Seminars per Email 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.

Fig. 2.10 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2012 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 174 Fragebögen, die die Meinung von 50% der Teilnehmer repräsentieren (2 346). Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

Auslastung des Zentrums

2.9

Die Verdichtung des Seminarprogramms trug dazu bei, dass Schloss Dagstuhl mit 12 898 Übernachtungen insgesamt und 10 256 im Rahmen von Seminaren und Workshops mehr Übernachtungen verzeichnen konnte als jemals zuvor. Das Zentrum beherbergte 126 Veranstaltungen mit 3 482

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

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. In 2012, seven research guests availed themselves of this opportunity.

Quality Assurance

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 via e-mail, 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.

Fig. 2.10 shows the satisfaction of responding participants in 2012 with regard to selected aspects of their stay. The results were compiled from 1 174 questionnaires, representing the responses of 50% of all participants (2 346). These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

Utilization of the Center

Thanks in part to the intensification of the Dagstuhl Seminar program following the opening of the new guest house, in 2012 Schloss Dagstuhl had more overnight stays (12 898) and more overnight stays in seminars and workshops (10 256) than ever before. The center hosted a total

Gästen. Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2012 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 November 2014 und bis einschließlich Februar 2014 für drei- und fünftägige Seminare mit 30 Teilnehmern ausgebucht.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2012 einschließlich Dagstuhl-Seminare, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminare und Veranstaltungen wie Tagungen und Sommer-schulen, bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Website (<http://www.dagstuhl.de>) kann unser Kalender mit anstehenden Veranstaltungen eingesehen werden ebenso wie weitere Informationen und Materialien wie z.B. Ziele und thematischer Fokus, Teilnehmerlisten und abschließende Berichte zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

of 126 events with 3 482 guests; see Chapter 13 for further details.

Weekends were kept free in 2012, 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 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 November of 2014 for 5-day, 45-person seminars and through February of 2014 for 3- and 5-day, 30-person seminars.

A comprehensive listing of all 2012 events at Schloss Dagstuhl in 2012, 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 (<http://www.dagstuhl.de>) to view our calendar 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.

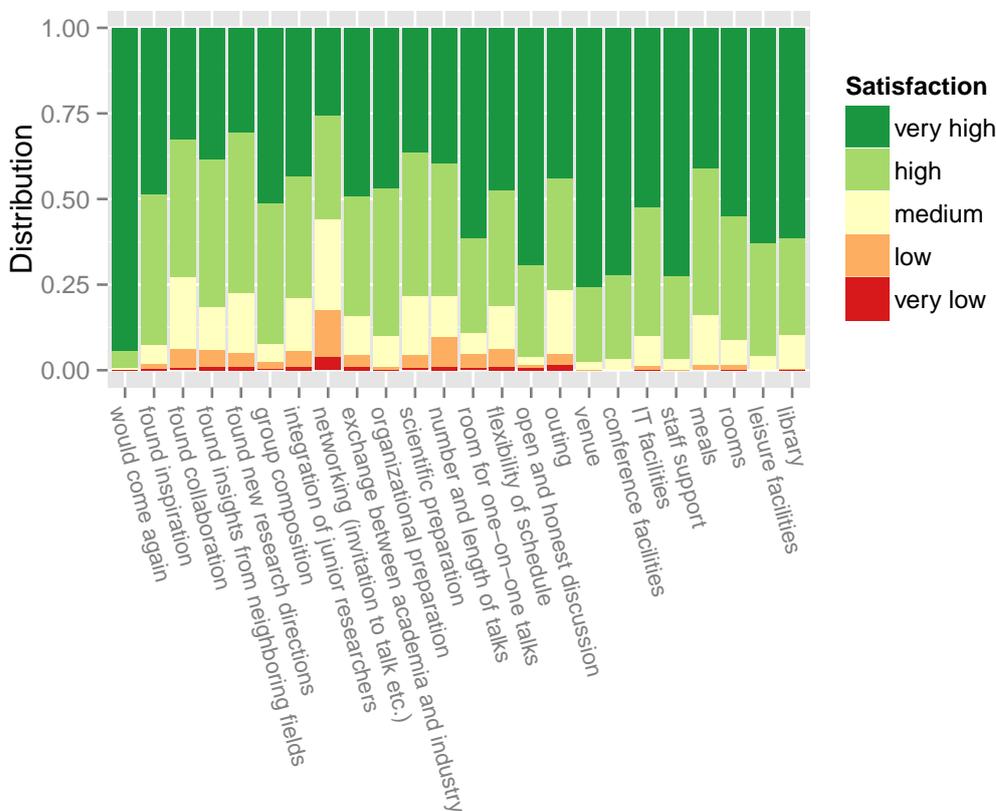


Fig. 2.10 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2012, according to survey results

3 Resonanz *Feedback*

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 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 and write us about their impressions. What follows is an excerpt of our large thank-you collection.

Albert Alonso (Hospital Clinic de Barcelona – Barcelona, ES)

12231 – Future Internet for eHealth | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

Thanks for your hard work in organising and running such an inspiring and thought provoking meeting in these wonderful Dagstuhl facilities. I deeply enjoyed it and look forward to future initiatives within the group.

Veronica Becher (University of Buenos Aires, AR)

12021 – Computability, Complexity and Randomness | Dagstuhl Seminar | <http://www.dagstuhl.de/12021>

The seminar had 50 participants, including the most recognized senior specialists as well as young researchers. The atmosphere was very stimulating and led to new research contacts and collaborations. The seminar was well received, as witnessed by the high rate of accepted invitations, and the exemplary degree of involvement by the participants. Due to the broad scope and depth of the problems on algorithmic randomness and information quantity that have been discussed in the presentations and informal discussions, the organizers regard the seminar as a great success. The organizers wish to express their gratitude towards the Scientific Directors of the Dagstuhl Center for their support of this seminar. We foresee the proposal of a new seminar focusing in the interplay between algorithmic randomness and computable analysis.

Robert Chansler (Linkedin, US)

12282 – Database Workload Management | Dagstuhl Seminar | <http://www.dagstuhl.de/12282>

I very much appreciated the opportunity to attend the Dagstuhl seminar! It was an excellent event, and I commend the staff for the well-managed facilities.

Marc Dacier (Symantec Research Labs – Sophia Antipolis, FR)

12502 – Securing Critical Infrastructures from Targeted Attacks | Dagstuhl Seminar | <http://www.dagstuhl.de/12502>

Thanks also for the great organization and the logistics. Always a pleasure to come to Dagstuhl and I look forward for my next visit.

Mirjam P. Eladhari (University of Malta, MT)

12191 – Artificial and Computational Intelligence in Games | Dagstuhl Seminar | <http://www.dagstuhl.de/12191>

It was a truly inspiring week. I found myself making connections and associations in my thoughts that will most likely guide my future work. It is as if getting together like this suddenly makes one function better, as if the smartness of the group somehow enhances the individual. The organisation team, Simon Lucas, Michael Mateas, Mike Preuss, Pieter Spronck and Julian Togelius, did an amazing job, aided by the smooth machinery of Schloss Dagstuhl. Thank you :)

Paul Groth (VU University Amsterdam, NL)

12091 – Principles of Provenance | Dagstuhl Seminar | <http://www.dagstuhl.de/12091>

<http://thinklinks.wordpress.com/2012/03/07/thoughts-from-the-dagstuhl-principles-of-provenance-workshop/>

Dagstuhl [is] a unique intellectually intense environment. It's one of the nicest traditions in computer science. This was a great workshop for me. I wanted to thank the organizers for putting it together. It's a lot of effort. Additionally, thanks to all of the participants for really great conversations.

Nick Guldemon (TU Delft, NL)12231 – Future Internet for eHealth | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

Thank you so much for the inspiring meeting. I never experienced such a synergetic multidisciplinary med-tech get-together. High spirits to make a difference.

Rainer Herzog (Siemens – München, DE)12231 – Future Internet for eHealth | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

Thanks again for very inspiring days at Dagstuhl. Well done! I am looking forward to future sessions.

Koen V. Hindriks (TU Delft, NL)12342 – Engineering Multi-Agent Systems | Dagstuhl Seminar | <http://www.dagstuhl.de/12342>

The unique atmosphere of Dagstuhl provided the perfect environment for leading researchers from a wide variety of backgrounds to discuss future directions in programming languages, tools and platforms for multiagent systems, and the road map produced by the seminar will have a timely and decisive impact on the future of this whole area of research.

Clare Hooper (University of Southampton, UK)12351 – Interaction Beyond the Desktop | Dagstuhl Seminar | <http://www.dagstuhl.de/12351>
<http://www.clarehooper.net/blog/2012/09/dagstuhl-initial-impressions/>

That was a very intense 5 days. It turns out that what everyone told me pre-seminar – Dagstuhl Seminars are really special, you must go – was absolutely right!

Clifford B. Jones (Newcastle University, GB)12271 – AI meets Formal Software Development | Dagstuhl Seminar | <http://www.dagstuhl.de/12271>

It is a pleasure to extend our thanks to everyone involved in the Dagstuhl organisation: they provide a supportive and friendly context in which such fruitful scientific exchanges can develop unhindered by distraction.

Jussi Kangasharju (University of Helsinki, FI)12363 – Software Defined Networking | Dagstuhl Seminar | <http://www.dagstuhl.de/12363>

The seminar was well received by the participants. Among the participants there were also organizers of future SDN workshops (IRTF SDN and DIMACS SDN) who signaled the intent of building their workshops around the similar discussion-oriented structure preferred at Dagstuhl.

Harumi Kuno (HP Labs – Palo Alto, US)12282 – Database Workload Management | Dagstuhl Seminar | <http://www.dagstuhl.de/12282>

I heard many compliments about the excellent Dagstuhl staff and facilities from the participants while we were there. Dagstuhl is really a rare and wonderful place, that we appreciate very much . . . Thank you!

Lenka Lhotská (Czech Technical University, CZ)12231 – Future Internet for eHealth | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

Thank you very much for organizing such a great seminar. It was really a great pleasure to be there together with all participants. I think we all enjoyed it.

Angela Lozano (UC Louvain-la-Neuve, BE)

12071 – Software Clone Management Towards Industrial Application | Dagstuhl Seminar | <http://www.dagstuhl.de/12071>
<http://sites.uclouvain.be/blog-ingi/2012/04/the-dagstuhl-experience/>

My experience at the Dagstuhl seminar has been inspiring, and energizing. It gave me time to reflect on problems that I find interesting, and to identify new research targets. It also allowed me to discuss with industry representatives on their priorities and issues related to source code clones. Finally, it allowed me to interact with several academics that I only knew by their publications.

Mark Manulis (University of Surrey – Guildford, GB)

12381 – Privacy-Oriented Cryptography | Dagstuhl Seminar | <http://www.dagstuhl.de/12381>

The organizers would like to thank all participants for accepting our invitations and attending the seminar, and for sharing their ideas and contributing to the interesting seminar program. We hope that discussions were fruitful and the opportunity to work face-to-face during the seminar helped to create impulses for exciting new research projects, paving the way for further progress and new discoveries in Privacy-Oriented Cryptography. Finally, the organizers, also on behalf of the participants, would like to thank the staff and the management of Schloss Dagstuhl for their support throughout the 1,5 years of preparations of this very pleasant and successful event.

Kevin Patrick (UCSD, US)

12231 – Future Internet for eHealth | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

My kudos as well to you for organizing a superb meeting. It was great to make so many new connections – and to spend a few days with such a wonderful group of fellow-travelers on the road to the future. I look forward to follow-up activities and will do my best to help in any way that I can.

Wolter Pieters (TU Delft, NL)

12501 – Organizational Processes for Supporting Sustainable Security | Dagstuhl Seminar | <http://www.dagstuhl.de/12501>

Thanks for another great Dagstuhl seminar! The discussions and interactions were excellent again, and I have a lot of ideas to think about.

Mike Preuss (TU Dortmund, DE)

12191 – Artificial and Computational Intelligence in Games | Dagstuhl Seminar | <http://www.dagstuhl.de/12191>

As organizers we are really pleased with how the Seminar turned out. It proved to be the stimulating and inspirational environment that we had hoped for. We found that most, if not all participants agreed with us on that. A lot of this success is due to the excellent facilities provided by the people of Schloss Dagstuhl. We are highly grateful for having had the opportunity to be their guests for the Seminar. We definitely hope to return in the future.

Georg Regensburger (RICAM – Linz, AT)

12462 – Symbolic Methods for Chemical Reaction Networks | Dagstuhl Seminar | <http://www.dagstuhl.de/12462>

Vielen Dank nochmal für die Mitorganisation des Seminars, es war wirklich eine sehr spannende Woche!

Fabrice Rossi (TELECOM-Paristech, FR)

12081 – Information Visualization, Visual Data Mining and Machine Learning | Dagstuhl Seminar | <http://www.dagstuhl.de/12081>

Judging by the liveliness of the discussions and the number of joint research projects proposed at the end of the seminar, this meeting between the machine learning and the information visualization communities was more than needed. The flexible format of the Dagstuhl seminars is perfectly adapted to this type of meeting and the only frustration perceivable at the end of the week was that it had indeed reached its end. It was clear that researchers from the two communities were starting to understand each other and were eager to share more thoughts and actually start working on joint projects. This calls for further seminars ...

Martijn Stam (University of Bristol, GB)

12031 – Symmetric Cryptography | Dagstuhl Seminar | <http://www.dagstuhl.de/12031>
<http://bristolcrypto.blogspot.de/2012/01/dagstuhl-symmetric-cryptography-seminar.html>

I had heard a lot about Dagstuhl (and its mathematical cousin Oberwolfach) in the past, so I was quite excited about going. The organizers had done a sterling job in picking a good mix of cryptographers, so that symmetric crypto could be covered from all relevant angles. The program was representative of this and what I really liked was the blend of presentations: there was work in various stages of the lifecycle (recently published work, accepted work awaiting publication, submitted work, and work still very much in progress); there were slick powerpoint/beamer presentations as well as black board presentations (prepared or impromptu); and the topics represented the attendees well.

Frank Steinicke (Universität Würzburg, DE)

12151 – Touching the 3rd Dimension | Dagstuhl Seminar | <http://www.dagstuhl.de/12151>

Ich möchte mich nochmals ausdrücklich bei Ihnen allen für die ausgezeichnete Unterstützung während der Vorbereitungen aber auch während der Durchführung unseres Seminars bedanken. Wir haben die Zeit sehr genossen und sehr spannend diskutieren können und es sind viele neue Ideen gesammelt worden.

Stefan Strecker (FernUniversität in Hagen, DE)

12131 – Open Models as a Foundation of Future Enterprise Systems | Dagstuhl Seminar | <http://www.dagstuhl.de/12131>

Im Namen der Teilnehmerinnen und Teilnehmer des Dagstuhlseminars 12131 möchte ich mich sehr herzlich bei Ihnen und dem gesamten Team in Dagstuhl für die hervorragende Betreuung bedanken. Wir haben die Zeit in Dagstuhl sehr genossen und fahren mit frischen Ideen nach Hause zurück.

Christian Theobalt (MPI für Informatik – Saarbücken, DE)

12431 – Time-of-Flight Imaging: Algorithms, Sensors and Applications | Dagstuhl Seminar | <http://www.dagstuhl.de/12431>

The seminar was very successful with respect to the set goals and initiated great interaction between researchers from different domains which had never happened in this way at other conferences or workshops.

Joseph F. Traub (Columbia University, US)

12391 – Algorithms and Complexity for Continuous Problems | Dagstuhl Seminar | <http://www.dagstuhl.de/12391>

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

Heribert Vollmer (Leibniz Universität Hannover, DE)

12471 – SAT Interactions | Dagstuhl Seminar | <http://www.dagstuhl.de/12471>

The organizers regard the workshop as a great success. Bringing together researchers from different areas of theoretical computer science fostered valuable interactions and led to fruitful discussion. 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 *SAT Interactions* in the future.

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.

4

Die Seminare in 2012

The 2012 Seminars

<p>Data Structures, Algorithms, Complexity</p> <ul style="list-style-type: none"> ■ Computability, Complexity and Randomness (12021) ■ Learning in Multiobjective Optimization (12041) ■ Computation and Incentives in Social Choice (12101) ■ Data Reduction and Problem Kernels (12241) ■ Algorithms and Complexity for Continuous Problems (12391) ■ Algebraic and Combinatorial Methods in Computational Complexity (12421) ■ The Constraint Satisfaction Problem: Complexity and Approximability (12451) ■ SAT Interactions (12471) 	<p>Verification, Logic, Formal Methods, Semantics</p> <ul style="list-style-type: none"> ■ Software Clone Management Towards Industrial Application (12071) ■ Software Synthesis (12152) ■ Verifying Reliability (12341) ■ Information Flow and Its Applications (12352) ■ Coalgebraic Logics (12411) ■ Games and Decisions for Rigorous Systems Engineering (12461) ■ Divide and Conquer: the Quest for Compositional Design and Analysis (12511) 	<p>Geometry, Image Processing, Graphics</p> <ul style="list-style-type: none"> ■ Information Visualization, Visual Data Mining and Machine Learning (12081) ■ Touching the 3rd Dimension (12151) ■ Putting Data on the Map (12261) ■ Biological Data Visualization (12372) ■ Time-of-Flight Imaging: Algorithms, Sensors and Applications (12431) ■ Representation, Analysis and Visualization of Moving Objects (12512)
<p>Artificial Intelligence, Computational Linguistics</p> <ul style="list-style-type: none"> ■ Normative Multi-Agent Systems (12111) ■ Artificial and Computational Intelligence in Games (12191) ■ Cognitive Approaches for the Semantic Web (12221) ■ Engineering multiagent Systems (12342) ■ Foundations and Challenges of Change and Evolution in Ontologies (12441) ■ Interpreting Observed Action (12491) 	<p>Software Technology</p> <ul style="list-style-type: none"> ■ Foundations for Scripting Languages (12011) ■ Analysis of Executables: Benefits and Challenges (12051) ■ Open Models as a Foundation of Future Enterprise Systems (12131) ■ Co-Design of Systems and Applications for Exascale (12212) ■ AI meets Formal Software Development (12271) ■ Architecture-Driven Semantic Analysis of Embedded Systems (12272) ■ Requirements Management – Novel Perspectives and Challenges (12442) 	<p>Distributed Computation, Networks, Architecture</p> <ul style="list-style-type: none"> ■ Network Attack Detection and Defense Early Warning Systems – Challenges and Perspectives (12061) ■ Abstractions for scalable multi-core computing (12161) ■ Future Internet for eHealth (12231) ■ Security and Dependability for Federated Cloud Platforms (12281) ■ Information-centric networking – Ready for the real world? (12361) ■ Software Defined Networking (12363) ■ Is the Future of Preservation Cloudy? (12472)
<p>Cryptography, Security</p> <ul style="list-style-type: none"> ■ Symmetric Cryptography (12031) ■ Machine Learning Methods for Computer Security (12371) ■ Privacy-Oriented Cryptography (12381) ■ Web Application Security (12401) ■ Quantitative Security Analysis (12481) ■ Analysis of Security APIs (12482) ■ Organizational Processes for Supporting Sustainable Security (12501) ■ Securing Critical Infrastructures from Targeted Attacks (12502) 	<p>Data Bases, Information Retrieval, Data Mining</p> <ul style="list-style-type: none"> ■ Principles of Provenance (12091) ■ Semantic Data Management (12171) ■ Database Workload Management (12282) ■ Robust Query Processing (12321) ■ Mobility Data Mining and Privacy (12331) ■ The Multilingual Semantic Web (12362) 	<p>Applications, Interdisciplinary Work</p> <ul style="list-style-type: none"> ■ Applications of Combinatorial Topology to Computer Science (12121) ■ Quality of Experience: From User Perception to Instrumental Metrics (12181) ■ Social, Supply-Chain, Administrative, Business, Commerce, Political Networks: a Multi-Discipline Perspective (12182) ■ Structure Discovery in Biology: Motifs, Networks & Phylogenies (12291) ■ Interaction Beyond the Desktop (12351) ■ Computation and Palaeography: Potentials and Limits (12382) ■ Publication Culture in Computing Research (12452) ■ Symbolic Methods for Chemical Reaction Networks (12462) ■ Human Activity Recognition in Smart Environments (12492)

4.1 Foundations for Scripting Languages

Organizers: Robert Hirschfeld, Shriram Krishnamurthi, and Jan Vitek
Seminar No. 12011

Date: 02.–06. January, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.1.1

© Creative Commons BY-NC-ND 3.0 Unported license

© Robert Hirschfeld, Shriram Krishnamurthi, and Jan Vitek

Participants: Amal Ahmed, Carl Friedrich Bolz, Ravi Chugh, William R. Cook, Theo D'Hondt, Matthias Felleisen, Robert Bruce Findler, Cormac Flanagan, Matthew Flatt, Jeffrey Foster, Andreas Gal, Philippa Gardner, Kathryn E. Gray, Arjun Guha, Stefan Udo Hanenberg, David Herman, Manuel Hermenegildo, Michael Hicks, Robert Hirschfeld, Atsushi Igarashi, Shriram Krishnamurthi, Ralf Lämmel, Benjamin Lerner, Jens Lincke, Hidehiko Masuhara, Mark S. Miller, Floreal Morandat, Oscar M. Nierstrasz, Nathaniel Nystrom, Yoshiki Ohshima, Joe Politz, Gregor Richards, Manuel Serrano, Jeremy G. Siek, Bastian Steinert, Ankur Taly, Eric Tanter, Sam Tobin-Hochstadt, Laurence Tratt, Herman Venter, Jan Vitek, Philip Wadler, Francesco Zappa Nardelli



Common characteristics of scripting languages include syntactic simplicity, a lack of onerous constraints for program construction and deployment, the ability to easily connect to and control systems processes, strong built-in interfaces to useful external objects, extensive library support, and lightweight (and embeddable) implementations. More broadly, these characteristics add up to strong support for effective software prototyping. Due to a combination of these characteristics, common scripting languages like Perl, Python, Ruby, JavaScript, Visual Basic, and Tcl have moved from the fringes to mainstream program development.

To academics, these languages do not appear that different from, say, Scheme or ML. Since languages like Scheme and ML have well-defined semantics and other formal attributes, the mainstream passion for scripting languages may appear to simply be the result of ignorance of better languages amongst mainstream developers. However, the properties that scripting language users claim to find most beneficial are often *not* found in their more academic counterparts, such as a strong orientation towards systems process management, easily extensible objects, specific but useful control operators, etc.

In short, the academic tendency towards reductionism appears to miss some important characteristics. In particular, properties that may appear incidental—and are ignored by the formalization of academic languages—may actually be essential. As a result, the formal study of scripting languages is a worthwhile research activity in its own right.

Not only does the study of scripting offer academics fresh problems, their results have the potential for widespread benefit. As scripts grow into programs, the very characteristics that seem an

advantage sometimes prove to be disadvantages. If any object can be extended by any other object, it is impossible to reason about its behavior. If any object can access any resources, it is impossible to bound security implications. If programmers can place values of any type into a variable, it is impossible to obtain type guarantees. And so on. In other words, the very flexibility that enables prototyping inhibits the reasoning necessary for programs to grow in scale.

In the early days of scripting, there was an expectation that scripts were not meant to “grow up”. Rather, as a prototype proved valuable, it would be turned into a program in a mainstream language, such as Java. However, reality does not match this vision. First, once a system becomes valuable to an organization, it is not possible to halt development on it while waiting for a full re-implementation. Second, even if the current version is converted to Java, the next version would probably still benefit from the benefits of prototyping. Thus, in both cases, programs that start in a scripting language are likely to remain in it. Finally, even if clients do want to rewrite the program in a more mature language, they would benefit from formal support to enable this conversion.

As a result, the formal study of scripting languages is a worthwhile research activity in its own right. In particular, we hope this seminar had both direct and indirect impact on academia and industry. We also hope that, based on our discussions, academics will identify concrete problems that need solutions and find scripting language experts who they can communicate with. In turn, we hope scripting experts identified knowledge, expertise, and interest from academia and are better aware of how to formulate problems for academics and map their solutions back to practice.

4.2 Computability, Complexity and Randomness

Organizers: Verónica Becher, Laurent Bienvenu, Rodney Downey, and Elvira Mayordomo
Seminar No. 12021

Date: 08.–13. January, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.1.19

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Verónica Becher, Laurent Bienvenu, Rodney Downey, and Elvira Mayordomo



Participants: Eric Allender, Klaus Ambos-Spies, George Barmaplias, Bruno Bauwens, Verónica Becher, Laurent Bienvenu, Harry Buhrman, Douglas Cenzer, Chris J. Conidis, Quinn Culver, David Diamondstone, Rodney Downey, Lance Fortnow, Johanna N. Y. Franklin, Cameron Freer, Noam Greenberg, Serge Grigorieff, Pablo A. Heiber, John Hitchcock, Rupert Hözl, Michal Koucký, Thorsten Kräling, Antonin Kucera, Sophie Laplante, Andrew Lewis, Bruno Loff, Elvira Mayordomo, Wolfgang Merkle, Joseph S. Miller, Benoit Monin, Philippe Moser, Satyadev Nandakumar, Andre Nies, Sylvain Perifel, Christopher P. Porter, Robert Rettinger, Andrej E. Romashchenko, Ronen Shaltiel, Alexander Shen, Theodore A. Slaman, Ludwig Staiger, Antoine Tavenaux, Leen Torenvliet, Daniel Turetsky, Vinodchandran N. Variyam, Stijn Vermeeren, Paul M. B. Vitanyi, Vladimir Viyugin, Osamu Watanabe, Marius Zimand

Randomness and information quantity are central notions in computer science that are still undeveloped. Although classical information theory and probability provide formalizations of these notions they do not allow us to measure the information of a specific string or say that a particular real number is random. The definition of the property of randomness and its connection with a measure of information content was obtained in the 1960s and combines different complexity measures.

As witnessed by the three seminars previously organized in Dagstuhl on complexity and randomness (Seminar 9318, *Descriptive complexity: a multidisciplinary perspective* in 1993; Seminar 03181, *Centennial Seminar on Kolmogorov Complexity and Applications* in 2003; and Seminar 06051 *Kolmogorov Complexity and Applications* in 2006) in recent years there has been an upsurge produced by the people in computability theory that resulted in rapid progress in our understanding of even the most basic notions in randomness, and the solution of old open questions. This has changed and is still changing the landscape and opened up new avenues of research. An evidence of this activity has been the publication of two new books in the area and the new edition of an already classical one: *Algorithmic Randomness and Complexity*, R. Downey and D. Hirschfeldt, Foundations on Computing, Springer, 2010; *Computability and Randomness*, A. Nies, Oxford University Press, 2009; and *An Introduction to Kolmogorov Complexity and Its Applications*, M. Li and P. Vitanyi, third Edition, Springer Verlag, 2008.

Seminar 12021 has celebrated significant recent research progress. New results connect the theory of algorithmic randomness with computable analysis. We consider them important because they lead to the naturalness of the notions of algorithmic randomness. For instance, Brattka, Miller, and Nies translated the theorem “every non-decreasing function is almost everywhere differentiable” to the computable world, by showing that a real x is computably random if and only if every computable non-decreasing function is differentiable at x (this work is has not yet appeared as a publi-

cation). Similar investigations identified the notions of randomness that correspond to the Lebesgue density and differentiation theorems. J. Franklin and the work of Gács, Hoyrup, and Rojas related Birkhoff’s pointwise ergodic theorem in connection with Schnorr randomness.

Considerable results have been obtained for problems on Kolmogorov complexity and computable enumerable sets, in particular, in the degree structure that arises from comparing the complexity of the initial segments of two reals. Barmaplias announced the solution of the already long standing open problem posed by Downey and Hirschfeldt *Is there a minimal pair of c.e. reals in the K -degrees?* The answer is no.

Since the start of the discipline, the notion of randomness was defined for infinite sequences, or real numbers. The problem posed by Kolmogorov on a notion of randomness of finite objects remains unsolved. This is also the case for arbitrary countable objects. C. Freer made significant progress on the questions *When is a graph random?* and *What is the connection between quasi-random graphs and pseudorandom bit strings?* He pointed to an emerging theory of continuous limits of finite combinatorial structures that connects graph limits, property testing, and exchangeable relations.

There was a general consensus on the fact that there is yet no adequate solution to the fundamental problem that high-quality independent random bits are in very short supply. And there are many practical applications rely on randomness (for instance, assigning keys to users of a public-key crypto-system). Randomness extractors are algorithms developed “extract” high-quality random bits from low-entropy sources. Construction of such algorithms is foreseen to be an active research area.

The aim of Seminar 12021 was to bring together researchers covering this spectrum of relevant areas, to report their advances and to discuss the relevant research open questions. The seminar had 50 participants, including the most recognized senior specialists as

well as young researchers. The atmosphere was very stimulating and led to new research contacts and collaborations.

Concluding remarks and future plans. The seminar was well received, as witnessed by the high rate of accepted invitations, and the exemplary degree of involvement by the participants. Due to the broad scope and depth of the problems on algorithmic randomness and information quantity that have been discussed in the presentations and informal discussions, the organizers regard the seminar as a great success. The organizers wish to express their gratitude towards the Scientific Directors of the Dagstuhl Center for their support of this seminar. We foresee the proposal of a new seminar focusing in the interplay between algorithmic randomness and computable analysis.

■ Description of the seminar topics

Anti-randomness The class of sequences with minimal prefix-free Kolmogorov complexity, dubbed K -trivial sequences, were understudied until five years ago. In the seventies, Solovay proved that there is a non computable K -trivial. They are now very well understood, with a number of surprising characterizations and applications. For instance, the “cost function” construction of a K -trivial gives simplest known example of a non computable incomplete computably enumerable set, they also appear in the Kucera-Slaman solution to a well-known question about Turing degrees in Scott sets, also K -triviality has led to a better understanding of the reverse mathematics of the regularity of Lebesgue measure. K -triviality one of the most technically deep subjects in algorithmic randomness, significant questions remain open.

Resource bounded versions Classical computational complexity theory comes into play defining resource-bounded versions of Kolmogorov complexity, measure, and dimension. This has led to new characterizations of complexity classes involving efficient reducibility to the set of Kolmogorov random strings. Resource-bounded measure and dimension have been used to gain understanding of properties of complexity classes and their complete sets. For instance, they can be used as a probabilistic methods to prove lower bounds on nonuniform complexity.

Derandomization and complexity hierarchies Derandomization is the study of how to replace probabilistic algorithms with deterministic algorithms. Earlier work by Allender et al. showed that the techniques of derandomization could be viewed through the lens of resource-bounded Kolmogorov complexity theory, and gave significant applications. More recently, they proved that every sufficiently dense set in $NP \cap coNP$ contains strings of low resource-bounded Kolmogorov complexity at every length. In still unpublished work, Allender and his co-authors show that if deterministic and nondeterministic exponential time coincide, this implies

a partial collapse of the exponential-time hierarchy, shedding light on a question that has been open for two decades.

Randomness extractors Randomness extractors have been used and to derive zero-one laws for the packing dimensions of complexity classes and Turing degrees. Recently it has been shown that the converse direction also holds and Kolmogorov extraction is in fact equivalent to randomness extraction.

Computational depth The computational depth of a string is roughly the difference between its time-bounded Kolmogorov complexity, and its (plain) Kolmogorov complexity. Quite recently, Antunes and Fortnow showed that, under a plausible complexity assumption, computational depth is the right notion to present a “universal” poly-time samplable distribution, in the same way that Kolmogorov complexity allows one to define universal computable semi-measures. They derive a new characterization of algorithms that run in polynomial time on average, and give a relation with their worst-case running time.

Algorithmic randomness and computable analysis The most accepted definition of randomness for infinite sequences, or real numbers, is based on constructive measure theory and was given by Martin Löf, 1965. It coincides with the maximal initial segment complexity. Other notions have been proposed since then, by Schnorr, Demuth, Kurtz and others, either via measure theory, or via martingale theory. Most of these definitions have been very well studied in the space of infinite binary sequences, but less in known for other spaces (although there has been some deep founding work by Levin and Gács). Some natural questions are: for a given randomness notion, to what kind of probability space can this notion be extended? To what extent does the chosen space affect the properties of random objects? Then, for every probability space to which we can extend randomness notions, it is interesting to look at classical theorems from a randomness perspective, and try to convert classical theorems of the form “property P holds for μ -almost every sequence” into “property P holds for every μ -random sequence”. This line of study has recently been investigated in a number of different settings: random closed sets, effective ergodic theory, effective brownian motion, etc.

■ Organization of the seminar and activities

The seminar consisted in nineteen talks, sessions on open questions, and informal discussions among the participants. The organizers selected the talks in order to have comprehensive lectures giving overview of main topics and communications of new research results. Each day consisted of talks and free time for informal gatherings among participants. There were two main sessions on open questions.

4.3 Symmetric Cryptography

Organizers: Frederik Armknecht, Stefan Lucks, Bart Preneel, and Phillip Rogaway
Seminar No. 12031

Date: 15.–20. January, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.1.39

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Frederik Armknecht, Stefan Lucks, Bart Preneel, and Philip Rogaway



Participants: Elena Andreeva, Frederik Armknecht, Jean-Philippe Aumasson, Daniel J. Bernstein, Eli Biham, Alex Biryukov, Andrey Bogdanov, Joan Daemen, Itai Dinur, Yevgeniy Dodis, Orr Dunkelman, Marc Fischlin, Ewan Fleischmann, Christian Forler, Matthias Hamann, Tetsu Iwata, Antoine Joux, Lars Ramkilde Knudsen, Matthias Krause, Rudolphe Lampe, Gregor Leander, Stefan Lucks, Florian Mendel, Vasily Mikhalev, Tilo Müller, Maria Naya-Plasencia, Kaisa Nyberg, Jacques Patarin, Kenneth G. Paterson, Bart Preneel, Christian Rechberger, Phillip Rogaway, Sondre Ronjom, Greg Rose, Yu Sasaki, Adi Shamir, Kyoji Shibutani, Martijn Stam, John Steinberger, Deniz Toz, Kerem Varici, Bogdan Warinschi, Jakob Wenzel, Kan Yasuda, Erik Zenner

Research in Symmetric Cryptography is quickly evolving. The seminar was the third of its kind, the first one took place in 2007, the second in 2009. We observe a steadily increasing interest in Symmetric Cryptography, as well as a growing practical demand for symmetric algorithms and protocols. The seminar was very successful in discussing recent results and sharing new ideas. Furthermore, it inspired the participants to consider how Symmetric Cryptography has evolved in the past, and how they would like it to evolve in the future.

Two intense discussions dealt with Authenticated Encryption and the issue of a ‘valid’ attack on a symmetric primitive. The participants agreed on Authenticated Encryption becoming a major research topic for Symmetric Cryptography in the next few years, because current Authenticated Encryption Schemes are not always suitable for practical demands – especially are the relevant attack modes and models not yet well-understood (e.g., misuse attacks, blockwise adaptive attacks, etc.). Regarding the issue of ‘valid’ attacks, the participants agreed that the current development of academic cryptanalysis with a growing number of increasingly ‘marginal’ attacks, is unsatisfactory.



Fig. 4.1
Mane Hellenthal – Bergbild. Part of the Dagstuhl art collection and donated by: Roland and Ute Vollmar, Marc Herbstritt, Reinhard Wilhelm, Holger Schlingloff, Kirk Pruhs, Kurt Mehlhorn, Rainer Koschke, Kathi Fisler, Shriram Krishnamurthi, and participants of the Dagstuhl Seminars 12271, 12281, 13051, 13062, and 13181.

4.4 Learning in Multiobjective Optimization

Organizers: Salvatore Greco, Joshua D. Knowles, Kaisa Miettinen, and Eckart Zitzler
Seminar No. 12041

Date: 23.–27. January, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.1.50

© Creative Commons BY-NC-ND 3.0 Unported license

© Salvatore Greco, Joshua D. Knowles, Kaisa Miettinen, and Eckart Zitzler



Participants: Richard Allmendinger, Anne Auger, Jürgen Branke, Heinrich Braun, Dimo Brockhoff, Susanne Buerklen, Carlos A. Coello Coello, Salvatore Corrente, Kalyanmoy Deb, Matthias Ehrgott, Michael Emmerich, Jose Rui Figueira, Joerg Fliege, Carlos M. Fonseca, Simon French, Tobias Friedrich, Xavier Gandibleux, Antonio Gaspar-Cunha, Martin Josef Geiger, Salvatore Greco, Jussi Hakanen, Markus E. Hartikainen, Hisao Ishibuchi, Milosz Kadzinski, Kathrin Klamroth, Joshua D. Knowles, Pekka Korhonen, Karlheinz Kuefer, Manuel López-Ibáñez, Mariano Luque, Kaisa Miettinen, Sanaz Mostaghim, Vincent Mousseau, Boris Naujoks, Amos H. C. Ng, Akira Oyama, Dmitry Podkopaev, Silvia Poles, Enrico Rigoni, Günter Rudolph, Pradyumn Kumar Shukla, Johannes Siebert, Karthik Sindhya, Roman Słowiński, Theodor J. Stewart, El-Ghazali Talbi, Tamara Ulrich, Jyrki Wallenius, Margaret M. Wiecek, Katrin Witting, Eckart Zitzler

Multiobjective optimization is the study of optimization under competing interests, goals or criteria; it concerns the search for *nondominated* solutions (or Pareto optima) that offer different trade-offs of the competing criteria, as well as methods for choosing among the alternative solutions by the consideration of *preferences*. Multiobjective optimization problems arise naturally in several areas: engineering, economics, operations research/management, and the natural sciences, and today a significant portion of research into optimization is concerned with these problems. The present seminar, the fourth in a series on Multiobjective Optimization (following 04461, 06501 and 09041) dating back to 2004, renewed its ambitions to unite researchers from the two main communities studying multiobjective optimization, MCDM (multiple criteria decision making) and EMO (evolutionary multiobjective optimization) to stimulate new research directions crossing these discipline boundaries.

As with earlier meetings in the series, we chose a strong theme for the seminar, which this time was *Learning*. In multiobjective optimization, learning has a key role to play because, uniquely to the multiobjective case, optimization involves both an exploration of trade-offs and a consideration of user (or decision maker) *preferences*, which are usually implicit in the mind(s) of decision maker(s) at the start of the solution process. Solving a problem therefore involves at least two simultaneous learning processes: the decision maker (DM) learning about the problem, and the optimization process itself learning about the DM's preferences (to achieve a steering of the search toward a preferred solution). Our aim in the seminar was to focus centrally on this learning aspect to give it, for the first time, due attention, as in previous seminars it arose rather peripherally to other themes.

The seminar took place January 22nd–27th 2012. The main goals of the seminar were to explore in depth three different aspects

of learning in multiobjective optimization which may be briefly summarized as:

Focus 1: User preferences What should be learnt from user interactions and how should user preferences be captured?

Focus 2: Problem understanding What should be learnt about the problem structure and how can useful information for the DM be extracted?

Focus 3: The problem solving process How do we know if a decision maker has learned? How does a decision maker learn? What factors influence how and what a decision maker learns?

Participants were given some written materials [1, 2] prior to the seminar to orient them to these different aspects and to help them prepare relevant contributions to the seminar programme.

During the seminar, the programme was updated on a daily basis to maintain flexibility and, through this system, we were able to give adequate time both to prepared material and to evolving discussions, mostly taking place in working groups. In particular, breakout working groups were organized initially by lottery (to be purposefully disruptive of existing groupings) and then by forming subtopics that individuals could sign up to for the remainder of the week. Six groups emerged in this way. (In the appendix, the complete list of topics suggested can be seen).

The prepared part of the programme included four invited talks of forty-five minutes each and sixteen contributed talks of twenty minutes each. These were spaced to allow time for discussion, and the evenings were kept free to allow further reflection and relaxation.

Other notable events during the week included: (i) an interactive demonstration given by Pekka Korhonen on rationality in decision making, which reminded us all of the limits of human (including our own “expert”) rationality in the face of complex data; (ii) a presentation session to allow us to share details of upcoming events in our research community; and (iii), rather less formally, a wine and

cheese party was offered by Dagstuhl in the name of ESTECO to express appreciation to ESTECO for giving a donation to the Dagstuhl Foundation.

■ Outcomes

The outcomes of each of the working groups can be seen in the sequel, but a number of key findings are worth brief mention:

DM Sense working group outlined the design for a system that could aid decision-makers rationalize their learning and decisions *in natural language* by pulling together both recent and older research in artificial intelligence and decision making systems.

Pareto Sense working group established a critical agenda of research to undertake in learning and knowledge representation of the combined spaces of Pareto sets and fronts.

Quantifying Learning working group formalized a method for quantifying the learning associated with decision makers steering a search process, and compared this with the algorithmic learning that occurs in some key model-learning MCDM methods.

Navigation working group developed a detailed understanding of search and decision making approaches to identify the most-preferred solution among the Pareto-set (termed “Navigation”), using this to categorize current methods, and identify applications.

Representation working group considered learning in multiobjective optimization from a machine perspective, proposing that learning could be viewed as the process of obtaining parsimonious representations that enable efficient query-answering in support of (particular) search algorithms or decision processes.

Algorithm Design Methods working group considered formally how algorithms for search and decision making should be se-

lected based on information about the decision maker, as well as the problem, and were able to produce first bounds on the number of function evaluations and queries to a decision maker needed to solve a problem.

These findings were reported to the main group during the seminar, and led to lively debate. Further work within the groups (by email correspondence) following the end of the seminar is planned, including several proposals for joint conference and journal papers.

At the wrap-up session of the seminar, we invited written comments from all the participants concerning how the seminar may be improved, what should be maintained, and inviting topics for future seminars. Comments included ‘working groups were a great opportunity to discuss [...] common features from different perspectives’, ‘Not too many talks — very good; staying in focus — very good; atmosphere — very good’, ‘atmosphere ... is very fruitful, encouraging’, and ‘maintain: the diversity of the experts / participants; good balance between presentations and group discussions, like this time’.

In summary, the seminar made for a very productive and enjoyable week. It has revealed a number of research problems that need careful consideration and detailed further study. It has allowed us to begin this work in earnest, and make some significant first steps.

■ Acknowledgments

Many thanks to the Dagstuhl office and its helpful and patient staff; huge thanks to the organizers of the previous seminars in the series for setting us up for success; and thanks to all the participants, who worked hard and were amiable company all week.

We also give special thanks to Kaisa Miettinen and Eckart Zitzler as they step down from the organizer role.

■ References

- 1 S. Greco, J. Knowles, K. Miettinen, E. Zitzler, Dagstuhl Seminar 12041: Learning in Multiobjective Optimization. Seminar Proposal Document, 2011.
- 2 V. Belton, J. Branke, P. Eskelinen, S. Greco, J. Molina, F. Ruiz, R. Słowiński, Interactive multiobjective optimization

from a learning perspective. Chapter 15 [in]: J. Branke, K. Deb, K. Miettinen, R. Słowiński (eds.), *Multiobjective Optimization: Interactive and Evolutionary Approaches*. Springer-Verlag, Berlin, pp. 405–434, 2008.

4.5 Analysis of Executables: Benefits and Challenges

Organizers: Andy M. King, Alan Mycroft, Thomas W. Reps, and Axel Simon

Seminar No. 12051

Date: 29. January | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.1.100

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Axel Simon and Andy King



Participants: Gogul Balakrishnan, Sébastien Bardin, Edward Barrett, Sebastian Biallas, Jörg Brauer, Doina Bucur, Mihai Christodorescu, Bjorn De Sutter, Thomas Dullien, Emmanuel Fleury, Andrea Flexeder, Roberto Giacobazzi, Sean Heelan, Paul Irofti, Johannes Kinder, Andy M. King, Arun Lakhotia, Jerome Leroux, Junghee Lim, Alexey Loginov, Florian Martin, David Melski, Bogdan Mihaila, Barton P. Miller, Martin Murfitt, Alan Mycroft, Magnus Myreen, Michael Petter, Thomas W. Reps, Xavier Rival, Edward Robbins, Daniel Roelker, Andrew Ruef, Alexander Sepp, Holger Siegel, Axel Simon, Fausto Spoto, Aditya Thakur, Christopher Vick, Aymeric Vincent, Andrew Walenstein, Florian Zuleger

The analysis of executables is concerned with extracting information from a binary program typically, though not exclusively, with program analysis techniques based on abstract interpretation. This topic has risen to prominence due to the need to audit code, developed by third parties for which the source is unavailable. Moreover, compilers are themselves a source of bugs, hence the need to scrutinise and systematically examine executables.

■ Seminar topics

The theme of the analysis of executables is an umbrella term adopted for this seminar, covers, among other things, the following topics:

- specifying the semantics of native instructions, intermediate languages and the synthesis of transfer functions from blocks of instructions;
- abstract domains for binary analysis and how to combine them; type synthesis;
- control-flow graph (CFG) reconstruction, which is a prerequisite for many program analysis, and CFG matching, which is useful for detecting piracy;
- self-modifying code, characterising its semantics and detecting malware.

■ Chronological overview of the discussion

For practical reasons, all talks on Monday were held by the four organizers, including an overview of various known tools created by Thomas Reps and his group. His talk was followed by synthesis of transfer functions (the semantics of basic blocks) using SAT solving by Andy King, type reconstruction by Alan Mycroft and the combination of several abstract domains by Axel Simon. These

rather varied topics gave a good introduction. Thomas Reps suggested that we identify common goals through a group discussion, which we could not complete on Monday due to the lack of time. Instead, we scheduled mostly industrial talks on Tuesday in order to find out about the problems that security engineers face in their everyday work and which tools they developed themselves. With this information, a group discussion on Tuesday afternoon quickly raised specific issues and their priorities: analyses must be scalable, preferably to some 12.5 billion instructions that large and vulnerable applications such as Adobe Reader are comprised of. This focus begs the question of whether we can afford a sound analysis or, as was suggested on the last day of the talk on CFG reconstruction, if an engineer can afford to work on a CFG in which not all indirect jumps are resolved precisely. In general, we should be aware of what assumptions we are making, for instance, about the correctness of CPU hardware, and possibly focus more on tools that are sound only under certain assumptions. This would still be an improvement since most security engineers nowadays even use unsound tools if they are helpful. A laudable long-term goal is the verification of a browser.

A more technical topic was the way we think about the control flow of a program, in the sense that associating a program counter address with a control flow graph node is inadequate in the presence of self-modifying code. Similarly, it is not clear what constitutes a function (due to for example, tail sharing) and how to reliably identify a function in the presence of obfuscated or optimized code that does not adhere to any standard ABI. It was pointed out that functions can have hundreds of entries with a large common body, implying that duplicating this body for each entry might create a considerable code size increase for an analysis.

To contrast the applied side of binary analysis with a theoretical view on static analysis, we scheduled the more theoretic talks on Wednesday morning. The speakers addressed how mutating malware could be classified (Roberto Giacobazzi) and how to treat me-

mory allocated from a static array as independent heap cells (Xavier Rival). These topics gave an outlook on the challenges that lie beyond the already complicated reconstruction of the control flow graph.

Thursday and Friday featured talks mostly from the academic community who presented their current state-of-the-art. One particular debate arose on how the semantics of assembler instructions are best expressed. During an informal meeting on Thursday evening we agreed that the community would benefit from a common infrastructure to decode executable code. The way in which we should specify the semantics of native instructions was more difficult to agree upon. Thus, we set up a mailing list to discuss a common decoder infrastructure that should be able to accommodate several platforms (say ARM and x86). The design of a decoder should feature a domain specific language that allows for a human readable specification of decoding instructions. This DSL should ideally be usable to also express the semantics of instructions, even if the various groups might want to implement their own semantic interpretation depending on their analysis needs.

■ Participation

In all, 42 researchers, both senior and more junior, from 10 countries attended the meeting. This high number shows the strong interest in this emerging field. The feedback from the participants was also very positive.

■ Directions for the future

Thus, one of the tangible outcomes is that the community set out to create a common piece of infrastructure. Beyond this, it was agreed that another seminar about the analysis of executables in two years time would be most welcome. We discussed what topics this new seminar should focus on and we distilled that malware, obfuscation, interpreters and self-modifying code should be major topics, as these constitute challenges that the community needs to address.

4.6 Network Attack Detection and Defense Early Warning Systems – Challenges and Perspectives

Organizers: Georg Carle, Hervé Debar, Falko Dressler, and Hartmut König
Seminar No. 12061

Date: 05.–10. February, 2012 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.2.2.1

© © © Creative Commons BY-NC-ND 3.0 Unported license
 © Georg Carle, Hervé Debar, Falko Dressler, and Hartmut König



Participants: Nils Aschenbruck, Lothar Braun, Roland Büschkes, Georg Carle, Hervé Debar, Sven Dietrich, Till Dörges, Gabi Dreo Rodosek, Falko Dressler, Ulrich Flegel, Felix C. Freiling, Elmar Gerhards-Padilla, Peter Herrmann, Marko Jahnke, Holger Kinkel, Hartmut König, Jan Kohlrausch, Pavel Laskov, Corrado Leita, Michael Meier, Simin Nadjm-Tehrani, Andreas Paul, Aiko Pras, Konrad Rieck, René Rietz, Sebastian Schmerl, Bettina Schnor, Franka Schuster, Robin Sommer, Radu State, Jens Tölle, Michael Vogel, Alexander von Gernler, Stephen Wolthusen

The objective of the seminar was to discuss new challenges, technologies, and architectures in the area of network attack detection and defense. The focus of this seminar laid in particular on *early warning systems*, *malware detection*, and the *protection of critical infrastructures*, but also other recently emerging topics were supposed to be discussed. On this account, the seminar consisted of plenary sessions with technical talks and various breakout sessions. Beside the topics mentioned above two other topics on recently emerging issues were added, namely *cyber crime versus cyber war* and the *protection of cyber-physical systems*.

The seminar started off with an introductory session in which all participants shortly introduced themselves and discussed the focus and the structure of the seminar. Thereafter the first topic *Challenges on Early Warning Systems and Malware Detection* was raised. Michael Meier gave a state of the art talk on the development of early warning systems in the last years and open issues. Felix C. Freiling and Falko Dressler reported on the results of their projects in this field with the German Federal Office for Information Security (BSI). Jan Kohlrausch gave an overview of the experience with the deployment of early warning systems in practice with the DFN-CERT. In the afternoon the first breakout sessions were held. The topics discussed were the *Future of Early Warning Systems*, *Cloud Security*, and *Teaching IT Security*.

Tuesday was devoted to the topic *Protection of Critical Infrastructures*. Introductory talks of the various aspects and challenges for protecting critical infrastructures were given by Stephen Wolthusen and Corrado Leita, followed by technical talks by Franka Schuster and Andreas Paul about a project for protecting supervisory control and data acquisition (SCADA) networks, by Simin Nadjm-Tehrani on the security of smart meters, and by Georg Carle, Lothar Braun and Holger Kinkel on large-scale vulnerability assessment. In the afternoon Jens Tölle spoke about the protection of IP infrastructures with model-based cyber defense situational awareness. After coffee break we continued with two further breakout

sessions on *Information Security for Novel Devices* and *Fighting against Botnets*.

Wednesday morning was devoted to two special topics which have emerged recently: *Security of Cyber-Physical Systems* and *Cyber Crime versus Cyber War*. Nils Aschenbruck gave an introductory talk to the first topic reflecting the evolution from sensor networks to cyber-physical systems. Falko Dressler addressed in his talk the security challenges for future nano communication. The discussion on this topic was continued in the breakout session on Thursday. The second topic was opened by Felix C. Freiling posing various questions about the differences between malware for the masses and exclusive malware, and how to detect them as basis for a longer discussion in the auditorium. Gabi Dreo Rodosek then elucidated at length the issue in her talk about cyber defense. In the afternoon we made a nice trip to the historic city of Trier. The pretty cold weather there gave many opportunities to continue the discussions in warm coffee shops.

On Thursday morning we commenced with two talks by Pavel Laskov and Konrad Rieck on *Malware Detection* which dealt especially with machine learning aspects. Sven Dietrich added a talk on his SkyNET project about the use of drones to launch attacks on wireless networks. Thereafter we continued the topic on the protection of critical infrastructures with the focus on new challenges in deep packet inspection. Radu State began with a talk on the semantic exploration of DNS domains. René Rietz continued with a talk on the increasing threat by attacks over the web. After lunch Robin Sommer introduced the new version of the intrusion detection system (IDS) Bro. Alexander von Gernler reported about the current practice of application level firewalling and virus scanning from the perspective of a firewall manufacturer. Finally, Michael Vogel presented an approach for a dynamically adapting multi-agent intrusion detection system which copes with the growing gap between the evolution of network bandwidth and the single-thread performance of today's CPU architectures. After the coffee break, two further

breakout sessions on cyber-physical systems and smart energy grids took place.

Friday morning hosted two talks by Bettina Schnor and Simin Nadjm-Tehrani on IPv6 security and anomaly detection in mobile networks. After that we concluded the seminar with a discussion about the seminar outcome and possible future seminars.

■ Conclusion

The seminar was well-received by all participants. It gave a good opportunity to inform about current challenges in the area of network attack detection and defense and discuss possible countermeasures. Especially the breakout sessions found a great acceptance. The participants further liked much the possibility to have detailed discussions with colleagues outside the official program. They regret that not all invited foreign scientist accepted the invitation. They will advertise more strongly for this seminar. All participants agreed that proposal for another seminar should be submitted. There are two concrete contributions of this seminar:

1. Current research results of eight participating groups were published in special issue of the journal PIK 1/2012 which is especially devoted to this Dagstuhl seminar.
2. The discussion during the breakout session on cyber-physical systems showed that there is still an unclear picture on the security challenges to these systems. This raised the idea to apply for a Dagstuhl perspective workshop to discuss in detail the security challenges for protecting cyber-physical systems and to define them in a manifesto as working base for further research activities. The proposal has been submitted meanwhile.

4.7 Software Clone Management Towards Industrial Application

Organizers: Rainer Koschke, Ira D. Baxter, Michael Conradt, and James R. Cordy
Seminar No. 12071

Date: 12.–17. February, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.2.21

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Rainer Koschke



Participants: Hamid Abdul Basit, Ira D. Baxter, Saman Bazrafshan, Michel Chilowicz, Michael Conradt, James R. Cordy, Yingnong Dang, Serge Demeyer, Stephan Diehl, Daniel M. German, Michael W. Godfrey, Nils Göde, Jan Harder, Armijn Hemel, Elmar Jürgens, Cory J. Kasper, Jindae Kim, Rainer Koschke, Jens Krinke, Thierry Lavoie, Angela Lozano, Douglas Martin, Ravindra Naik, Jochen Quante, Martin P. Robillard, Sandro Schulze, Niko Schwarz, Werner Teppe, Rebecca Tiarks, Gunther Vogel, Andrew Walenstein, Minhaz Zibran

Software clones are identical or similar pieces of code or design. They are often a result of copying and pasting as an act of ad-hoc reuse by programmers. Software clone research is of high relevance for software engineering research and practice today. Several studies have shown that there is a high degree of redundancy in software both in industrial and open-source systems. This redundancy bears the risk of update anomalies and increased maintenance effort.

Many techniques exist that try to detect clones. Some of them are already available in open-source (e.g., PMD) as well as commercial tools (e.g., CloneDr). There are also lines of research in clone detection that evaluate these approaches, reason about ways to remove clones, assess the effect of clones on maintainability, track their evolution, and investigate root causes of clones. Today, research in software clones is an established field with more than 100 publications in various conferences and journals.

The purpose of this seminar was to solidify and give shape to this research area and community. Unlike previous similar events, this Dagstuhl seminar put a particular emphasis on industrial application of software clone management methods and tools and aimed at gathering concrete usage scenarios of clone management in industry, which will help to identify new industrially relevant aspects in order to shape the future research. Research in software clones is very close to industrial application. Among other things, we focused on issues of industrial adoption of our methods and tools.

To achieve our goals, we invited many participants from industry. We managed to reach a percentage of about 30 % industrial participation. Talks were given mostly by industrial participants who shared their experiences with us and posed their problem statements. Academic participants were allowed to give a talk if their talk had a clear focus on industrial experiences, needs, problems, and applications of software clone management and related research fields. The focus, however, was on interaction in form of plenary discussions and smaller working groups. The topics for workings

groups were gathered by clustering issues the participants wanted to discuss at the seminar. The seminar wiki was used intensively to record the results of the working groups. This agile format was very much appreciated by the participants.

The following working groups were formed:

- **Detection/Use cases:** This working group discussed issues in detecting clones. Because there are already many clone detectors, the focus of this working group was to gather use cases for these. The particularities of a use case dictates what kinds of features a suitable clone detector should have. The group's result was a list of different use cases for clone detection and an enumeration of distinct features a clone detector should have to support the respective use case. An overview of known limitations and issues of actual clone detectors is also provided along with some research questions oriented towards the improvement of clone detection techniques.
- **Presentation:** Because clone detectors typically find many clones in large systems, the user faces a huge amount of data he or she needs to make sense of. Visualization is a means of presenting large and complex data that takes advantage of a human's ability for visual pattern matching. This working group dealt with presentation issues of clone information. Again, use cases were enumerated because suitability of visualization is task dependent. The group connected the identified use cases with different existing types of software visualization suitable for these.
- **Interoperability:** To foster collaboration among researchers it is helpful to build interoperable tools. Then, for instance, the result of one researcher's clone detector could be fed into the visualization tool of another researcher. Interoperable tools are also needed to serve practitioners' diverse needs. This working group created a common model to represent clone information that addresses the needs of a wide range of use cases in research and practice.

- **Refactoring:** Contrary to the abundance of available clone detectors, there are relatively few tools that help in removing clones. The purpose of this working group was to consider the mechanics and utility of forming clone abstractions and achieving clone refactoring.

The group identified various means of eliminating clones that are either provided by the languages the clones are written in or by abstraction outside of the language (e.g., code generation). It also delved into managerial aspects of clone refactoring and particularities of clones in software product lines.

- **Clone management (process):** Clone management is the set of activities to detect, track, assess, handle, and avoid clones. This working group went into the matter of where clone management may play a role in the development and maintenance process.

The group discussed how clone analysis fits into the overall software development process (requirements engineering, development, testing, after deployment). They broached the issue of relation of code search and clone detection and how clone detection could be used in recommender systems.

- **Provenance and clones in artifacts that are not source code:** Most research in software clones focuses on source code, but as it has been shown by several researchers, clones can also be found in other software artifacts such as models and requirement specifications. This working group investigated needs to extend our research into these fields and the particularities of these fields with respect to clone detection. In addition to that, this working group dealt with provenance of clones, that is, the question where the clone comes from. Although the issues of provenance and clones in other artifacts appear to be largely independent, this working group worked on them jointly for organizational issues.

The group elaborated how clones could be detected and handled in binaries, models, and bug reports.

For the remainder of this report, it is important to know the following current categorization of clones:

- **Type-1 clone:** Identical fragments only.
- **Type-2 clone:** Lexically identical fragments except for variations in identifiers, literals, types, whitespace, layout, and comments
- **Type-3 clone:** Gapped clones, that is, clones where statements have been added, removed, or modified.
- **Type-4 clone:** Semantic clones, that is, clones with similar semantics but different implementations in code.

4.8 Information Visualization, Visual Data Mining and Machine Learning

Organizers: Daniel A. Keim, Fabrice Rossi, Thomas Seidl, Michel Verleysen, and Stefan Wrobel
Seminar No. 12081

Date: 19.–24. February, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.2.58

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Daniel A. Keim, Fabrice Rossi, Thomas Seidl, Michel Verleysen, and Stefan Wrobel



Participants: Daniel Archambault, Michael Aupetit, Michael Biehl, Kerstin Bunte, Etienne Come, Dianne Cook, Jean-Daniel Fekete, Brian D. Fisher, Ksenia Genova, Andrej Gisbrecht, Hans Hagen, Barbara Hammer, Helwig Hauser, Heike Hofmann, Ata Kaban, Samuel Kaski, Johannes Kehrer, Daniel A. Keim, Andreas Kerren, Jörn Kohlhammer, Bongshin Lee, John A. Lee, Marcus A. Magnor, Florian Mansmann, Ian Nabney, Jaakko Peltonen, Gabriele Peters, Nathalie Riche-Henry, Fabrice Rossi, Frank-Michael Schleif, Tobias Schreck, Marc Strickert, Holger Theisel, Peter Tino, Laurens van der Maaten, Jarke J. Van Wijk, Michel Verleysen, Nathalie Villa-Vialaneix, Thomas Villmann, Daniel Weiskopf, Hadley Wickham, Leishi Zhang

Information visualization and visual data mining leverage the human visual system to provide insight and understanding of unorganized data. Visualizing data in a way that is appropriate for the user's needs proves essential in a number of situations: getting insights about data before a further more quantitative analysis, presenting data to a user through well-chosen table, graph or other structured representations, relying on the cognitive skills of humans to show them extended information in a compact way, etc.

Machine learning enables computers to automatically discover complex patterns in data and, when examples of such patterns are available, to learn automatically from the examples how to recognize occurrences of those patterns in new data. Machine learning has proven itself quite successful in day to day tasks such as SPAM filtering and optical character recognition.

Both research fields share a focus on data and information, and it might seem at first that the main difference between the two fields is the predominance of visual representations of the data in information visualization compared to its relatively low presence in machine learning. However, it should be noted that visual representations are used in a quite systematic way in machine learning, for instance to summarize predictive performances, i.e., whether a given system is performing well in detecting some pattern. This can be traced back to a long tradition of statistical graphics for instance. Dimensionality reduction is also a major topic in machine learning: one aims here at describing as accurately as possible some data with a small number of variables rather than with their original possibly numerous variables. Principal component analysis is the simplest and most well known example of such a method. In the extreme case where one uses only two or three variables, dimensionality reduction is a form of information visualization as the new variables can be used to directly display the original data.

The main difference between both fields is the role of the user in the data exploration and modeling. The ultimate goal of machine learning is somehow to get rid of the user: everything should

be completely automated and done by a computer. While the user could still play a role by, e.g., choosing the data description or the type of algorithm to use, his/her influence should be limited to a strict minimum. In information visualization, a quite opposite point of view is put forward as visual representations are designed to be leveraged by a human to extract knowledge from the data. Patterns are discovered by the user, models are adjusted to the data under user steering, etc.

This major difference in philosophy probably explains why machine learning and information visualization communities have remained relatively disconnected. Both research fields are mature and well structured around major conferences and journals. There is also a strong tradition of Dagstuhl seminars about both topics. Yet, despite some well known success, collaboration has been scarce among researchers coming from the two fields. Some success stories are the use of state-of-the-art results from one field in the other. For instance, Kohonen's Self Organizing Map, a well known dimensionality reduction technique, has been successful partly because of its visualization capabilities which were inspired by information visualization results. In the opposite direction, information visualization techniques often use classical methods from machine learning, for instance, clustering or multidimensional scaling.

The seminar was organized in this context with the specific goal of bringing together researchers from both communities in order to tighten the loose links between them. To limit the risk of misunderstandings induced by the different backgrounds of researchers from the two communities, the seminar started with introductory talks about both domains. It was then mainly organized as a series of thematic talks with a significant portion of the time dedicated to questions and discussions. After the first two days of meeting, understanding between both communities reached a sufficient level to organize, in addition to the plenary talks, working group focusing on specific issues.

Several research topics emerged from the initial discussions and

lead to the creation of the working groups. The subject that raised probably the largest number of questions and discussions is Evaluation. It is not very surprising as differences between the communities about evaluation (or quality assessment) might be considered as the concrete technical manifestation of cultural and philosophical differences between them. Indeed, in machine learning, automatic methods are mostly designed according to the following general principle: Given a quality measure for a possible solution of the problem under study, one devises an algorithm that searches the solution space efficiently for the optimal solution with respect to this measure. For instance, in SPAM filtering a possible quality measure is the classification accuracy of the filter: it has to sort unsolicited bulk messages correctly into the SPAM class and all other emails in the HAM class. In a simple setting, the best filter could be considered as the one with the smallest number of errors. However, counting only the number of errors is usually too naive, and better quality measures have to be used, such as the area under the ROC curve: the Receiver Operating Characteristic curve shows the dependency between the true positive rate (the percentage of unsolicited bulk messages classified as SPAM) and the false positive rate (the percentage of correct emails classified as SPAM).

In information visualization, evaluation cannot rely only on mathematical quality measures as the user is always part of the story. A successful visualization is a solution, with which the user is able to perform better, in a general sense, compared to existing solutions. As in machine learning, a method is therefore evaluated according to some goal and with some quality metric, but the evaluation process and the quality metrics have to take the user into account. For instance, one display can be used to help the user assess the correlation between variables. Then, a quality metric might be the time needed to find a pair of highly correlated variables, or the time needed to decide that there is no such pair. Another metric might be the percentage of accurate decisions about the correlation of some pairs of variables. In general, a visualization system can be evaluated with respect to numerous tasks and according to various metrics. This should be done in a controlled environment and with different users, to limit the influence of interpersonal variations.

Among the discussions between members of the two communities about evaluation, questions were raised about the so-called unsupervised problems in machine learning. These problems, such as clustering or dimensionality reduction, are ill-posed in a machine learning sense: there is no unquestionable quality metric associated to e.g. clustering but rather a large number of such metrics. Some of those metrics lead to very difficult optimization problems (from a computational point of view) that are addressed via approximate heuristic solutions. In the end, machine learning has produced dozens of clustering methods and dimensionality reduction methods, and evaluations with respect to user needs remain an open problem. An important outcome of the seminar was to reposition this problem in the global picture of collaboration between information visualization and machine learning. For instance, if many quality measures are possible, one way to compare them would be to measure their link to user performances in different tasks. If several methods seem to perform equally well in a machine learning sense, then the user feedback could help to identify the “best” method. It was also noted that many methods that are studied in machine learning and linked to information visualization, in particular dimensionality reduction and embedding techniques, would benefit from more interaction between the communities. At minimum, state-of-the-art methods from machine learning should be known by information visualization researchers and state-of-the-art visualization techniques should be deployed by machine learning researchers.

Another topic discussed thoroughly at the seminar was the visualization of specific types of objects. Relational data were discussed, for instance, as a general model for heterogeneous complex da-

ta as stored in a relational database. Graph visualization techniques provide a possible starting point, but it is clear that for large databases, summarization is needed, which brought back the discussion of the ill defined clustering problem mentioned above. Among complex objects, models obtained by a machine learning algorithms were also considered, in particular as good candidates for interactive visualizations. Decision trees give a good example of such objects: Given a proper visualization of the current tree, of some possible simplified or more complex versions and of the effect of the tree(s) on some dataset, an expert user can adapt the tree to his/her specific goals that are not directly expressible in a quality criterion. The extreme case of visualizing the dynamic evolution of a self learning process was discussed as a prototype of complex objects representation: The system is evolving through time, it learns decision rules, and it evolves using complex (and evolving) decision tables.

Finally, it became clear that a large effort is still needed at the algorithmic and software levels. First, fast machine learning techniques are needed that can be embedded in interactive visualization systems. Second, there is the need for a standard software environment that can be used in both communities. The unavailability of such a system hurts research to some extent as some active system environments in one field do not include even basic facilities from the other. One typical example is the R statistical environment with which a large part of machine learning research is conducted and whose interactive visualization capabilities are limited, in particular in comparison to the state-of-the-art static visualization possibilities. One possible solution foreseen at the seminar was the development of some dynamic data sharing standard that can be implemented in several software environments, allowing fast communication between those environments and facilitating software reuse.

Judging by the liveliness of the discussions and the number of joint research projects proposed at the end of the seminar, this meeting between the machine learning and the information visualization communities was more than needed. The flexible format of the Dagstuhl seminars is perfectly adapted to this type of meeting and the only frustration perceivable at the end of the week was that it had indeed reached its end. It was clear that researchers from the two communities were starting to understand each other and were eager to share more thoughts and actually start working on joint projects. This calls for further seminars ...

4.9 Principles of Provenance

**Organizers: James Cheney, Anthony Finkelstein, Bertram Ludäscher, and Stijn Vansummeren
Seminar No. 12091**

Date: 26 February | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.2.84

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© James Cheney, Bertram Ludäscher, and Stijn Vansummeren



Participants: Umut A. Acar, Shawn Bowers, Peter Buneman, Adriane Chapman, James Cheney, Stephen Chong, Sarah Cohen-Boulakia, Victor Cuevas-Vicenttin, Lois Delcambre, Kai Eckert, Nate Foster, Juliana Freire, James Frew, Irini Fundulaki, Daniel Garijo, Floris Geerts, Ashish Gehani, Carole Goble, Todd J. Green, Paul Groth, Torsten Grust, Olaf Hartig, Melanie Herschel, Bertram Ludaescher, Andrew Martin, Simon Miles, Paolo Missier, Luc Moreau, Leon J. Osterweil, Christopher Re, Vladimiro Sassone, Martin Schäler, Margo Seltzer, Christian Skalka, Perdita Stevens, Wang-Chiew Tan, Jan Van den Bussche, Stijn Vansummeren, Marianne Winslett, Steve Zdancewic, Jun Zhao

The term “provenance” refers to information about the origin, context, derivation, ownership or history of some artifact. In both art and science, provenance information is crucial for establishing the value of a real-world artifact, guaranteeing for example that the artifact is an original work produced by an important artist, or that a stated scientific conclusion is reproducible. Even in everyday situations, we unconsciously use provenance to judge the quality of an artifact or process. For example, we often decide what food to buy based on freshness, origin and “organic” labels; and we decide whether or not to believe an online news article based on its source, author, and timeliness.

Maintaining good records of provenance that are sufficient to convince skeptics of the value of an artifact is difficult. It requires reflection or monitoring actions as they are performed. Every step in the chain of ownership of an important work of art needs to be recorded in a secure way, for example, in order to defend against forgery and deter attempts to sell stolen artwork.

Since it is much easier to copy or alter digital information than to alter real-world artifacts, there are even more opportunities for misinformation, forgery and error in the digital world than there are in the traditional physical world. For this reason, the need for provenance is now widely appreciated. Simple and unreliable forms of automatic provenance tracking, such as version numbering, ownership, creation and modification timestamps in file systems, have long been supported as a basic services on which more sophisticated tools can rely. In today’s increasingly networked and decentralized world, however, we anticipate the need for richer provenance recording and management capabilities to be built into a wide variety of systems.

For example, “grid” or “cloud” computing infrastructures are frequently used for scientific computing, as part of a widespread trend towards “eScience”, “cyberinfrastructure” or more recently the data-intensive “fourth paradigm” of science popularized by Jim Gray and others. These systems are complex and opaque. The cor-

rectness and repeatability of scientific conclusions (about, for example, climate change) is increasingly being questioned because of the lack of transparency of the complex computer systems used to derive the results. Provenance technology can help to restore transparency and increase the robustness of eScience, countering increasing skepticism of scientific results as evidenced by the so-called “Climategate” controversy in 2009.

This problem is already widely appreciated in scientific settings but is increasingly recognized as a problem in business, industrial and Web settings. Until recently, work on provenance has mostly taken place in relatively isolated parts of existing research communities, such as databases, scientific workflow-based distributed computing, or file systems, or the Semantic Web. However, we believe that to make real progress it will be necessary to form a broader research community focusing on provenance.

In this respect, the aims of Dagstuhl Seminar 12091 were to:

- bring together researchers from databases, security, scientific workflows, software engineering, programming languages, and other areas to identify the commonalities and differences of provenance in these areas;
- improve the mutual understanding of these communities;
- identify main areas for further foundational provenance research.

The seminar hosted 41 participants in total from the above communities, and included representatives from the W3C Provenance Working group that is in the process of standardizing a common data model for representing and exchanging provenance information.

To improve the mutual understanding of the various communities, the first day of the seminar was devoted to tutorial talks from well-respected members of each community.

The rest of the seminar consisted of presentations of recent ongoing provenance research in the various communities, as well as break-out sessions aimed at deepening discussions and identifying open problems.



Fig. 4.2
Marlene Reucher – Hommage an Nicola. Part of the Dagstuhl art collection and donated by: Alfons Simon, Roland and Ute Vollmar, Reinhard Wilhelm, and participants of the Dagstuhl Seminars 11492 and 12091.

4.10 Computation and Incentives in Social Choice

Organizers: Edith Elkind, Christian Klamler, Jeffrey S. Rosenschein, and M. Remzi Sanver
Seminar No. 12101

Date: 04.–09. March, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.3.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Edith Elkind, Christian Klamler, Jeffrey S. Rosenschein, and M. Remzi Sanver



Participants: Craig Boutilier, Steven J. Brams, Felix Brandt, Robert Brederbeck, Markus Brill, Ioannis Caragiannis, Jiehua Chen, Yann Chevaleyre, Vincent Conitzer, Andreas Darmann, Britta Dorn, Daniel Eckert, Edith Elkind, Ulle Endriss, Gábor Erdélyi, Piotr Faliszewski, Felix Fischer, Umberto Grandi, Paul Harrenstein, Edith Hemaspaandra, Lane Hemaspaandra, Marc Kilgour, Christian Klamler, Sascha Kurz, Jérôme Lang, Kate Larson, Omer Lev, Vangelis Markakis, Nicolas Maudet, Vincent Merlin, Ilan Nehama, Svetlana Obraztsova, Ulrich Pferschy, Francesca Rossi, Jörg Rothe, Lena Schend, Arkadii Slinko, Isabelle Stanton, Toby Walsh, Virginia Vassilevska Williams, Gerhard Woeginger, Michael Zuckerman

The aim of classic social choice theory is to explain how groups of agents can come to a joint decision that reflects the heterogeneous preferences of individual agents. This covers a wide range of scenarios, such as, for example, voting, fair division and ranking. As such, social choice theory enhances our understanding of human societies and can be used as a theoretical foundation for the design of multiagent systems.

In recent years, the study of computational aspects of social choice received a lot of attention from AI and theoretical computer science communities. This interest was motivated by existing and potential applications of social choice ideas in AI settings, which, in turn, highlighted the importance of understanding which of the recommendations of social choice theory are computationally feasible.

The value of algorithmic analysis in the context of social choice stems from the fact that, to be practically applicable, a decision-making rule needs to be efficiently implementable. Indeed, the analysis of computational complexity of well-known voting rules, both in the general case, and in interesting special cases (such as, e.g., single-peaked preferences) is one of the most actively studied topics in computational social choice, with a number of impressive results obtained so far.

However, computational tractability is not the only criterion for selecting a social choice procedure: an equally desirable feature is *incentive compatibility*, i.e., resilience to dishonest behavior by self-interested participants, who may want to manipulate the outcome of the procedure in their favor. There is an exciting interplay between incentive compatibility and computational tractability: in many settings of interest, computing one's optimal strategy requires solving a hard optimization problem, while acting honestly is computationally easy. Thus, one may view computational complexity as a barrier against strategic behavior, and try to design or identify social choice procedures that make strategizing difficult. This research di-

rection was initiated more than 20 years ago and remains a major research focus of the computational social choice community.

Alternatively, one can deal with manipulative agents in the context of social choice by embracing the strategic behavior rather than trying to prevent it. This can be done either by investigating the outcomes of standard social choice procedures under the assumption that all agents act strategically, or, more ambitiously, by designing social choice procedures that result in desirable outcomes even if agents are not truthful; these two approaches are associated, respectively, with game theory and mechanism design. Both game-theoretic and mechanism design approaches are widely used by the classic social choice community; however, their computational aspects have received relatively little attention so far.

In contrast, algorithmic aspects of strategic behavior in other settings, such as, e.g., matrix games or auctions, have been studied extensively in the last few years. Indeed, computational game theory and algorithmic mechanism design are among the fastest-growing subfields of both AI and theoretical computer science. Thus, in organizing this seminar, we aimed to bring together the researchers in the areas of computational and classic social choice and those in the area of algorithmic game theory. Our goal was to foster a discussion of computational aspects of various forms of strategic behavior in social choice contexts.

■ Outcomes

The seminar took place on March 4–9, 2012. It was interdisciplinary in nature: among the participants, there were computer scientists, mathematicians, social choice theorists and political scientists. There were 32 regular talks, as well as an after-dinner talk by Virginia Vassilevska-Williams, who spoke about her groundbreaking work on algorithms for matrix multiplication. The seminar talks covered a broad range of topics, such as, e.g., the complexity of dishonest behavior in voting, judgement aggregation, coalitional

game theory, and fair division. The program also featured a rump session consisting of short (5–8 minute) talks; these included announcements about events that were likely to be of interest to the seminar participants, short research talks, and presentations of open problems. The participants also used the seminar as an opportunity to continue ongoing research projects or start new ones. We are aware of two research papers that are largely based on discussions that happened during this Dagstuhl seminar; both of them have been recently submitted to the 4th International Workshop on Computational Social Choice. Moreover, several speakers who presented work in progress received useful feedback from other seminar participants, and, as a result, were able to improve or extend their papers significantly. To summarize, the participants of the seminar benefited from it in a variety of ways: by being exposed to new research results and directions, by getting fresh perspectives on their work, by learning about open problems and initiating new collaborations, and by having an opportunity to work with their co-authors from all over the world on ongoing research projects.

4.11 Normative Multi-Agent Systems

Organizers: Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leon van der Torre
Seminar No. 12111

Date: 11.–16. March, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.3.23

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Giulia Andrighetto, Guido Governatori, Pablo Noriega, and Leon van der Torre



Participants: Natasha Alechina, Giulia Andrighetto, Tina Balke, Jan M. Broersen, Cristiano Castelfranchi, Amit K. Chopra, Rob Christiaanse, Silvano Colombo-Tosatto, Stephen Cranefield, Natalia Criado, Célia da Costa Pereira, Mehdi Dastani, Marina De Vos, Gennaro Di Tosto, Frank Dignum, Yehia Elrakaiby, Nicoletta Fornara, Dov M. Gabbay, Aditya K. Ghose, Guido Governatori, Joris Hulstijn, Max Knobbout, Brian Logan, Henrique Lopes-Cardoso, Maite Lopez-Sanchez, Emiliano Lorini, Samhar Mahmoud, Eunat Mayor Villalba, John McBreen, Pablo Noriega, Mario Paolucci, Xavier Parent, Simon Parsons, David Pearce, Antonino Rotolo, Bastin Tony Roy Savarimuthu, Fernando Schapachnik, Francois Schwarzentruher, Munindar Singh, Paolo Turrini, Leon van der Torre, Wamberto Vasconcelos, Harko Verhagen, Serena Villata

The multi-disciplinary workshop on Normative Multi Agents attracted leading international scholars from different research fields (e.g., theoretical computer science, programming languages, cognitive sciences and social sciences).

The workshop was organised as follows: the organisers identified several relevant areas of research covering a wide and comprehensive spectrum of topics in the field of Normative Agents. For each area, a prominent researcher was appointed as chair for the area. In the months preceding the workshop the chairs collected material from the participants. During the first day they presented an overview of the areas they were in charge with special emphasis on some open questions and direction for future research.

The participants were divided in groups corresponding to the areas (due to some last minute cancellations some topics were under-represented and the scholars in those areas joined groups for closely related topics). Each group was allocated a morning session during which each member of the group had five minutes to provide an overview of their personal contribution to Normative Multi-Agents (plus some time for QA).

The format was well received by the participants and conducive to discussion. It gave them the opportunity to give very focused presentations while keeping the audience attention. The afternoon sessions, other the contrary, were dedicated to group work and group discussions. The aim of these sessions was to build consensus material of the specific topics and to identify fundamental research directions. The material is expected to be refined and to be articulated in chapters intended as a first step for the development for a road-map for this emerging area of computer-science with close interactions with other disciplines.

■ Results

During the seminar, participants split in seven working groups, centered around seven discussion themes. In the following paragraphs

there is a summary of the discussion held by each working group.

Normative MAS: An Introduction. This working group first focused on three definitions and some related requirements for normative MAS. For each of such definitions, some guidelines for developing normative MAS have been proposed. Second, it has been discussed how to relate the concept of normative MAS to different conceptions of norms and how norms can be used within the systems. Finally, some specific issues that open research questions or that exhibit interesting overlaps with other disciplines have been identified.

Normative Consequence. This working group first provided a definition what deontic logic and normative reasoning is. Second, it discussed why normative reasoning is relevant for normative multi-agent systems and pointed out the advantages of formal methods in multi-agent systems. Finally, it focused on the specificity of normative reasoning in comparison to other kinds of reasoning.

Computational NorMAS. This working group considered normative systems from the computational perspective, proposing the following themes as challenging for the domain: 1) trade-offs in expressive power of the languages for representing deontic notions (such as norms, conflicts of norms, violations of norms, etc.); 2) complexity of algorithms required for a) implementing tools capable of analysing and verifying norms, b) implementing normative system platforms capable of monitoring norm violations and finally c) implementing agents capable of deliberating about norms.

Regulated MAS: Social Perspective. This working group addressed the problem of building normative multi-agent systems. It developed a static conceptual model through which a normative multi-agent system may be specified along with a dynamic conceptual model through which the operation of a normative

system can be captured. A demonstration of how the proposed approach may be applied in prototypical applications of normative systems has been proposed.

Norm Compliance in MAS. This working group aimed to understand how norms regulate agent conduct and how norms impact on agent reasoning and behavior. Agents must be endowed with abilities to be able to reason about, process and otherwise manage norms in some appropriate fashion. In short, it demands that agent architectures are considered in terms of their ability to address these concerns, and that suitable architectures are developed.

(Social) Norm Dynamics. The working group aimed to identify the main steps in the dynamics of norms - i.e., generation, spreading, stabilization and evolution - as well as some of the relevant factors or determinants of such a process. The need for a deep understanding of these dynamics is becoming a compelling task for the NorMAS community due to the growing interest in open, evolving and flexible norm regulated and socio-technical systems. The working group pointed out that for a well-founded and innovative study of norms, it is necessary on the one hand to look at the cognitive mechanisms underlying the dynamics of norms and on the other hand to consider the role played by trust and cultural dimensions.

Norms and Simulation. This working group focused on the application of agent-based modeling and simulation to the issue of norm emergence, modification, and change. For the NorMAS community, agent-based simulations offer a platform to evaluate the behaviour of different models of norms and normative processes in a dynamic environment. Vice versa, the NorMAS community can supply (social) agent-based simulation studies with formal models of social concepts and mechanisms, especially those related to normative concepts, such as norms, roles, values, morals and conventions, and their transmission within a society.

The findings of the working groups were reported and discussed during the morning plenary sessions, and led to lively debate. During the seminar, each working group drafted a document presenting the main outputs achieved. Further work within the groups (by email correspondence) followed the end of the seminar, allowing finalizing the documents.

After a review process, the contributions of the working groups will be collected in a volume of the novel Dagstuhl Follow-up Series titled *A Prospective view of Normative-Multi Agent Systems*, aimed to become a standard reference in the field and to provide guidelines for future research in normative multi-agent systems.

In addition, *The Journal of Logic and Computation* and *Artificial Intelligence and Law* have agreed to publish special issues based on expanded and revised versions of the material presented at the seminar.

4.12 Applications of Combinatorial Topology to Computer Science

Organizers: Lisbeth Fajstrup, Dmitry Feichtner-Kozlov, and Maurice Herlihy

Seminar No. 12121

Date: 18.–23. March, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.3.50

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Lisbeth Fajstrup, Dmitry Feichtner-Kozlov, and Maurice Herlihy



Participants: Henry Adams, Sergio Cabello, Armando Castaneda, Bernadette Charron-Bost, Frederic Chazal, Frederick R. Cohen, Armindo Emanuel Costa, Carole Delporte, Jean-Marie Droz, Herbert Edelsbrunner, Ulrich Fahrenberg, Lisbeth Fajstrup, Michael Farber, Hugues Fauconnier, Eva-Maria Feichtner, Dmitry Feichtner-Kozlov, Pierre Fraigniaud, Patrizio Frosini, Eric Goubault, Tobias Heindel, Maurice Herlihy, Anil N. Hirani, Marc Jeanmougin, Mateusz Juda, Matthew Kahle, Petr Kuznetsov, Claudia Landi, Frank H. Lutz, Facundo Memoli, Roy Meshulam, Alessia Milani, Samuel Mimram, Dmitriy Morozov, Thomas Nowak, Ami Paz, Sergio Rajsbaum, Martin Raussen, Srivatsan Ravi, Matthieu Roy, Primož Skraba, Christine Tasson, Corentin Travers, Mikael Vejdemo-Johansson, Peter Widmayer, Krzysztof Ziemianski

In recent years, concepts and techniques adapted from combinatorial and algebraic topology have led to a variety of promising new results in several areas of Computer Science, including distributed computing, sensor networks, semantics of concurrency, robotics, and vision.

The recent Dagstuhl seminar *Applications of Combinatorial Topology to Computer Science* (12121), brought together researchers in these fields, both to share ideas and experiences, and to establish the basis for a common research community. Because of differences in terminology and academic culture, it is often difficult for researchers in one area to become aware of work in other areas that may rely on similar mathematical techniques, sometimes resulting in duplication of effort. This Dagstuhl seminar provided a valuable opportunity to bring together researchers in both computer science and mathematics who share a common interest in emerging applications of combinatorial topology.



Fig. 4.3

Franz Jenull – DA/III. Part of the Dagstuhl art collection and donated by Britta Schinzel.

4.13 Open Models as a Foundation of Future Enterprise Systems

Organizers: Robert B. France, Ulrich Frank, Andreas Oberweis, Matti Rossi, and Stefan Strecker
Seminar No. 12131

Date: 25.–30. March, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.3.67

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Ulrich Frank, Andreas Oberweis, Matti Rossi, Robert B. France, and Stefan Strecker



Participants: Jörg Becker, Tony Clark, Patrick Delfmann, Jörg Desel, Gregor Engels, Werner Esswein, Hans-Georg Fill, Robert B. France, Ulrich Frank, Andreas Hess, Dimitris Karagiannis, Mogens Kühn-Pedersen, Marc Lankhorst, Sina Lehmann, Peter Loos, Florian Matthes, Andreas Oberweis, Erik Proper, Mirja Pulkkinen, Wolfgang Reisig, Dirk Riehle, Matti Rossi, Stefan Strecker, Bernhard Thalheim, Michael zur Muehlen

To effectively support business operations and managerial decision-making, future enterprise systems require an elaborate conceptual foundation that promotes a tight mutual alignment of information systems and the business. Enterprise models provide such a foundation. They integrate conceptual models of an information system (e. g. an object model) with models of the surrounding action system (e. g. business process models or strategy models). Thereby, they relax the notorious cultural chasm between business and IT experts and provide a versatile instrument for the conjoint development of large-scale, mission-critical enterprise systems and for analyzing and (re-) designing the corporation.

However, the development of comprehensive enterprise models requires efforts, expertise, and resources beyond the capabilities of even large corporations. Therefore, the development and dissemination of reference enterprise models that can be adapted to a wide range of companies is a pivotal success factor. Enterprise models are usually specified by domain-specific modelling languages (DSML). The development and evaluation of reference enterprise models and corresponding DSML is an attractive scientific challenge. It corresponds to the development of theories: Reference models and DSML are linguistic constructions (on different levels of abstraction) that come with the claim for general validity or suitability respectively—not just for one particular occurrence but for an entire class of organizations.

They integrate and consolidate contributions from several scientific disciplines such as Computer Science, Information Systems, and Management Science. Both, reference models and DSML provide a reification of an attractive vision: Higher quality of software systems at lower cost. It is the complexity of modern organizations and the diversity of involved perspectives that renders the development of reference enterprise models and corresponding DSML a particular research challenge. Inspired by the remarkable results of the free/open source movement, recent work on reference enterprise models has resulted in the notion of open reference enterprise mo-

dels (open models for short). Research into open models does not only address the feasibility issue. Furthermore, it introduces a new model of collaboration among researchers, developers, and prospective users of reference enterprise models—leading to the prospect of shaping future enterprise systems. Recent initiatives on joint, collaborative modeling of open licensed conceptual models, thus, provide a new, innovative model for research on reference enterprise models that served as the starting point to this Dagstuhl seminar. It links to research on collaborative modeling, modeling tool development, model management, models@run.time, enterprise systems, and model-driven engineering.

This Dagstuhl seminar was aimed at bringing together a multi-disciplinary group of academic and industry researchers from the disciplines of Wirtschaftsinformatik, Computer Science, Information Systems, and Software Engineering, specifically those working in Requirements Analysis, Conceptual Modelling, and Enterprise Modelling to foster our understanding of how to develop, evaluate, disseminate, and promote the use of open reference enterprise models. The primary emphasis of the seminar was to determine the present state-of-the-art in this multi-disciplinary research field, and to establish a research agenda for future work towards solving theoretical and practical challenges related to the development of open reference enterprise models. The following overview describes more particular questions/objectives and related achievements:

1. **What are key characteristics of future ES?** The analysis of this question started with assumptions about relevant changes to be expected for the use of future ES. On the one hand, it was commonly expected that in many industries there will be a growing need for adapting the ES quickly to changing demands, e. g. to benefit from sudden opportunities or to build effective protection against threats. On the other hand, it was assumed that a growing number of managers will have received professional training in sophisticated uses of information systems. As a consequence, it was concluded that future ES should not only

be based on an elaborate conceptual foundation but should also make this foundation, e. g. an enterprise model, accessible to prospective users—on various levels of abstraction and detail. This would not only empower users to perform more advanced analyses, but also to modify the ES to a certain extent by applying changes to certain parts of the underlying conceptual model. From a software engineering perspective such a conception of future ES creates the challenge to allow for using models at run time—and to synchronize models and code. It was concluded that programming languages which allow for an arbitrary number of abstraction layers provide a promising approach to address this challenge.

2. **What is a promising strategy for the development of a common modeling platform?** A platform for enterprise modeling needs to integrate an extensible set of DSML editors. Also, it should support the specification of DSML and the development of corresponding model editors. Furthermore, it should enable model analysis and support the use of models at run time. The participants agreed that there is no environment available that would satisfy all these demands. At the same time, developing such an environment would require a substantial amount of resources and would take years. During that time, the intended modeling activities would be compromised, since they lacked the required tool platform. Therefore, it was concluded that only an evolutionary approach to developing a common modeling platform is a realistic option. It should start with existing modeling tools that are gradually extended or replaced with more advanced systems.
3. **What are key features to be offered by a repository to integrate contributions from a wide range of participants?** Since a common modeling environment cannot be expected at the beginning of an open model initiative, there is need to integrate contributions (models, meta models etc.) from various sources. That puts emphasis on a versatile repository that allows handling a wide range of representations on a level of semantics that enables model integration and various forms of retrieval and analysis. A working group focused on a corresponding architecture and presented an elaborate proposal.
4. **What are appropriate guidelines to establishing and sustaining initiatives and corresponding processes of collaborative modeling of open models?** Apart from incentives, discussions centered on organisational issues involving considerations of the economics of open models and success factors related to community aspects, procedural aspects, stakeholder aspects and infrastructure aspects. A life-cycle and a maturity model were proposed together with an initial process model aimed at guiding the steps to establish and sustain open model initiatives. The concluding plenary discussions corroborated the need for a guided and concerted division of labor.

A joint publication by the organizers is currently in preparation to reflect the seminar's key results. It is to appear in 2013.

4.14 Touching the 3rd Dimension

Organizers: Daniel Keefe, Antonio Krüger, Frank Steinicke, and Jean-Baptiste de la Rivière
Seminar No. 12151

Date: 09.–12. April, 2011 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.4.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Daniel Keefe, Antonio Krüger, Frank Steinicke, and Jean-Baptiste de la Rivière



Participants: Elisabeth André, Gerd Bruder, Andreas Butz, Sheelagh Carpendale, Raimund Dachsel, Florian Daiber, Jean-Baptiste de la Rivière, Tanja Döring, Steven K. Feiner, Mark Hancock, Petra Isenberg, Tobias Isenberg, Joaquim A. Jorge, Daniel Keefe, Sebastian Knödel, Sven Kratz, Antonio Krüger, Alexander Kulik, Marc Erich Latoschik, Karljohan Lundin Palmerius, Sina Mostafawy, Patrick L. Olivier, Antti Oulasvirta, Bernhard Preim, Bruno Rodrigues De Araujo, Johannes Schöning, David Selosse, Anthony Steed, Frank Steinicke, Sophie Stellmach, Wolfgang Stuerzlinger, Dimitar Valkov

■ Touching the 3rd Dimension

In recent years interactive visualization of 3D data has become important and widespread due to the requirements of several application areas. However, current user interfaces often lack adequate support for 3D interactions: 2D desktop systems are often limited in cases where natural interaction with 3D content is required, and 3D user interfaces consisting of stereoscopic projections and tracked input devices are rarely adopted by ordinary users.

Touch interaction has received considerable attention for 2D interfaces, and more recently for 3D interfaces. Many touch devices now support multiple degrees of freedom input by capturing multiple 2D contact positions on the surface as well as varying levels of pressure and even depth. There is great potential for multi-touch interfaces to provide the traditionally difficult to achieve combination of natural 3D interaction without any instrumentation. When combined with a stereoscopic display as well as depth cameras, multi-touch technology have the potential to form the basis for a next generation of 3D user interfaces.

Our Dagstuhl seminar *Touching the 3rd Dimension* focussed on bringing together researchers from a diverse set of fields of Computer Science to discuss the next generation of user interfaces based on multi-touch technology as well as 3D visualization. It is envisaged that such user interfaces will be fundamentally different to the current multi-touch systems that are being deployed in mobile devices, desktop environments as well as entertainment systems. First, future stereoscopic displays will work without the need for users to wear cumbersome 3D stereo glasses. In this context autostereoscopic displays have already been deployed for personal working spaces as well as entertainment systems, but are still rarely used due to technical limitations such as low resolution and lacking multi-viewpoint support. Second, future multi-touch technology will not be limited by planar screens, but touch surfaces may be of arbitrary shapes. Furthermore, touch interactions in space will be possible,

while still providing the user haptic feedback, for example, by the usage of air pressure devices.

■ Multi-touch Interaction

Tangible user interfaces (TUIs) augment digital information and computation by physical form and thus help to activate the sophisticated skills, which people have developed for sensing and manipulating their physical environments, for the exploration of virtual worlds. In this context multi-touch technology is one of the most interesting current developments in human-computer interaction which allows users to interact via a touch screen without the need of conventional input devices (e. g., mouse, keyboard). In contrast to standard touch screens that recognize only a single touch point, multi-touch screens recognize multiple simultaneous touch points. Recently presented user interfaces such as Apple's iPhone or iPad, Microsoft's Surface or Kinect, or Jeff Han's work on multi-touch sensing based on frustrated total internal reflection (FTIR) prove the relevance of the emerging field of multi-touch enabled surfaces. Currently, most of people's sensing and manipulation skills are not used when interacting with digital worlds. Multi-touch builds upon those skills and situates the physically-embodied digital information on the multi-touch enabled screen, thus helping the user to interact in a natural way while supporting tactile feedback.

Several research groups have analyzed how humans interact with multi-touch surfaces, and new interaction techniques, in particular for 2D interactions (e. g., panning, rotating, zooming) have been proposed. Although some researchers have implicitly addressed the interaction with 3D data, 3D interaction, in particular with stereoscopic displayed data, has not been examined in depth until now.

■ 3D Visualization

Most 3D user interfaces are applied only in highly specific application scenarios within some virtual reality (VR) laboratories. To some extent this is certainly due to the fact that interaction of humans with synthetic 3D environments still suffers from many drawbacks, and numerous problems have not finally been resolved. Devices with three or more degrees of freedom may provide a more direct interface to 3D manipulations than their 2D counterparts, but using multiple DoFs simultaneously still involves challenges. Since most often 2D interactions are performed best with the WIMP metaphor and 2D devices supporting only two DoFs (e. g., mouse and keyboard), 3D user interfaces are usually inappropriate to accomplish tasks requiring exclusively or mainly two-dimensional control. However, two-dimensional desktop systems are often limited in cases where natural interfaces are desired, which enable users to interact like in the real world. In such cases, VR systems using stereoscopic projections of three-dimensional synthetic worlds and tracking technologies support a better exploration of complex data sets. When stereoscopic display is used, each eye of the user perceives a different perspective of the same scene. With using stereoscopic displays, objects can be displayed with different stereoscopic *parallax paradigms*, i. e., negative, zero, and positive parallax, resulting in different stereoscopic effects. Objects with identical and congruent projections in both half-images have *zero parallax* and will appear to be at the same depth as the screen, objects with *negative parallax* appear in front and objects with *positive parallax* behind the projection screen.

However, although binocular disparity provides important depth cues, interaction with stereoscopically displayed objects is still a challenging task, in particular when the interaction is restricted to a 2D touch surface. Objects having zero parallax are displayed monoscopically and therefore are ideally suited for multi-touch interaction. Objects with positive or negative parallax appear behind respectively in front the touch screen and therefore cannot be accessed directly; since interaction is restricted to the screen plane either the screen limits the reach of the user or the hands may interfere with the stereoscopic effect.

However, several research groups have begun to explore the potential, limitations, and challenges of this and other 3D touch environments, and first commercial systems are already available. The goal of this seminar is to address the research and industrial challenges involved in exploring the space where the flat digital world of surface computing meets the physical, spatially complex, 3D space in which we live. The seminar provided a common forum to foster discussion among scientists, designers, manufactures and other participants working in several areas of HCI addressing the associated research questions, such as design and manufacturing of stereoscopic touch surfaces, 3D interaction techniques, evaluation methodologies, social impact, collaborative scenarios, and emerging application areas, to share their visions of the future and recent results in the area of improving 3D interaction and visualization by taking advantage of the strengths of advanced multi-touch computing.

4.15 Software Synthesis

Organizers: Rastislav Bodik, Sumit Gulwani, and Eran Yahav
Seminar No. 12152

Date: 09.– 13.04, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.4.21

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Rastislav Bodik, Sumit Gulwani, and Eran Yahav



Participants: Rajeev Alur, Don Batory, Rastislav Bodik, Krzysztof Czarnecki, Jyotirmoy Deshmukh, Laurent Doyen, Bernd Finkbeiner, Pierre Flener, Franz Franchetti, Sumit Gulwani, Amey Karkare, Hadas Kress-Gazit, Rupak Majumdar, Oded Maler, Mark Marron, Vu Minh Le, Alon Mishne, Jean-Noël Monette, Georg Ofenbeck, Doron A. Peled, Ruzica Piskac, Nir Piterman, Corneliu Popeea, Subhajit Roy, Roopsha Samanta, Sven Schewe, Sanjit A. Seshia, Armando Solar-Lezama, Saurabh Srivastava, Philippe Suter, Nathalie Sznajder, Emina Torlak, Stavros Tripakis, Martin T. Vechev, Christian von Essen, Eran Yahav, Jean Yang

Software verification and synthesis are founded on similar principles, yet verification has become industrial reality while successes of synthesis remain confined to a handful of domains. Still, recent years witnessed increased interest in software synthesis—a trend spurred by growing software complexity and simultaneously enabled by advances in verification, decision procedures, and machine learning. The goal of the seminar is to help the revival of software synthesis through intellectual exchange among experts in deductive synthesis, controller synthesis and the diverse spectrum of new synthesis efforts in inductive synthesis, auto-tuning, programming by demonstration and partial programming.

This is an opportune moment for software synthesis. First, multi-core processors are likely to make software development harder, motivating automatic construction of synchronization and communication code. Second, software verification and checking reached industrial maturity through judicious use of linguistic support, decision procedures, and dynamic analyses, inspiring solutions to open synthesis problems. Third, by incorporating verification into synthesis, we may be able to synthesize programs that are easier to verify than handwritten programs. Fourth, parallel computers enable search powerful enough for synthesis of well-tuned programs, as demonstrated by auto-tuners and super-optimizers. Finally, recent systems built on programming by demonstration make us hope that specification will be easier to write.

The seminar organizers hope to achieve the following goals:

- Offer brief tutorials on techniques developed by communities participating in the seminar.
- Develop a set of challenge problems for practical synthesis, a collection of practical problems solvable by (semi-)automatic synthesis in five years.
- Deepen the understanding of the relationships between the various approaches to synthesis. In particular, to what extent are the techniques developed by the respective communities inde-

pendent from their driving applications? Understand strengths of the alternative approaches.

- Understand relationships and applicability of verification technology to software synthesis.
- Outline a syllabus for a graduate course in software synthesis.

4.16 Abstractions for scalable multi-core computing

Organizers: Faith Ellen, Christof Fetzer, Tim Harris, and Nir Shavit
Seminar No. 12161

Date: 15.–20. April, 2012 | Dagstuhl Seminar

Participants: Yehuda Afek, Dan Alistarh, Hagit Attiya, David F. Bacon, Annette Bieniusa, Irina Calciu, Brian Demsky, Dave Dice, Sandhya Dwarkadas, Faith Ellen, Pascal Felber, Christof Fetzer, Tim Harris, Danny Hendler, Lisa Higham, Matt Horsnell, Michael Isard, Ryan Johnson, Milind Kulkarni, Konrad Lai, Christian Lengauer, Yossi Lev, Jean-Pierre Lozi, Mikel Lujan, Maged M. Michael, Alessia Milani, Mark Moir, Adam Morrison, J. Eliot B. Moss, Rotem Oshman, Victor Pankratius, Dmitri Perelman, Erez Petrank, Michael Philippsen, Martin Pohlack, Ravi Rajwar, Vijaya Ramachandran, Thomas Rauber, Srivatsan Ravi, Torvald Riegel, Noam Rinetzky, Gudula Rünger, Michael L. Scott, Nir Shavit, Gurindar S. Sohi, Michael F. Spear, Kevin Streit, Philippos Tsigas, Osman Unsal, Martin T. Vechev, Jons-Tobias Wamhoff



(No documentation available)

4.17 Semantic Data Management

Organizers: Grigoris Antoniou, Oscar Corcho, Karl Aberer, Elena Simperl, and Rudi Studer
Seminar No. 12171

Date: 22.–27. April, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.4.39

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Grigoris Antoniou, Oscar Corcho, Karl Aberer, Elena Simperl, and Rudi Studer



Participants: Karl Aberer, Grigoris Antoniou, Marcelo Arenas, Wolf-Tilo Balke, James Cheney, Oscar Corcho, Philippe Cudré-Mauroux, Gianluca Demartini, Orri Erling, Dieter Fensel, Norbert Fuhr, Avigdor Gal, José Manuel Gomez-Perez, Alasdair J G Gray, Marko Grobelnik, Paul Groth, Andrey Gubichev, Peter Haase, Stephen Harris, Olaf Hartig, Manfred Hauswirth, Jeff Heflin, Spyros Kotoulas, Paolo Missier, Luc Moreau, Charalampos Nikolaou, Ivana Podnar Zarko, Edna Ruckhaus, Satya S. Sahoo, Manuel Salvadores, Ralf Schenkel, Juan F. Sequeda, Wolf Siberski, Elena Simperl, Kavitha Srinivas, Rudi Studer, Kerry Taylor, Martin Theobald, Bryan Thompson, Duc Thanh Tran, Frank van Harmelen, Maria-Esther Vidal, Valentin Zacharias

The Semantic Web represents the next generation World Wide Web, where information is published and interlinked in order to facilitate the exploitation of its structure and semantics (meaning) for both humans and machines. To foster the realization of the Semantic Web, the World Wide Web Consortium (W3C) developed a set of metadata (RDF), ontology languages (RDF Schema and OWL variants), and query languages (e.g., SPARQL). Research in the past years has been mostly concerned with the definition and implementation of these languages, the development of accompanying ontology technologies, and applications in various domains. This work has been very successful, and semantic web technologies are being increasingly adopted by mainstream corporations and governments (for example by the UK and USA governments) and by several science communities (for example, Life Sciences or Astronomy). Moreover, semantic technologies are at the core of future developments, e.g. in the UK Open Data Institute. However, compared to more traditional solutions, semantic technologies often appear to be immature, and current tools lag behind in terms of efficiently handling of large data sets. What are additionally needed include solid data management concepts, architectures, and tools that follow the paradigms of more traditional database (DB) and information retrieval (IR) systems. Semantic data management refers to a range of techniques for the manipulation and usage of data based on its meaning. The aim of this workshop was to discuss in-depth a number of crucial issues, with particular emphasis on the fruitful exchange of ideas between the semantic web, database systems and information retrieval communities. Relevant key questions cutting across all topics covered were: (i) how can existing DB and IR solutions be adapted to manage semantic data; and (ii) are there new challenges that arise for the DB and IR communities (i.e. are radically new techniques required)? For the purposes of this workshop, and for this report, we understand semantic data simply as data expressed in RDF, the lingua franca of linked open data and hence the

default data model for annotating data on the Web. The workshop was organized along the following key themes:

1. **Scalability.** In order to make semantic technologies take on the targeted market share, it is indispensable that technological progress allows semantic repositories to scale to the large amount of semantic data that is already available and keeps growing. It is essential to come close to performance parity with some of the best DB solutions without having to omit the advantages of a higher schema flexibility compared to the relational model. Moreover, the exploitation of semantic data on the Web requires managing the scale that so far can only be handled by the major search engine providers. However, this should be possible without losing the advantages of a higher query expressivity compared to basic key-value stores and IR solutions.
2. **Provenance.** An important aspect when integrating data from a large number of heterogeneous sources under diverse ownership is the provenance of data or parts thereof; provenance denotes the origin of data and can also include information on processing or reasoning operations carried out on the data. In addition, provenance allows for effectively supporting trust mechanisms and policies for privacy and rights management.
3. **Dynamicity.** Another important property of many (semantic) data is its dynamicity. While some data, such as public administration archives or collections of text documents might not change too frequently, other data, coming from sensors, RSS, user-generated content (e.g. microblogging), etc., might evolve on a per millisecond basis. The effects of such changes have to be addressed through a combination of stream processing, mining, and semantics-based techniques.
4. **Search and Ranking.** The large and growing amount of semantic data enables new kinds of applications. At the same time, more data means that ultimately, there might be more results produced from it that one can or desires to inspect. Data and results to concrete information needs vary in the degree of re-

levance. Effective ranking mechanisms that incorporate the information needs as well as contextual information into account can deliver and rank pertinent results and help the users to focus on the part of the data that is relevant.

Each day of the workshop was scheduled to deal with one of the four key themes, in the order that has been presented above. For every theme there was an initial tutorial presentation that aimed at setting the context and vocabulary to be used throughout the day, as well as expose the current state of the art in the theme and the open challenges. Such tutorial presentation also helped to introduce the more in-depth research topics dealt with later on by the other presenters. Some of the main results of the discussions held during the workshop were the following, separated again by themes:

■ Scalability (day 1)

Issues that have been dealt with so far in the semantic data management community are how to store semantic data, how to index it and how to deal with processing and optimizing complex queries. Solutions that have been proposed work for the small datasets and scale that we know from the previous years. Meanwhile, the amount of data published on the Web including linked data and RDFa data associated with Web pages has exploded, resulting in billions of RDF triples – a number that is increasing very fast. Existing techniques and solutions no longer scale to this scenario. It requires the adoption of mainstream and proven DB and IR technologies and possibly, new radical solutions. Some of the main questions addressed at the workshop were (a) how DB solutions capable of dealing with large amounts of data with flexible schema or even schema-less data can be adopted for semantic data management. Likewise, IR solutions to data management have proven to be robust and scale to the Web. (b) How can this body of work be adopted and extended to deal with not only keyword search but also complex queries? In particular, the specific topics that require interests are (c) storage and indexing of large-scale semantic data, (d) parallel SPARQL query processing and optimization, (e) federated query processing on linked data, (f) top-k and approximate query processing, (g) workload-adaptive dynamic load balancing and optimization and (h) semantic data management in the cloud.

The theme was introduced by Philippe Cudre-Mauroux (University of Fribourg) through a tutorial that reviewed some of the recent techniques to support vertical (scale-up) and horizontal (scale-out) scalability for Semantic Web systems. Several presentations followed that were focused not only on demonstrating how some scalability aspects had been addressed in different systems, but also at clarifying what scalability means for semantic data management, and how it relates to the DB community. These were presentations from Avigdor Gal (Technion), Frank van Harmelen (VU), Peter Haase (fluidOps), Juan Sequeda (UT) and Charalampos Nikolau (NKUA). These were complemented with presentations from use cases in the financial and medical domains, from Steve Harris (Garlik) and Satya Sahoo (Case Western Reserve University).

The rest of presentations and discussions of the day, including part of the research groups, were focused on benchmarking. Two specific presentations on benchmarking were delivered by Kavitha Srinivas (IBM) on the characterisation of benchmarks, and Jeff Heflin (Lehigh University), on the experience with LUBM. These were the basis of the working group that was held later on benchmarking, and were important for several discussions throughout the week.

Some of the conclusions obtained from the discussions were related to the relationships between relational models and graph data models, and how difficult it was to express this difference clearly across communities. In particular, it was rather clear that the semantic web community is still too constrained in the work related

to trying to make the underlying technology store and access RDF, rather than exposing the benefits of semantics. In fact, work and discussions should not be focused on comparing both worlds, but on characterising the workloads and usage that drive the use of each type of technology, and the tipping point of when to use each. A message of obtaining scalability not only in seconds, but also in months (in terms of the benefits that semantic technology can bring in in issues like schema heterogeneity and data integration) were also pointed out.

■ Provenance (day 2)

Provenance is an important theme in the database systems and Web communities. There are two main research trends related to provenance management: workflow and data provenance. In the case of data provenance, the underlying model deals with declarative queries that include operations such as joins and unions, whereas workflow provenance models typically describe procedural workflows and involve operations that are treated as black boxes because of their complexity. The W3C Linking Open Data effort has boosted the publication and interlinkage of large amounts of heterogeneous RDF datasets on the Semantic Web and the result of this work is the Linked Data paradigm. In this context, the big challenges that must be addressed are related to the evaluation of linked data properties such as quality, trustworthiness, reliability to support a variety of applications including data integration and fusion, and mashups. In order to be able to evaluate the aforementioned qualities it is critical to record the provenance which denotes the origin of data that can also include information on processing or reasoning operations carried out on the data.

Paul Groth (VU) was in charge of the initial tutorial where the basic concepts of provenance were explained, and was accompanied next by Olaf Hartig (HU Berlin), Luc Moreau (University of Southampton) and Paolo Missier (Newcastle University) on the introduction to the W3C PROV family of specifications. Such tutorial and introduction to W3C activities were a good starting point to understand the current status of work on provenance by the semantic data management community. Besides, results from the recent Dagstuhl seminar on Provenance were also shared with the attendees.

The first presentation after the tutorial was focused on how provenance is used in BioPortal to provide access control in SPARQL. This was presented by Manuel Salvadores (Stanford University). Then the rest of presentations were focused on technologies and approaches to represent and exploit provenance in different contexts. These were presentations by Bryan Thompson (SYSTAP), which proposed a novel manner to represent provenance for triples or quads, and James Cheney (University of Edinburgh), who presented database wiki. There was also one presentation on workflow provenance done by Kerry Taylor (CSIRO).

Some of the conclusions obtained from the provenance sessions were related to the fact that this is a clear community with clear problems, and already quite well organised (for instance, through the corresponding W3C group on provenance). Some of the current limitations, and calls for action were on the need to make provenance data sets available and understandable, and address some of the interoperability issues that currently exist between different languages and models for describing provenance. Some of the open research issues that need to be addressed are (a) the definition of provenance models that take into account schema information to represent provenance information independently of the underlying applications (b) the extension of the existing query languages to support implicit provenance (d) querying and reasoning with provenance data (e) efficient storage of provenance data in the presence of data evolution (f) the use of provenance in applications such as view maintenance and updates, access control, debugging, attribution of responsibility.

ty, assessment of trustworthiness and information quality (g) presentation and visualization of provenance information, (h) provenance for probabilistic RDF data, and (i) provenance in scientific workflows in order to get insights on the quality and integrity of results, and to better understand scientific results, while also reaching the more general business workflow community. The possibility of organising a panel at one of the subsequent BPM conferences was considered as a nice follow-up. Another nice analysis is provided by Paul Groth in his blog, at <http://thinklinks.wordpress.com/2012/05/01/dagstuhl-semantic-data-management/>

■ Dynamicity (day 3)

Several areas of work address the problems related to the dynamicity of data sources, many of them common across types of data sources and leading to general research issues such as: (a) generating models for describing and understanding the data sources and the evolution of their data in time, considering not only primary data sources, but also derived ones, whose changes normally depend on the ones that they are based on. There are also specific problems related to the intrinsic characteristics of each data source. For example, problems in sensor and stream data management are mainly related to: (b) the efficient processing and storage of this type of data, (c) to the standardization of their representation using semantic models, (d) to the integration of heterogeneous sources, and to their analysis in order to detect patterns, trends, complex events, but also anomalies, by means of different types of techniques, including learning and summarization techniques.

Kerry Taylor (CSIRO) was in charge of presenting the initial tutorial explaining what we normally understand by dynamicity and how it is tackled on the Web (mainly Data Stream Management Systems and Complex Event Processing) and on the Semantic Web (with approaches focused on ontology versioning as streams of Blow"changes, and streaming extensions of SPARQL that have appeared lately). She also described current efforts in providing semantics to sensor networks and their related data, such as the W3C Semantic Sensor Network (SSN) Ontology, and finished her keynote with a discussion on some of the upcoming challenges in this area: pushing semantics into little devices, catching up with the work on semantics that is being performed by the Open Geospatial Consortium (OGC) and working on effecting change, not just responding to it.

Then we had several sessions on presentations about dynamicity in the context of data streams and sensor data, from Spyros Kotoulas (IBM Research), Manfred Hauswirth (DERI), Oscar Corcho (UPM) and Ivana Podnar Zarko (University of Zagreb), on the need to consider the dynamic nature of endpoints and Linked Data URIs while querying, from María Esther Vidal (Universidad Simón Bolívar) and Olaf Hartig (HU Berlin), and on dynamicity in scientific data, services and workflows, from José Manuel Gómez-Pérez (iSOCO).

In the context of semantic sensor data and data stream processing, several use cases were proposed, many of which where in the context of Smart Cities (traffic, emergency management, black water distribution, energy prediction, etc.) and in environmental sensing. Some of the open challenges in this area are related to how to make sure that most of the processing and RDF management could be done at the sensor level, instead of the middleware level, given the limitations of these devices. how to handle private and public data, static and dynamic data, provenance, access rights, etc. The need for appropriate benchmarks and experimental facilities was also highlighted, and the need to standardise the semantic stream query language to be used and the Linked stream standards.

In the context of dynamicity of endpoints and Linked Data, the need for adaptive semantic data management techniques was high-

lighted, and several current proposals in this direction were presented. Besides, a novel approach for querying the Web of Linked Data, focused on dereferencing URIs instead of using directly endpoints (e.g., as in the SQUIN system), was presented and heavily discussed among attendees.

Finally, in the context of dynamicity in scientific domains (workflows, data sources, services), the need to identify and handle appropriate evolution and lifecycle models for all these artifacts, so as to ensure reproducibility in Science, was presented and discussed, together with lessons learned from some ongoing work on scientific workflow preservation, such as the need to monitor decay, to understand data with a purpose, and to encapsulate parts of workflows and services.

Some of the main conclusions of the discussions held during this day were related to **the need to cover all types of streaming data sources** (from sensor data to social media) as new types of data sources that may benefit from the use of semantics in different forms. In the context of sensor data sources, work is clearly underway by different groups and through joint initiatives at W3C, and some of the next challenges are related to the needs of the Internet of Things community, in scenarios like, for instance, Smart Cities. Another conclusion was that the Web is dynamic, and we have to be able to **provide solutions that are able to adapt to this dynamicity**, including query support when no triple stores are available in Web servers, we have to provide a better characterisation of the balance between caching and non-caching, and intermediate solutions, and we have to cope with run-time discovery of semantic data sources, query planning to them, etc., as well as data and processes decay.

■ Search and Ranking (day 4)

In the IR community, the central question has always been the one about relevance. Which results are relevant to the query posed by the user? Recently, this question has attracted interests of DB researchers. Instead of retrieving documents, which constitute the primary subject in the IR field, query processing in the DB community is about computing complex results tuples. Correspondingly, different notions of relevance have been investigated. There exists a broad range of applications for semantic data which might involve simple IR-like queries as well as complex tasks in the fashion of DB-style relational queries, online analytical processing and reasoning. Also from the data perspective, semantic data shares both DB- and IR-like characteristics. Semantic data found on the Web might be clean and highly-structured just like relational data in the database or messy, schema-less and possibly, contain a large-amount of text just like data in an IR system. For ranking results and enabling users to focus on relevant semantic data, we need to study, adopt and extend existing notions of relevance proposed by IR and DB researchers. Some of the specific questions discussed during the workshop were (a) how can IR-approaches such as the probabilistic model or the most recent state of the art of generative models be applied to the problem of semantic data ranking? Is this enough or (b) do we also need to take the explicit structure and semantics of the data into account, and (c) can we leverage some existing DB-approaches for that? Do we need hybrid models combining IR and DB solutions to deal with the structured and unstructured nature of semantic data? (d) What other new challenges are introduced by semantic data and correspondingly, what new concepts are required to address them?

Duc Thanh Tran (KIT) was in charge of the initial tutorial explaining what we understand by semantic search (using semantics in structured data, conceptual and lexical models) and why it is important to consider it. He explained several of the techniques currently used for semantic search in the state of the art and discussed about

several research directions, among which he pointed out to keyword search over structured/semantic data, supporting complex information needs (long tail), and requiring efficient traversal algorithms and specialized indexes, and various models for ranking). Among the selected challenges identified in this talk we can cite the need for benchmarking, the need to deal with hybrid content management, and ranking hybrid results, the need to study the querying paradigm to use for complex retrieval tasks (keywords, natural language, facets), and the possibility of using semantics or providing a broader search context.

This presentation was followed by several other presentations on the use of probabilistic models to check consistency between links of different datasets, from Edna Ruckhaus (Universidad Simón Bolívar), on ranking SPARQL results, from Ralf Schenkel (Universität de Saarlandes), and an example application of search mechanisms applied to a concrete case study on crystals, from Kerry Taylor (CSIRO). Later on some additional presentations were done on some of the current limitations of SPARQL in terms of complexity, from Marcelo Arenas (PUC).

The last two presentations, and part of the discussion, went about the role of crowdsourcing for search and for the provisioning of metadata, with talks from Gianluca Demartini (University of Fribourg) and Wolf-Tilo Balke (TU Braunschweig).

One of the main conclusions from this day and the discussions held during the day were related to the fact that this theme is the one that shows a stronger relationship between the three communities at which this workshop was addressing (IR/DB/SW), and hence it is one of the areas where there is a larger possibility of more **cross-fertilisation** and of finding more research and further development opportunities. Some additional examples to semantic search will be gathering evidences from documents, combining data sources while answering queries, applying provenance in IE pipelines, etc.

Another important conclusion and step forward that was identified was related to the context of **crowdsourcing and semantics**, which is well extended in the IR community. It was clear throughout discussions that we cannot talk yet about a Semantic Read/Write Web in a general manner, since the SW community is still operating in a Read/Web1.0 context (a few publishers only) instead of in a Read/Write Web2.0 model (lots of producers). Some opportunities may arise in this context in the application of this paradigm for complex tasks like link validation, improving the quality of Linked Data, etc. What it was not so clear was whether it was easy as well to provide support from semantic technologies to some of the existing crowdsourcing activities that are being held now outside of the SW community. In this sense, there was a proposal to organise in the future a workshop or roadmapping activity on this topic.

4.18 Quality of Experience: From User Perception to Instrumental Metrics

Organizers: Markus Fiedler, Sebastian Möller, and Peter Reichl
Seminar No. 12181

Date: 01.–04.May, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.5.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Markus Fiedler, Sebastian Möller, and Peter Reichl



Participants: Patrik Arlos, Ake Arvidsson, Sergio Beker, Katrien R. De Moor, Marcus Eckert, Sebastian Egger, Markus Fiedler, Marie-Neige Garcia, Riccardo Guerzoni, Dennis Guse, Richard John Harris, Helmut Hlavacs, Tobias Hoßfeld, Selim Ickin, Lucjan Janowski, Maria Kihl, Khalil Laghari, Patrick Le Callet, Sebastian Möller, Alexander Raake, Peter Reichl, Gerardo Rubino, Junaid Shaikh, Lea Skorin-Kapov, Martin Varela, Katarzyna Wac, Ina Wechsung

During the recent years, Quality of Experience (QoE) has established itself as a topic of its own for both industrial and academic research. With its focus on the user in terms of acceptability, delight and performance, it is about to take over the role of Quality of Service as key paradigm for provisioning and managing services and networks. As one of the follow-up activities of the Dagstuhl Seminar 09192 “From Quality of Service to Quality of Experience”, this Dagstuhl Seminar 12181 focused on the relation between quality perception and QoE quantification, which is among the most challenging tasks for bringing together the three essential corner stones, i.e. user, technology, and business. In particular, qualitative user perception needs to be translated into quantitative input to dimensioning and control of networks and services. Further, different kinds of feedback flows (acceptance, usage, cost, quality) need to be taken into account. Considering the multidisciplinary nature of this problem with complementary and potentially controversial views, the seminar worked towards metrics and measurement techniques aimed at improving QoE prediction and control. The outcomes are expected to become visible in the future QoE research agenda and corresponding standardisation efforts.

■ Introduction

Dagstuhl seminars strongly depend on the delegates and their input. In order to give room for both presentations and group discussions during a three-day seminar, the presentations were confined to five minutes and one slide. Each presentation was followed by a short block of questions and answers. In order to truly reflect the delegates’ positions with regards to the topic of the seminar, the abstracts have been included in the sequel as-is and in alphabetical order.

The presentation round was followed by the presentation of a QoE White Paper around a QoE definition that has emanated from the Dagstuhl Seminar 09192.

The related discussions of QoE-related definitions and notions were continued and deepened during the first group work entitled “Key aspects of experience perception and their subjective evaluation”. The other two groups discussed “Measurable aspects of QoE” and “Identification of QoE-related feedback loops”. The outcomes of the group works were presented and discussed in the plenum, and excerpts are presented below.

The seminar was concluded with a plenary discussion of follow-up activities.

4.19 Social, Supply-Chain, Administrative, Business, Commerce, Political Networks: a Multi-Discipline Perspective

Organizers: Matthias Häsel, Thorsten Quandt, and Gottfried Vossen
Seminar No. 12182

Date: 02.–04. May, 2012 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.2.5.26

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Matthias Häsel, Thorsten Quandt, and Gottfried Vossen

Participants: Jörg Becker, Daniel Beverungen, François Bry, Clemens Cap, Ingo Dahm, Stuart Dillon, Emese Domahidi, Matthias Häsel, Bernd Hellingrath, James A. Hendler, Stefan Klein, Nicolas Pflanzl, Thorsten Quandt, Michael Räckers, Gottfried Vossen



The information society is shaped by an increasing presence of networks in various manifestations. Efficient computer networks are regarded as a significant enabler for the process of change towards networks of any size and complexity. They serve as an administrative and technological basis for social network structures, with the result that online networks connect people all around the world at day and night, and allow to communicate and to work collaboratively, efficiently, and without recognizable time delay. Companies reduce their in-house production depth, join forces in supply chain networks and establish cooperation with their suppliers, with their customers, and even with their competitors. By now, social networks like Facebook, Google+, LinkedIn or XING are seen as the de facto standard of “social networking” in the information society. Companies are mimicking their effects internally, allow overlays of networking applications with regular business ones, and a use of social networks for enterprise purposes including and beyond advertising has become common. Public administrations create and improve shared services and establish “Private Public Partnerships (PPP)” to benefit from synergetic effects of cooperation with private and public organizations.

As the interactions between people in these networks increase at various levels, new approaches are needed to analyze and study networks and their effects in such a way that individuals as well as organizations and enterprises can benefit from them. This Perspectives Workshops has convincingly shown that more interaction and collaboration between fields such as information systems, computer science, social sciences, economics, communication sciences and others is needed. The fields need to identify a common level of language, tools and set of methodologies so that the various aspects of networking can be addressed and jointly developed further. The most important point is the need for a renewed multi-disciplinarity. To a great extent, networks are driven and further developed by practitioners; which also means that they are evolving in a very fast manner and not emanating from a single scientific discipline. To

be able to both understand them and contribute to the state of art, true inter- or multi-disciplinary research is needed that involves the fields mentioned. As these distinct disciplines grow together and embark on collaborative research, it is also important to convince funding agencies that multi-disciplinary research should arrive on their agendas. Finally, Web sciences need to be developed as a field, and also need to be integrated into teaching. This will most likely lead to novel curricula which receive their content from multiple disciplines in a balanced way.

4.20 Artificial and Computational Intelligence in Games

Organizers: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius
Seminar No. 12191

Date: 6.–11. May, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.5.43

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius



Participants: Elisabeth André, Ruth Aylett, Christian Bauckhage, Michal Bida, Adi Botea, Bruno Bouzy, Paolo Burelli, Michael Buro, Martin V. Butz, Alex J. Champandard, Clare Bates Congdon, Peter Cowling, Marc Ebner, Mirjam P. Eladhari, Richard Evans, Philip F. Hingston, Graham Kendall, Pier Luca Lanzi, John M. Levine, Daniele Loiaco, Simon M. Lucas, Michael Mateas, Risto Miikkulainen, Hector Munoz-Avila, Dana S. Nau, Ana Paiva, Mike Preuss, Günter Rudolph, Tom Schaul, Moshe Sipper, Pieter Spronck, Kenneth O. Stanley, Tommy Thompson, Julian Togelius, Georgios N. Yannakakis, R. Michael Young

The video game industry is the largest of the entertainment industries and growing rapidly. The foundations of this industry are techniques from computer science. New developments within video games pose fresh challenges to computer scientists. Around the world, the number of dedicated study programs producing the workforce of the game industry is increasing steadily, as is the number of computer science academics dedicating their careers to solving problems and developing algorithms related to video games. Such problems often require domain knowledge from various research domains, such as psychology and the arts, leading to an inherently interdisciplinary research field.

Artificial intelligence (AI) and computational intelligence (CI), in one form or another, can be found at the heart of almost any video game, controlling the non-player characters (NPCs) as well as many aspects of the game world. They are also used throughout the game design and development process. Academic research within these domains in games aims to solve problems and enable innovation, pertaining to game design, game development, and gameplay. A main focus is on solving algorithmic problems to make game mechanisms more intelligent and efficient, thus making games more immersive, interesting, and entertaining. In the context of serious and educational games, such improvements enable these games to fulfill their societal objectives better.

Artificial intelligence seeks to simulate intelligent behavior in any possible way with human intelligence as a paradigm. Computational intelligence is an umbrella term for nature-inspired computational methods for optimization, learning and controlling. The main methods are evolutionary algorithms, artificial neural networks, fuzzy logic, swarm intelligence, and artificial immune systems. Nowadays, the borders between both disciplines are blurred, and state-of-the-art solutions use hybrid techniques combining elements of symbolical AI systems, CI algorithms and methods from statistical machine learning.

The aim for the *Dagstuhl Seminar on Artificial and Computa-*

tional Intelligence in Games was to bring together creative experts in an intensive meeting with the common goals of gaining a deeper understanding of various aspects of games, and of further improving games. It was meant to enforce the communication of different communities and the collaboration with the games industry. The exchange of different views and competencies was to help identify the main challenges in game AI research and the most promising venues to deal with them. This could lead to a common vision of what kind of games could be made possible in the future.

The Seminar was held from Sunday, May 6, 2012, until Friday, May 11, 2012. Over 40 researchers came together at Schloß Dagstuhl, many of them highly-respected and well-known researchers in their field, but also several talented young researchers and even a few representatives from the AI specialists of the game industry. In contrast to what is common for such gatherings, very little time was spent on plenary talks. Instead, the focus was on workgroups which discussed particular topics. However, several plenary sessions were held in which the workgroups reported on their results, and new topics for discussion were brought up. To allow researchers to present their recent work, a poster session was held during the second day of the Seminar, and the posters remained up until the end.

The topics of the workgroups, in alphabetical order, were the following:

- AI Architectures for Commercial Games
- AI Clearing House
- AI for Modern Board Games
- Computational Narratives
- Evaluating Game Research
- Game AI for Mobile Devices
- General Game Playing
- Learning in Games
- Pathfinding
- Player Modeling
- Procedural Content Generation

- Search
- Social Simulation Games
- Video Game Description Languages

Several researchers wrote a short report on their poster, which are included too. We aim to bring a full reports of all the workgroups in the form of proceedings later.

As organizers we are really pleased with how the Seminar turned out. It proved to be the stimulating and inspirational environment that we had hoped for. We found that most, if not all participants agreed with us on that. A lot of this success is due to the excellent facilities provided by the people of Schloss Dagstuhl. We are highly grateful for having had the opportunity to be their guests for the Seminar. We definitely hope to return in the future.

4.21 Co-Design of Systems and Applications for Exascale

Organizers: Arndt Bode, Adolfy Hoisie, Dieter Kranzlmüller, and Wolfgang E. Nagel
Seminar No. 12212

Date: 22.–25. May, 2012 | Dagstuhl Perspectives Workshop
 Full report – DOI: 10.4230/DagRep.2.5.71

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Arndt Bode, Adolfy Hoisie, Dieter Kranzlmüller, and Wolfgang E. Nagel



Participants: Arndt Bode, Hans-Joachim Bungartz, Kent Czechowski, Karl Furlinger, Hans Michael Gerndt, Vladimir Getov, Adolfy Hoisie, Jeffrey K. Hollingsworth, Darren Kerbyson, Alice E. Koniges, Bettina Krammer, Dieter Kranzlmüller, Erwin Laure, Allen Malony, John Mellor-Crummey, Barton P. Miller, Bernd Mohr, Daniel Molka, Wolfgang E. Nagel, Cherri M. Pancake, Michael M. Resch, Sabine Roller, Martin Schulz, Stephen L. Scott, Christian Straube, Achim Streit, Shinji Sumimoto, Jan Treibig, Felix Wolf, Zhiwei Xu

With Petascale computing being a reality today, the focus of the computational science community is already on the next barrier – exascale computing. With systems even more powerful by orders of magnitude, scientists start thinking about the possibilities and challenges. This workshop addressed the many scientific, technological, and financial challenges of exascale level computing with the hypothesis that exascale computing is only possible by co-designing across different levels of software, hardware, and the surrounding infrastructure.

The workshop program has been composed of a series of short talks, less than 20 minutes on average, and extensive time for discussions. Starting with an overview of the workshop motivation and the general methodologies for co-design, different aspects of co-design have been addressed. This has been followed by talks on modeling, simulation and tools, as well as programming models, runtime support and compilers. The second part addressed the specific problems of system-software for performance, power and reliability and the resulting system architectures, while finally application level aspects of exascale co-design have been discussed between the participating experts from different areas of high performance computing. In all discussions it has been important to tackle a multidimensional combination of major challenges associated with the development of exascale systems and applications from different angles instead of addressing an isolated aspect.

The results of the workshop are manifold¹: Firstly, the vision based on the requirements of the scientific community is thus “to provide exascale capabilities to scientific and engineering applications”, where it is important to notice that exascale means extreme scale or large scale, not the particular barrier of exaflop performance looming ahead. With this vision at hand, the participating experts identified their particular role and mission as follows: “to co-design systems such that they reach exascale capabilities within the given technological and non-technical (social, ...) boundaries”. Each ex-

pert has been knowledgeable on a distinct layer of the exascale architecture, the mission requires expertise across all layers, and exascale computing requires involvement from all relevant areas of computer science in order to perform exascale co-design of hardware and software, including also different levels of software working closely together with hardware and the interfacing to the environmental infrastructure. This has led to the definition of co-design, where two or more distinct activities collaborate on and across different layers to design a system architecture for a specific goal.

In summary, the workshop has reflected on the current state of petascale machines providing multiple examples from world-leading machines and using them to derive the barriers on the road towards exascale computing. Looking beyond the current research into the future, where exascale computing will become feasible, we have been trying to identify the exascale roadmap with intermediate goals and pitfalls on the way to exascale, and leveraging the combined forces of computer science to overcome them.

¹ A scientific paper will be created within the next months.



Fig. 4.4

Rolf Zimmermann – Apfel, Birne. Part of the Dagstuhl art collection and donated by Barbara und Michael Marhöfer.

4.22 Cognitive Approaches for the Semantic Web

Organizers: Dedre Gentner, Frank van Harmelen, Pascal Hitzler, Krzysztof Janowicz, and Kai-Uwe Kühnberger

Seminar No. 12221

Date: 28. May–01. June, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.5.93

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Dedre Gentner, Pascal Hitzler, Kai-Uwe Kühnberger, Frank van Harmelen, and Krzysztof Janowicz



Participants: Benjamin Adams, Sören Auer, Alan Bundy, Claudia d'Amato, Jérôme Euzenat, Kenneth D. Forbus, Andrew U. Frank, Christian Freksa, Aldo Gangemi, Dedre Gentner, Robert L. Goldstone, Giancarlo Guizzardi, Helmar Gust, Cory Henson, Pascal Hitzler, Zhisheng Huang, Frank Jäkel, Krzysztof Janowicz, Kai-Uwe Kühnberger, Werner Kuhn, David M. Mark, Alexander Mehler, Jens Ortmann, Sebastian Rudolph, Simon Scheider, Christoph Schlieder, Ute Schmid, Lael Schooler, Willem van Hage, Frank van Harmelen, Cong Wang, Stephan Winter, Wei Lee Woon, Gudrun Ziegler

The Dagstuhl Seminar 12221 on *Cognitive Approaches for the Semantic Web* was held from May 28th to June 1st, co-organized by Dedre Gentner, Frank van Harmelen, Pascal Hitzler, Krzysztof Janowicz and Kai-Uwe Kühnberger. The motivation of this seminar was to gather people from Semantic Web and Cognitive Science in order to determine the most promising ways to move forward on the vision of bringing findings from cognitive science to the Semantic Web, and to create synergies between the different areas of research. The seminar mainly focused on the use of cognitive engineering methods towards a more user-centric Semantic Web. However, the reverse direction, i.e., how Semantic Web research on knowledge representation and reasoning can feed back to the cognitive science community, was also discussed. Besides core members of the Semantic Web, artificial intelligence, and cognitive science communities, the researchers from fields that would benefit most from a more human-centric Semantic Web were also present. This especially included experts on Geographic Information Science (GIScience), the bioinformatics, as well as the digital humanities. While the invitations were balanced, most attending participants were from the Semantic Web, cognitive science, and GIScience communities.

The seminar consisted of three alternating blocks, short talks by the participants, work in breakout groups, and reports by the breakout groups followed by discussions among all participants. The short talks presented the participants' research or future ideas and were the inspiration for the topics discussed in the breakout groups. Each day had a distinct subtopic with respect to the combination of presenters and the formed breakout groups. While the task of the breakout groups differed, it was ensured that each of the 5-7 groups consists of members of all major research domains present at the meeting.

On May 29th, the first day of the seminar, Krzysztof Janowicz gave a short opening talk about the structure of the seminar. Next, Frank van Harmelen gave an overview talk about the Semantic Web, while Dedre Gentner introduced the cognitive science perspective

focusing on work on analogies. After lunch, the participants, Rob Goldstone, Christian Freksa, Ken Forbus, Kai-Uwe Kühnberger, Alexander Mehler, Ute Schmid, Gudrun Ziegler, and Helmar Gust, all involved in cognitive science research, presented their work in short talks of 10 minutes. After these talks, breakout groups were formed. The task of each group was to develop a research proposal outline and present it to all participants.

On May 30th, the participants presented their results from the breakout groups. This second day was devoted to researchers from GIScience, bioinformatics, and the digital humanities, as well as work of researchers that already bridged between the Semantic Web and cognitive science. The presenters were Andrew Frank, Werner Kuhn, Aldo Gangemi, Cory Henson, David Mark, Krzysztof Janowicz, Giancarlo Guizzardi and Simon Scheider. In the afternoon, the participants formed new breakout groups based on the presented topic. The task was to develop user interfaces and user interaction paradigms that exploit Semantic Web reasoning on the one side and analogy and similarity-based reasoning on the other side. Finally, the groups reported back to all participants and discuss synergies.

May 31st, started with additional domain talks and was then followed by presentations of core Semantic Web researchers. Presentations were given by Sören Auer, Lael Schooler, Willem van Hage, Zhisheng Huang, Stephan Winter, Christoph Schlieder, Jens Ortmann, Ken Forbus, Alan Bundy, Benjamin Adams, Jérôme Euzenat, Claudia d'Amato, Sebastian Rudolph, Wei Lee Woon and Pascal Hitzler. In the afternoon, the breakout groups were formed to discuss how Cognitive Science can benefit from Semantic Web research. The task was to design an experiment (in most cases involving human participants). Afterwards the breakout group reported back to all participants.

June 1st, last day of the seminar, started with two longer talks (each about 30 min.) that reported back on what Semantic Web researchers learned from cognitive scientists during the meeting as well

as the other way around. The first presenter was Jérôme Euzenat representing his view as Semantic Web researcher on the lessons learned. The second presentation was given by Rob Goldstone to illustrate the lessons learned by the cognitive science community. Finally, the seminar concluded with general discussions on future research and feedback about the seminar.

4.23 Future Internet for eHealth

Organizers: Katarzyna Wac, David Hausheer, Markus Fiedler, and Paolo Bonato
Seminar No. 12231

Date: 03.–06. June, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.6.1

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Katarzyna Wac, David Hausheer, Markus Fiedler, and Paolo Bonato



Participants: Albert Alonso, Gerald Bieber, Doris M. Bohman, Paolo Bonato, Sara Eriksén, Markus Fiedler, Geraldine Fitzpatrick, Caroline Franck, Stefan Göbel, Nick Guldemond, Mattia Gustarini, Jody Hausmann, Rainer Herzog, Hannes Kaufmann, Zviad Kirtava, Willy Kostucki, Erno Kovacs, Lenka Lhotská, Maria Martini, Goran Martinovic, Dave Marvit, Kevin Patrick, Terje Peetso, Pawel Swiatek, Halina Tarasiuk, Vicente Traver Salcedo, Muhammad Ullah, Marc van Anderlecht, Katarzyna Wac

The paradox of life in the 21st century is that while advancements in technology and medicine enable us to live longer, our lifestyle choices increase the probability of becoming chronically ill earlier in our life and experience long-term limitations, requiring long-term social support. In 2005, 78% of European medical care spending was on chronic disease management, while 86% of deaths were due to such a disease. Yet, current health systems are designed for an acute cure rather than for a chronic care, leading to a continuous increase in healthcare costs. To achieve economically sustainable and affordable healthcare system, efficient and effective solutions are needed integrating technological advancements, and empowering the patients for better self-management, as well as healthcare teams for better decisions.

Recently, multiple initiatives have been established to shape the Internet of the future, supporting key application sectors such as healthcare, transportation, and energy, amongst others. At the same time, the emergence of next generation high bandwidth public wireless networks and miniaturized personal mobile devices have given rise to new mobile healthcare (mHealth) services. For example, highly customizable vital sign tele-monitoring of chronically ill patients can be provided based on body area networks (BAN) and mHealth applications. Such applications enable live-transmission of the data to healthcare providers, and real-time feedback to the patient, enabling her to self-manage her disease and health, respectively. Additionally, elderly people can benefit from applications that help them to stay in contact with their care teams, which are provided with valuable hints on the state of the elderly, thus in the long run facilitating economically sustainable care combined with an improved quality of life.

However, such applications do not emerge by themselves, but need to be carefully designed to support in an evolutionary way the existing healthcare workflows, fulfilling their duties at the given quality level and cost. Such a task can only be tackled in a multi-disciplinary way as it was a goal of this seminar; experts from heal-

thcare, elderly care, insurance experts, together with experts from domains such as human-computer interaction, interactive application design, telecommunications, networking and economy teamed up to understand and support each other in designing and deploying future-proof eHealth services and applications based on Future Internet technology.

At large, the seminar addressed the following questions:

1. Which will be the key eHealth applications and services in the Future Internet?
2. Which are current and future quality requirements of eHealth applications and services?
3. Which business models are viable for future eHealth applications?
4. Which methodological support is required to design economically sustainable network-supported eHealth services?

Question 1 teamed up the participants around relevant use cases and facilitated discussions on the technical question 2 and the economical question 3, respectively. Question 4 addressed research needs from different domains and fertilized corresponding activities for advancing the topic of Future Internet for eHealth.



Fig. 4.5
Juliana Hümpfner – ohne Titel. Part of the Dagstuhl art collection and donated by Roland and Ute Vollmar, Reinhard Wilhelm, Angelika Mueller-von Brochowski, participants of the Lehrerfortbildung Informatik 2009 (09053), and further people.

4.24 Data Reduction and Problem Kernels

Organizers: Michael R. Fellows, Jiong Guo, Dániel Marx, and Saket Saurabh
Seminar No. 12241

Date: 10.–15. June, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.6.26

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Michael R. Fellows, Jiong Guo, Dániel Marx, and Saket Saurabh



Participants: Faisal Abu-Khazam, Cristina Bazgan, Hans L. Bodlaender, Bruno Courcelle, Marek Cygan, Rodney Downey, Andrew Drucker, Michael R. Fellows, Henning Fernau, Rudolf Fleischer, Fedor V. Fomin, Robert Ganian, Serge Gaspers, Archontia Giannopoulou, Jiong Guo, Gregory Z. Gutin, MohammadTaghi Hajiaghayi, Pinar Heggernes, Danny Hermelin, Petr Hlineny, Juraj Hromkovic, Falk Hüffner, Bart Jansen, Mark Jones, Iyad A. Kanj, Eun Jung Kim, Christian Komusiewicz, Stefan Kratsch, Michael A. Langston, Daniel Lokshtanov, Dániel Marx, Kurt Mehlhorn, Daniel Meister, Neeldhara Misra, Matthias Mnich, Anil Nerode, Rolf Niedermeier, Christophe Paul, Geevarghese Philip, Marcin Pilipczuk, Michal Pilipczuk, Venkatesh Raman, Fahimeh Ramezani, Felix Reidl, Frances A. Rosamond, Peter Rossmanith, Noy Rotbart, Ignasi Sau Valls, Saket Saurabh, Ildiko Schlotter, Hadas Shachnai, Somnath Sikdar, Karolina Soltys, Ulrike Stege, Ondrej Suchy, Stefan Szeider, Jan Arne Telle, Dimitrios M. Thilikos, Yngve Villanger, Magnus Wahlström, Gerhard J. Woeginger, Anders Yeo

Preprocessing (data reduction or kernelization) is used universally in almost every practical computer implementation that aims to deal with an NP-hard problem. The history of preprocessing, such as applying reduction rules to simplify truth functions, can be traced back to the origins of Computer Science — the 1950’s work of Quine, and much more. A modern example showing the striking power of efficient preprocessing is the commercial integer linear program solver CPLEX. The goal of a preprocessing subroutine is to solve efficiently the “easy parts” of a problem instance and reduce it (shrinking it) to its computationally difficult “core” structure (the *problem kernel* of the instance).

How can we measure the efficiency of such a kernelization subroutine? For a long time, the mathematical analysis of polynomial time preprocessing algorithms was neglected. The basic reason for this anomalous development of theoretical computer science, was that if we seek to start with an instance I of an NP-hard problem and try to find an efficient (P-time) subroutine to replace I with an equivalent instance I' with $|I'| < |I|$ then success would imply $P=NP$ — discouraging efforts in this research direction, from a mathematically-powered point of view.

The situation in regards the systematic, mathematically sophisticated investigation of preprocessing subroutines has changed drastically with advent of parameterized complexity, where the issues are naturally framed. More specifically, we ask for upper bounds on the reduced instance sizes as a function of a parameter of the input, assuming a polynomial time reduction/preprocessing algorithm.

A typical example is the famous Nemhauser-Trotter kernel for the Vertex Cover problem, showing that a “kernel” of at most $2k$ vertices can be obtained, with k the requested maximum size of a solution. A large number of results have been obtained in the past years, and the research in this area shows a rapid growth, not only

in terms of number of papers appearing in top Theoretical Computer Science and Algorithms conferences and journals, but also in terms of techniques. Importantly, very recent developments were the introduction of new lower bound techniques, showing (under complexity theoretic assumptions) that certain problems must have kernels of at least certain sizes, meta-results that show that large classes of problems all have small (e.g., linear) kernels — these include a large collection of problems on planar graphs and matroid based techniques to obtain randomized kernels.

Kernelization is a vibrant and rapidly developing area. This meeting on kernelization consolidated the results achieved in the recent years, discussed future research directions, and explored further the applications potential of kernelization algorithms, and gave excellent opportunities for the participants to engage in joint research and discussions on open problems and future directions. This workshop was also special as we celebrated the 60th birthday of one of the founder of parameterized complexity, Prof. Michael R. Fellows. We organised a special day in which we remembered his contributions to parameterized complexity, science in general and mathematics for children.

The main highlights of the workshop were talks on the solution to two main open problems in the area of kernelization. We give a brief overview of these new developments below.

The AND Conjecture. The OR-SAT problem asks if, given m formulas each of size n , at least one of them is satisfiable. In 2008, Fortnow and Santhanam showed that if there is a reduction from OR-SAT to any language L with the property that the reduction reduces to instances of size polynomial in n (independent of m) then the polynomial-time hierarchy collapses. Such a reduction is called an OR-distillation, and this work motivated the notion of an

OR-composition, which produces a boolean OR of parameterized instances of a given problem, without any restriction on the size. It was then established that an OR-composition and a polynomial kernel cannot co-exist, because these ingredients can be combined to lead to an OR-distillation. Thus, an OR-composition counts as evidence against the existence of a polynomial kernel, and it has turned into a very successful framework for establishing kernel lower bounds.

The question of whether there is similar evidence against the existence of an AND-distillation (defined analogously) has since been open. Such a result would imply that problems that have AND-compositions are also unlikely to admit polynomial kernels, and would therefore be a significant addition to the kernel lower bound toolkit. The question has been a central open problem for the kernelization community and was settled by Drucker in his work on classical and quantum instance compression. The route to the result is quite involved, and forges new connections between classical and parameterized complexity.

Tools from Matroid and Odd Cycle Traversal. The ODD CYCLE TRAVERSAL problem asks if, given a graph G , there is a subset S of size at most k whose removal makes the graph bipartite. Equivalently, the question is if there is a subset S of size at most k that intersects every odd cycle in G . The problem was first shown to be FPT by Reed, Smith, and Vetta in 2004, and this was also the first illustration of the technique of iterative compression. However, the question of whether the problem admits a polynomial kernel was among the main open questions in the study of kernelization.

A breakthrough was recently made in work by Kratsch and Wahlström, providing the first (randomized) polynomial kernelization for the problem. It is a novel approach based on matroid theory, where all relevant information about a problem instance is encoded into a matroid with a representation of size polynomial in k .

Organization of the seminar and activities. The seminar consisted of twenty two talks, a session on open questions, and informal discussions among the participants. The organizers selected the talks in order to have comprehensive lectures giving overview of main topics and communications of new research results. Each day consisted of talks and free time for informal gatherings among participants. On the fourth day of the seminar we celebrated the 60th birthday of Mike Fellows, one of the founder of parameterized complexity. On this day we had several talks on the origin, history and the current developments in the field of parameterized complexity.

4.25 Putting Data on the Map

Organizers: Stephen Kobourov, Alexander Wolff, and Frank van Ham

Seminar No. 12261

Date: June 24–29, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.6.51

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Stephen Kobourov, Alexander Wolff, and Frank van Ham



Participants: Muhammad Jawaherul Alam, Daniel Archambault, David Auber, Peter Bak, Therese Biedl, Robert P. Biuk-Aghai, Sergio Cabello, Sheelagh Carpendale, Walter Didimo, Stephan Diehl, Pierre Dragicevic, Tim Dwyer, Jason Dykes, Sara Irina Fabrikant, Stefan Felsner, Martin Fink, Joe Fowler, Fabrizio Frati, Martin Gronemann, Jan-Henrik Haunert, Herman J. Haverkort, Nathalie Henry Riche, Seok-Hee Hong, Ferran Hurtado, Christophe Hurter, Bernhard Jenny, Michael Kaufmann, Philipp Kindermann, Stephen Kobourov, Jan Kratochvil, Giuseppe Liotta, Anna Lubiw, William A. Mackaness, Petra Mutzel, Quan Nguyen, Arlind Nocaj, Martin Nöllenburg, Yoshio Okamoto, Sergey Pupyrev, Andreas Reimer, Maxwell J. Roberts, Ignaz Rutter, Falko Schmid, Monika Sester, Bettina Speckmann, Claudio Squarcella, Alexandru C. Telea, Barbara Tversky, Torsten Ueckerdt, Pavel Valtr, Frank van Ham, Marc van Kreveld, Kevin Verbeek, Dorothea Wagner, Alexander Wolff, Jing Yang

Visualization allows us to perceive relationships in large sets of interconnected data. While statistical techniques may determine correlations among the data, visualization helps us frame what questions to ask about the data. The design and implementation of algorithms for modeling, visualizing and interacting with large relational data is an active research area in data mining, information visualization, human-computer interaction, and graph drawing.

Map representations provide a way to visualize relational data with the help of conceptual maps as a data representation metaphor. In a narrow sense, a map representation of a graph is a contact graph representation where the adjacency of vertices is expressed by regions that share borders. Such representations are, however, limited to planar graphs by definition. We can extend the notion of a map representation to non-planar graphs by generalizing the idea as follows: clusters of well-connected vertices form countries, and countries share borders when neighboring clusters are tightly interconnected.

Information spatialization and cartograms also connect the notions of data with those of maps. Cartograms redraw an existing geographic map such that the country areas are proportional to some metric (e.g., population), an idea that dates back to a paper by Raisz in 1934 and is still popular today. *Spatialization* is the process of assigning two- or three-dimensional coordinates to abstract data points, ideally such that the spatial mapping has much of the characteristics of the original high-dimensional space. Multi-dimensional scaling or principal component analysis are techniques that allow us to spatialize high-dimensional data. Techniques like information landscapes can then be used to convert the resulting two-dimensional coordinates into meaningful three-dimensional landscapes.

Providing efficient and effective data visualization is a difficult challenge in many real-world software systems. One challenge lies in developing algorithmically efficient methods to visualize large and complex data sets. Another challenge is to develop effective vi-

ualizations that make the underlying patterns and trends easy to see. And finally, we need to allow users to interactively access, analyze, and filter these patterns in an intuitive manner. All of these tasks are becoming increasingly more difficult due to the growth of the data sets arising in modern applications, as well as due to their highly dynamic nature.

■ Topics of the Seminar

Graph representations of side-to-side touching regions tend to be visually appealing and have the added advantage that they suggest the familiar metaphor of a geographical map. Traditional maps offer a natural way to present geographical data (continents, countries, states) and additional properties defined with the help of contours (topography, geology, rainfall).

An important difference between drawings of graphs and maps is the following: graphs are usually *drawn on* the plane (using small placeholder symbols for vertices and curves for edges), whereas maps *fill* the plane (or a sufficiently large area). We want to explore this new paradigm.

In the process of data mining and data analysis, clustering is an extremely important step. It turns out that maps are very helpful in dealing with clustered data. There are several reasons why a map representation of clusters can be helpful. First, by explicitly defining the boundary of the clusters and coloring the regions, we make the clustering information clear. Second, as most dimensionality-reduction techniques lead to a two-dimensional positioning of the data points, a map is a natural generalization. Finally, while it often takes us considerable effort to understand graphs, charts, and tables, a map representation is intuitive, as most people are very familiar with maps and even enjoy carefully examining maps.

When designing algorithms to produce maps for abstract da-

ta, we can leverage cartography and GIS expertise in order to answer critical questions such as how regions and geographic networks (such as street or river networks) are represented on traditional geographic maps, how they are labeled (an interesting problem in its own right) and how (boundary) lines are simplified (through a process called *cartographic generalization*), or even schematized, in order to focus on important features. Therefore, participation of people from several diverse areas is essential for the success of our seminar.

■ Aims of the Seminar

The main goal of this seminar was to foster co-operation between researchers with interests in data visualization coming from the information visualization, human-computer interaction, data mining, graph drawing, and GIS communities. The specific aims of the Dagstuhl seminar were:

1. To bring together researchers working on visualization from a theoretical point of view (graph theory, computational geometry), from a practical point of view (information visualization, HCI), and from a map point of view (cartography, GIS).
2. To identify specific theoretical and practical problems that need to be solved in order to make it possible to create full-fledged conceptual maps as an interactive and scalable data-representation metaphor and to begin working on these problems at the seminar.
3. To formulate the findings as a first step to the solutions of the problems under consideration and to define future research directions.

In order to promote the communication and cooperation between the diverse set of participants, we used a non-traditional format, which included survey presentations, open problem sessions, demo sessions, open mic sessions, problem solving sessions, as well as an exhibition of map-based visualizations. The exhibition entitled “Beyond the Landscape” was organized by seminar participant Maxwell Roberts and by seminar co-organizer Alexander Wolff. It was opened on June 26 by the scientific director of Schloss Dagstuhl, Prof. Reinhard Wilhelm.

■ Achievements of the Seminar

The achievements in the seminar were numerous and varied. Some of the more important ones can be summarized as follows:

1. On Monday and Tuesday, we enjoyed five survey lectures. Jason Dykes discussed geographic data visualization. Sara Fabrikant presented the cartographic and geovisual perspective. Stephen Kobourov talked about visualizing relational data with the help of the map metaphor. Stefan Felsner illustrated connections with geometry and graph theory. Falko Schmid discussed maps and the interaction with geographic data on small mobile devices. Beyond the survey lectures, a highlight of the seminar was the Friday morning lecture by psychology and perception expert Barbara Tversky.
2. We also had a number of stimulating presentations and demos of new software. In particular, new approaches to the layout of large and/or dynamic graphs as well as new visualization paradigms were presented.
3. A number of relevant open problems were formulated early in the seminar and working groups formed around related open 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. Geometric properties of cartograms; convex cartograms
 - b. Evaluation of maps and graphs

- c. Metro map visualization
- d. Semantic word cloud visualization
- e. Edge bundling problems
- f. Multi-dimensional temporal data on maps
- g. Map distortion based on (dis)similarity
- h. Work flow for creating maps out of relational data
- i. Maps based on space-filling curve ordering
- j. Multi-scale map generalizations

The last three days of the seminar were dedicated to working group effort. 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 but two of the groups. We are expecting several 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 maps. The aforementioned exhibition “Beyond the Landscape” made topics of the seminar visible and raised new questions.

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. We thank Philipp Kindermann for helping us to put this report together.

4.26 AI meets Formal Software Development

Organizers: Alan Bundy, Dieter Hutter, Cliff B. Jones, and J Strother Moore
Seminar No. 12271

Date: 01.–06. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.7.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Cliff B. Jones



Participants: Rob Arthan, Serge Autexier, Alan Bundy, Simon Colton, David Crocker, Jorge Cuellar, Ewen W. Denney, Leo Freitas, Dimitra Giannakopoulou, Gudmund Grov, Reiner Hähnle, Dieter Hutter, Andrew Ireland, Moa Johansson, Cliff B. Jones, Ekaterina Komendantskaya, Thierry Lecomte, K. Rustan M. Leino, Michael Leuschel, Yuhui Lin, Maria Teresa Llano Rodriguez, Christoph Lüth, Ursula Martin, Stephan Merz, Rosemary Monahan, J Strother Moore, Michał Moskal, Yannick Moy, José Nuno Oliveira, Thomas Santen, Stephan Schulz, Volker Sorge, Mark Staples, Kurt Stenzel, Werner Stephan, Helen Treharne, Josef Urban, Andrius Velykis, Laurent Voisin, Martin Wehrle

This seminar brought together researchers from formal methods and AI. The participants addressed the issue of how AI can aid the formal software development process, including modelling and proof. There was a pleasing number of participants from industry and this made it possible to ground the discussions on industrial-scale problems.

■ Background

Industrial use of formal methods is certainly increasing but in order to make it more mainstream, the cost of applying formal methods, in terms of mathematical skill level and development time, must be reduced — and we believe that AI can help with these issues.

Rigorous software development using formal methods allows the construction of an accurate characterisation of a problem domain that is firmly based on mathematics; by applying standard mathematical analyses, these methods can be used to prove that systems satisfy formal specifications. A recent ACM computing survey [1] describes over sixty industrial projects and discusses the effect formal methods have on time, cost, and quality. It shows that with tools backed by mature theory, formal methods are becoming cost effective and their use is easier to justify, not as an academic exercise or legal requirement, but as part of a business case. Furthermore, the use of such formal methods is no longer confined to safety critical systems: the list of industrial partners in the DEPLOY project² is one indication of this broader use. Most methods tend to be posit-and-prove, where the user posits a development step (expressed in terms of specifications of yet-to-be-realised components) that has to be justified by proofs. The associated properties that must be

verified are often called proof obligations (POs) or verification conditions. In most cases, such proofs require mechanical support by theorem provers.

One can distinguish between automatic and interactive provers, where the latter are generally more expressive but require user interaction. Examples of state-of-the-art interactive theorem provers are ACL2, Isabelle, HOL, Coq and PVS, while E, SPASS, Vampire and Z3 are examples of automatic provers.

AI has had a large impact on the development of provers. In fact, one of the first AI application was a theorem prover and all theorem provers now contain heuristics to reduce the search space that can be attributed to AI. Nevertheless, theorem proving research and (pure) AI research have diverged, and theorem proving is barely considered to be AI-related anymore.

There follows a list of background references.

■ Organisation of the seminar

It might be useful to organisers of future seminars to record some organisational issues. We asked participants to prepare only short talks that introduced topics and –just as we wished– a number of the talks were actually prepared at the seminar location and with the benefit of having heard other talks. This free format worked well for our exchange of ideas and in most regards we were pleased that we started with only the Monday morning actually scheduled. Perhaps the biggest casualty of the fluid organisation (coupled with so many interesting participants) was that there was no time left for Panel Sessions. However, the differing lengths of discussions (and liberal use of breaks and a “hike” for people to establish new links) led to intensive interaction.

² DEPLOY was an EU-funded “IP” led by Newcastle University; a four year project with a budget of about 18M Euros; the industrial collaborators include Siemens Transport, Bosch and SAP.

It is a pleasure to extend our thanks to everyone involved in the Dagstuhl organisation: they provide a supportive and friendly context in which such fruitful scientific exchanges can develop unhindered by distraction.

■ Results

It is possible to address the results under the phases of the development cycle. Requirements capture is traditionally a pre-formal exercise and is the phase where one would expect least impact from formal ideas. There is certainly scope here for the use of ontologies and some hope for help in detecting inconsistencies in requirements but little time was spent in the seminar on these topics.

Once development moves to the creation of a specification, the scope for formalism increases and with it the hope for a greater contribution from AI. Essentially, a formal specification is a model. Formal proof can be used to establish internal consistency properties or to prove that properties match expectations about the required system. Model checking approaches are often the most efficient way of detecting inconsistencies.

Steps of development (in the posit and prove approaches) essentially introduce further models which should relate in precise ways to each other. The technical details vary between development methods but the overall implications for the use of proof and the contribution of AI are similar. It is perhaps worth reemphasising here that the seminar was trying to address problems of an industrial scale.

An interesting dichotomy was explored at the seminar concerning POs that fail to discharge. One school of thought is to interpose extra models in order to cause the generation of simpler POs; the

alternative is to take the POs as fixed and develop “theories” (collections of auxiliary functions and lemmas) to complete the proof process. Suffice it to say here that AI was seen to have a role in both approaches.

More generally, the whole task of refactoring models and reusing libraries of established material is another area seen as being in need of help from AI thinking.

Turning to the richest area of collaboration –that of proof itself– a prominent theme was on the ways in which machine learning can help. There are many facets of this question including analogy with previous proofs, data mining of proofs (and failures) and proof strategy languages.

One particularly important aspect of the cost of proof in an industrial setting is proof maintenance. In practical settings, many things change and it is unlikely to be acceptable to have to repeat the whole proof process after each change.

Another area that led to useful interactions between participants was the subject of failure analysis and repair. It was observed that it is useful to have strong expectations as to how proofs were meant to succeed.

In conclusion many points of contact can be seen in the material presented below. Unsurprisingly, the material ranges from hopes for future research to mature results that can be readily applied. It is not only a hope that the links between ideas and researchers made at the seminar will persist — we already have clear proof of collaborative work.

The four organisers are extremely grateful to Andrius Velykis who took on the whole of the task of collecting and tidying the contributions of the speakers.

■ References

- 1 J. Woodcock, P. G. Larsen, J. Bicarregui, and J. S. Fitzgerald. Formal methods: Practice and experience. *ACM Computing Surveys*, 41(4), 2009.
- 2 A. Bundy. *The Computer Modelling of Mathematical Reasoning*. Academic Press, 1983. (2nd edition 1986).
- 3 A. Bundy and A. Smaill. *A Catalogue of Artificial Intelligence Techniques*. Springer, 1984. (2nd edition 1986, 3rd edition 1990, 4th revised edition 1997).
- 4 A. Bundy, F. van Harmelen, J. Hesketh, and A. Smaill. Experiments with proof plans for induction. *J. Autom. Reasoning*, 7(3):303–324, 1991.
- 5 A. Bundy. A science of reasoning. In J.-L. Lassez and G. Plotkin, editors, *Computational Logic: Essays in Honor of Alan Robinson*, pages 178–198. MIT Press, 1991.
- 6 A. Bundy. The automation of proof by mathematical induction. In J. A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, pages 847–911. Elsevier and MIT Press, 2001.
- 7 A. Bundy, D. Basin, D. Hutter, and A. Ireland. *Rippling: Meta-level Guidance for Mathematical Reasoning*, volume 56 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, June 2005.
- 8 D. Hutter and W. Stephan, editors. *Mechanizing Mathematical Reasoning, Essays in Honor of Jörg H. Siekmann on the Occasion of His 60th Birthday*, volume 2605 of *LNCS*. Springer, 2005.
- 9 D. Hutter, W. Stephan, P. Traverso, and M. Ullmann, editors. *Proceedings of Current Trends in Applied Formal Methods (FM-Trends 98)*, volume 1641 of *LNCS*, Boppard, Germany, 1999. Springer.
- 10 C. B. Jones, K. D. Jones, P. A. Lindsay, and R. Moore. *mural: A Formal Development Support System*. Springer, 1991.
- 11 R. S. Boyer and J. S. Moore. *A Computational Logic Handbook*. Formal Methods Series. Academic Press, second edition, 1997.
- 12 M. Kaufmann, P. Manolios, and J. S. Moore. *Computer-Aided Reasoning: An Approach*, volume 3 of *Advances in Formal Methods*. Kluwer Academic Publishers, 2000.
- 13 M. Kaufmann and J. S. Moore. Some key research problems in automated theorem proving for hardware and software verification. *Revista de la Real Academia de Ciencias (RACSAM)*, 98(1):181–196, 2004.

4.27 Architecture-Driven Semantic Analysis of Embedded Systems

Organizers: Peter Feiler, Jérôme Hugues, and Oleg Sokolsky

Seminar No. 12272

Date: 01.–06. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.7.30

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Peter Feiler, Jérôme Hugues, and Oleg Sokolsky



Participants: De-Jiu Chen, Silvano Dal Zilio, Patricia Derler, Mamoun Filali-Amine, Laurent Fournier, Serban Gheorghe, Lars Grunske, Jérôme Hugues, Naoki Ishihama, Gabor Karsai, Alexey Khoroshilov, Fabrice Kordon, Brian Larson, Bruce Lewis, Ivano Malavolta, Frédéric Mallet, Pieter J. Mosterman, Viet Yen Nguyen, Thomas Noll, Peter Csaba Ölveczky, Paul Pettersson, Ramon Serna Oliver, Shin'ichi Shiraishi, Frank Singhoff, Oleg Sokolsky, Jean-Pierre Talpin, Michael W. Whalen

Architectural modeling of complex embedded systems is gaining prominence in recent years, both in academia and in industry. An architectural model represents components in a distributed system as boxes with well-defined interfaces, connections between ports on component interfaces, and specifies component properties that can be used in analytical reasoning about the model. Models are hierarchically organized, so that each box can contain another system inside, with its own set of boxes and connections between them. An architecture description language for embedded systems, for which timing and resource availability form an important part of the requirements, must describe resources of the system platform, such as processors, memories, communication links, etc. Several architectural modeling languages for embedded systems have emerged in recent years, including AADL, SysML, EAST-ADL, and the MARTE profile for UML.

In the context of model-based engineering (MBE) architectural modeling serves several important purposes:

An architectural model allows us to break the system into manageable parts and establish clear interfaces between these parts. In this way, we can manage complexity of the system by hiding the details that are unimportant at a given level of consideration; Clear interfaces between the components allow us to avoid integration problems at the implementation phase. Connections between components, which specify how components affect each other, help propagate the effects of change in one component to the affected components. Most importantly, an architectural model can be seen as a repository of the knowledge about the system, represented as requirements, design, and implementation artifacts, held together by the architecture. Such a repository enables automatic generation of analytical models for different aspects of the system, such as timing, reliability, security, performance, etc. Since all the models are generated from the same source, ensuring consistency of assumptions and abstractions used in different analyses becomes easier. The first two uses of architectural modeling have been studied in the research

literature for a number of years. However, the coordination role of architectural modeling in MBE is just currently emerging. We expect this role to gain importance in the coming years. It is clear that realizing this vision of “single-source” MBE with an architectural model at its core is impossible without having first a clear semantics of the architecture description language.

The goal of the seminar is to bring together researchers who are interested in defining precise semantics of an architecture description language and using it for building tools that generate analytical models from architectural ones, as well as generate code and configuration scripts for the system. Despite recent research activity in this area to use semantic interpretation of architectural models for analytical model generation, we observe a significant gap between current state of the art and the practical need to handle complex models. In practice, most approaches cover a limited subset of the language and target a small number of modeling patterns. A more general approach would most likely require an interpretation of the semantics of the language by the tool, instead of hard-coding of the semantics and patterns into the model generator.

4.28 Security and Dependability for Federated Cloud Platforms

Organizers: Rüdiger Kapitza, Matthias Schunter, Marc Shapiro, Paulo Verissimo, and Michael Waidner

Seminar No. 12281

Date: 8.–13. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.7.56

© ⓘ ⓘ Creative Commons BY-NC-ND 3.0 Unported license

© Rüdiger Kapitza, Matthias Schunter, Marc Shapiro, Paulo Verissimo, and Michael Waidner

Participants: Artur Andrzejak, Alysson Neves Bessani, Herbert Bos, Sara Bouchenak, Gregory Chockler, Miguel Pupo Correia, Maria Couceiro, Alexandra Dmitrienko, Dan Dobre, Kurt Geihs, Nicolas Geoffray, Gabriela Gheorghe, Stephan Groß, Flavio Paiva Junqueira, Rüdiger Kapitza, Nikola Knezevic, Guillaume Pierre, Peter R. Pietzuch, Hans Peter Reiser, Luis Rodrigues, Ahmad-Reza Sadeghi, Nuno Santos, Matthias Schunter, Marc Shapiro, Alexander Shraer, Radu Sion, Jan Stoess, Paulo Jorge Verissimo, Marko Vukolic, Michael Waidner, Alexander Wiesmaier



Computing services are increasingly pooled within global utility computing infrastructures offered by providers such as Amazon, Google and IBM. Infrastructure clouds provide virtual machines and resources. These infrastructure clouds are used to enable “platforms as a service” that simplify implementation of arbitrary scalable services.

The seminar targeted the management and protection of individual clouds and addressed the trend towards cloud federation by bringing together researchers from systems management, security, and dependability. The idea was that only such an integrated approach is able to guarantee security and dependability while preserving the essential cost and efficiency benefits of today’s emerging solutions.

The challenge to address was how to provide secure and dependable services on such federated cloud platforms. Selected research questions were: How can clouds securely interoperate, how can service availability be guaranteed despite failures or attacks by individual clouds, how can existing algorithms be adjusted to provide scalable eventual consistency, and finally whether cloud-of-cloud infrastructures can provide such benefits at costs that are competitive with single cloud solutions.

While these questions were addressed during the seminar it got also clear that dependability and security of single clouds is by far not solved and therefore was also discussed in depth.

4.29 Database Workload Management

Organizers: Shivnath Babu, Goetz Graefe, and Harumi Anne Kuno
Seminar No. 12282

Date: 09.–13. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.7.73

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Shivnath Babu, Goetz Graefe, and Harumi Anne Kuno



Participants: Ashraf Aboulnaga, Awny Al-Omari, Shivnath Babu, Michael A. Bender, Douglas Brown, Robert J. Chansler, Yanpei Chen, Jens Dittrich, Archana Ganapathi, Goetz Graefe, Hakan Haciguemus, Rao Kakarlamudi, Harumi A. Kuno, Bradley C. Kuszmaul, Barzan Mozafari, Sivaramakrishnan Narayanan, Wendy Powley, Norbert Ritter, Kai-Uwe Sattler, Russell Sears, Michael Seibold, Y. C. Tay, Kostas Tzoumas, Jingren Zhou

Much database research focuses on improving the performance of individual queries. Workload management focuses on a larger question – how to optimize the performance of the entire workload, as a whole. Workload management is one of the most expensive components of system administration. Gartner listed workload management as the first of two key challenges to emerge from the data warehouse market in 2009. However, we believe that even while both researchers and industry are building and experimenting with increasingly large-scale workloads, there is a disconnect between the OLTP/OLAP/Mixed/Hadoop/Map-Reduce workloads used in experimental research and the complex workloads that practitioners actually manage on large-scale data management systems.

One goal of this seminar was to bridge this gap between research and practice. Dagstuhl Seminar 12282 provided a venue where researchers can engage in dialogue with industrial participants for an in-depth exploration of challenging industrial workloads, where industrial participants can challenge researchers to apply the lessons-learned from their large-scale experiments to multiple real systems, and that would facilitate the release of real workloads that can be used to drive future research, and concrete measures to evaluate and compare workload management techniques in the context of these workloads.

With regard to seminar participants, we took a system-centric focus, and invited participants who could speak to the management of workloads in a variety of systems. Seminar participants came from a variety of academic and commercial institutions: Cloudera, EMC/Greenplum, LinkedIn, Microsoft, MIT, National University of Singapore, NEC, Queen's University, Stony Brook University, Teradata, Tokutek, TU Berlin, TU Ilmenau, TU München, UC Berkeley, Universität des Saarlandes, Universität Hamburg, University of Waterloo, and Yahoo.

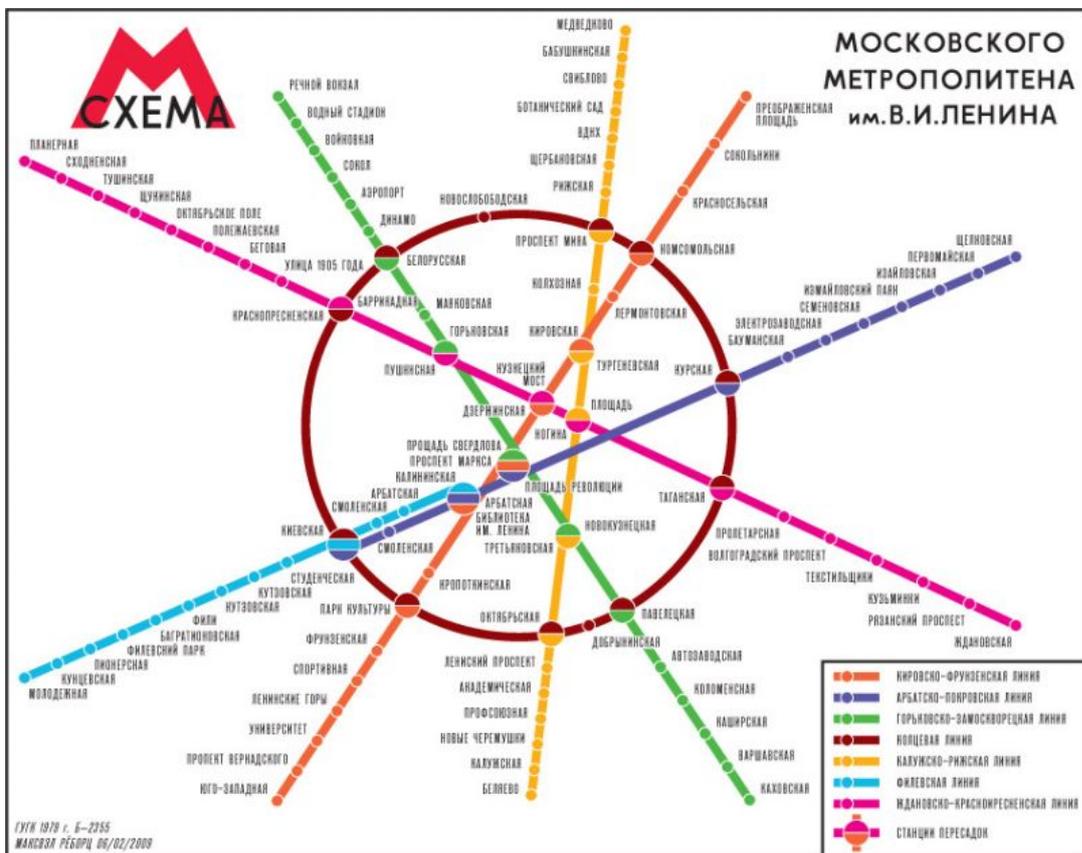


Fig. 4.6

Maxwell J. Robert – Moscow. Part of the Dagstuhl art collection and donated by participants of the Dagstuhl Perspectives Workshop 10482.

4.30 Structure Discovery in Biology: Motifs, Networks & Phylogenies

Organizers: Alberto Apostolico, Andreas Dress, and Laxmi Parida

Seminar No. 12291

Date: 15.–20. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.7.92

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Alberto Apostolico, Andreas Dress, and Laxmi Parida



Participants: Jörg Ackermann, Jonas Almeida, Alberto Apostolico, Benny Chor, Matteo Comin, Eduardo Corel, Yupeng Cun, Fabio Cunial, Gilles Didier, Andreas Dress, Péter L. Erdős, Ayse Funda Ergun, Mareike Fischer, Matthias Gallé, Raffaele Giancarlo, Alex Grossmann, Stefan Gruenewald, Concettina Guerra, Katharina T. Huber, Deok-Soo Kim, Jack Koolen, James A. Lake, Qiang Li, Matthias Löwe, Tobias Marschall, Christoph Mayer, Burkhard Morgenstern, Axel Mosig, Laxmi Parida, Alex Pothen, Sven Rahmann, Walter Schubert, Rahul Siddharthan, Andreas Spillner, Peter F. Stadler, Mike Steel, Gabriel Valiente, Susana Vinga, Shuhua Xu, Jun Yan

In biological systems, similarly to the tenet of modern architecture, form and function are solidly intertwined. Thus to gain complete understanding in various contexts, the curation and study of form turns out to be a mandatory first phase.

Biology is in the era of the “Omes”: Genome, Proteome, Toponome, Transcriptome, Metabolome, Interactome, ORFeome, Recombinome, and so on. Each Ome refers to carefully gathered data in a specific domain. While biotechnology provides the data for most of the Omes (sequencing technology for genomes, mass spectrometry and toponome screening for proteomes and metabolomes, high throughput DNA microarray technology for transcriptomes, protein chips for interactomes), bioinformatics algorithms often help to process the raw data, and sometimes even produce the basic data such as the ORFeome and the recombinome.

The problem is: biological data are accumulating at a much faster rate than the resulting datasets can be understood. For example, the 1000-genomes project alone will produce more than 10^{12} raw nucleic acid bases to make sense of. Thus, databases in the terabytes, even petabytes (10^{15} bytes) range are the norm of the day. One of the issues today is that our ability to analyze and understand massive datasets lags far behind our ability to gather and store the data with the ever advancing bio- and computing technologies. So, while the sheer size of data can be daunting, this provides a golden opportunity for testing (bioinformatic) structure-discovery primitives and methods.

Almost all of the repositories mentioned here are accompanied by intelligent sifting tools. In spite of the difficulties of structure discovery, supervised or unsupervised, there are reasons to believe that evolution endowed biological systems with some underlying principles of organization (based on optimization, redundancy, similarity, and so on) that appear to be present across the board. Correspondingly, using evolutionary thoughts as a “guiding light”, it should be possible to identify a number of primitive characteristics of the various embodiments of form and structure (for instance, sim-

ply notions of maximality, irredundancy, etc.) and to build similarly unified discovery tools around them. Again, the forms may be organized as linear strings (say, as in the genome), graphs (say, as in the interactome), or even just conglomerates (say, as in the transcriptome). And the fact that even the rate of data accumulation increases continuously becomes rather a blessing in this context than a curse. It is therefore a worthwhile effort to try and identify these primitives. This seminar was intended to focus on combinatorial and algorithmic techniques of structure discovery relating to biological data that are at the core of understanding a coherent body of such data, small or large. The goal of the seminar was twofold: on one hand to identify concise characterizations of biological structure that span across multiple domains; on the other to develop combinatorial insight and algorithmic techniques to effectively unearth structure from data.

The seminar began with a town-hall, round-table style meeting where each participant shared with the others a glimpse of their work and questions that they were most excited about. This formed the basis of the program that was drawn up democratically. As the days progressed, the program evolved organically to make an optimal fit of lectures to the interest of the participants.

The first session was on population genomics, covered by Shuhua Xu and Laxmi Parida. The second was on methods on genomic sequences, covered by Rahul Siddharthan and Jonas Almeida. The next talks were on clinical medicine: an interesting perspective from a practicing physician, Walter Schubert, on treatment of chronic diseases, and Yupeng Cun spoke about prognostic biomarker discovery. Algorithms and problems in strings or genomic sequences were covered in an after-dinner session on Monday and in two sessions on Tuesday morning and late afternoon. The speakers were Sven Rahmann, Burkhard Morgenstern, Eduardo Corel, Fabio Cunial, Gilles Didier, Tobias Marschall, Matthias Gallé, Susana Vinga and Gabriel Valiente. The last speaker presented a system called “Tango” on metagenomics, and in a bizarre twist concluded the session and the

day with a surprise live Argentine Tango dance performance with one of the organizers of the seminar. The early afternoon session was on metabolic networks, with lectures by Jörg Ackermann, Jun Yan and Qiang Li.

The Wednesday morning session was loosely on proteomics, with lectures by Alex Pothen, Benny Chor, Axel Mosig, Alex Grossmann, and Deok-Soo Kim. Coincidentally, three lecturers of this session shared very similar first names, leading to some gaffes and some light moments at the otherwise solemn meeting.

The Thursday sessions were on phylogenies and networks, with lectures by Mareike Fischer, Mike Steel, Katharina T. Huber, Christoph Mayer, James A. Lake, Péter L. Erdős, Stefan Gruenewald and Peter F. Stadler. James A. Lake presented an interesting shift in paradigm, based in biology, called *cooperation and competition in phylogeny*. Péter L. Erdős gave a fascinating talk on the realization of degree sequences. Yet another session on strings was covered by Matteo Comin and Funda Ergun on Thursday. The day concluded with a lecture by Andreas Dress on pandemic modeling.

There were a few after-dinner sessions on big data, thanks to Jonas Almeida. An eclectic set of lectures were given on the last session on Friday, by Raffaele Giancarlo on clustering and by Concettina Guerra on network motifs. The meeting concluded with a fascinating lecture by Matthias Löwe on the combinatorics of graph sceneries. The impact of this on biology may not be immediately clear, but such is the intent of these far-reaching, outward-looking seminars.

4.31 Robust Query Processing

Organizers: Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley
Seminar No. 12321

Date: 09.–13. July, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.8.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn Paulley



Participants: Martina-Cezara Albutiu, Peter A. Boncz, Renata Borovica, Surajit Chaudhuri, Campbell Fraser, Johann Christoph Freytag, Goetz Graefe, Ralf Hartmut Güting, Wey Guy, Theo Härder, Fabian Hüske, Stratos Idreos, Ihab Francis Ilyas, Alekh Jindal, Martin L. Kersten, Harumi Anne Kuno, Andrew Lamb, Allison Lee, Stefan Manegold, Anisoara Nica, Glenn Paulley, Ilia Petrov, Meikel Poess, Ken Salem, Bernhard Seeger, Krzysztof Stencel, Knut Stolze, Florian M. Waas, Jianliang Xu, Marcin Zukowski

■ Introduction

In early August 2012 researchers from both academia and industry assembled in Dagstuhl at the 2012 Dagstuhl Workshop on Robust Query Processing, Workshop 12321. An earlier Workshop—Dagstuhl Workshop 10381—held in September 2010 [16] had supplied an opportunity to look at issues of Robust Query Processing but had failed to make significant progress in exploring the topic to any significant depth. In 2012, 12321 Workshop participants looked afresh at some of the issues surrounding Robust Query Processing with greater success and with the strong possibility of future publications in the area that would advance the state-of-the-art in query processing technology.

■ Background and related research

A considerable amount of query processing research over the past 20 years has focused on improving relational database system optimization and execution techniques for complex queries and complex, ever-changing workloads. Complex queries provide optimization challenges because selectivity and cardinality estimation errors multiply, and so there is a large body of work on improving cardinality estimation techniques and doing so in an automatic fashion: from capturing histogram information at run time [1, 17], to mitigating the effects of correlation on the independence assumption [21], to utilizing constraints to bound estimation error [9, 10, 15, 18], to permitting various query rewritings to simplify the original statement [11, 19, 23, 26–28]. Studies of the feasibility of query re-optimization [7, 8], or deferring optimization to execution time [24], have until recently largely been based on the premise that the need for such techniques is due either to recovering from estimation errors at optimization time in the former case, or avoiding the problem entirely by performing all optimization

on-the-fly, such as with Eddies [6] rather than in a staged, ‘water-fall’ kind of paradigm.

More recent work on adaptive query processing [13, 14, 24, 25] has considered techniques to handle the interaction of query workloads [3–5], coupled with the realization that changes to environmental conditions can significantly impact a query’s chosen execution plan. These environmental conditions include:

- changes to the amount of memory available (buffer pool, heap memory);
- changes to i/o bandwidth due to concurrent disk activity;
- locking and other waits caused by concurrency control mechanisms;
- detected flaws in the currently executing plan;
- number of available CPU cores;
- changes to the server’s multiprogramming level [2];
- changes to physical access paths, such as the availability of indexes, which could be created on the fly;
- congestion with the telecommunications network;
- contents of the server’s buffer pool;
- inter-query interaction (contention on the server’s transaction log, ‘hot’ rows, and so on).

■ Background – Dagstuhl seminar 10381

Self-managing database technology, which includes automatic index tuning, automatic database statistics, self-correcting cardinality estimation in query optimization, dynamic resource management, adaptive workload management, and many other approaches, while both interesting and promising, tends to be studied in isolation of other server components. At the 2010 Dagstuhl Workshop on Robust Query Processing (Dagstuhl seminar 10381) held on 19–24 September 2010, seminar attendees tried to unify the study of these technologies in three fundamental ways:

1. determine approaches for evaluating these technologies in the

‘real’ environment where these independently-developed components would interact;

2. establish a metric with which to measure the ‘robustness’ of a database server, making quantitative evaluations feasible so as to compare the worthiness of particular approaches. For example, is dynamic join reordering during query execution worth more than cardinality estimation feedback from query execution to query optimization?
3. utilize a metric, or metrics, to permit the construction of regression tests for particular systems. The development of suitable metrics could lead to the development of a new, possibly industry-standard benchmark, that could be used to measure self-managing database systems by industry analysts, customers, vendors, and academic researchers and thus lead to better improvements in robust operation.

At the 2010 Dagstuhl seminar, attendees struggled somewhat with trying to define the notion of robustness, let alone trying to measure or quantify it. Robustness is, arguably, somewhat orthogonal to absolute performance; what we are trying to assess is a system’s ability to continue to operate in the face of changing workloads, system parameters and environmental conditions.

An example of the sorts of problems encountered in trying to define robustness is illustrated in Figure 4.7. Figure 4.7 shows the throughput rates of two systems, System A (blue line) and System B (red line), over time, for the same workload. The *Y*-axis represents the throughput rate, and the *X*-axis is elapsed time. Over time, the workload steadily increases.

Three areas of the graph are highlighted in Figure 4.7. The first, in green, shows that as the workload is increased, System A outperforms System B by some margin. That peak performance cannot be maintained, however, as the load continues to be increased. The area in blue shows that once System A becomes overloaded, performance drops precipitously. On the other hand, System B shows a much more gradual degradation (circled in red), offering more robust behaviour than System A but with the tradeoff of not being able to match System A’s peak performance.

One can argue that Figure 4.7 mixes the notions of query processing and workload management. In Figure 4.8 we simplify the problem further, and consider only simple range queries (using two columns) over a single table, where the (logarithmic) *X*-axis denotes the size of the result set.

In Figure 4.8, the yellow line illustrates a table scan: it is robust—it delivers identical performance over all result sets—but with relatively poor performance. The dashed red line is a traditional index-to-index lookup plan: that is, search in secondary index, row fetch out of the primary (clustered) index. This plan is considerably faster for very small selectivities, but becomes considerably poorer with only a marginal decrease in selectivity. The solid red line shows, in comparison, substantial-but-imperfect improvements over the index-to-index technique, due to asynchronous pre-fetch coupled with sorting batches of row pointers obtained from the secondary index. This query execution strategy is available in

Microsoft SQL Server 2005. While Figure 4.8 is just one simple query—one table, range predicates on two columns—Figure 4.8 illustrates both the magnitude of the problem and the opportunity for improving the robustness of such plans.

At the 2010 Dagstuhl seminar, seminar attendees explored a number of different ways in which to define robustness. One idea was to define a metric for robustness as the accumulated variance in the wall clock times of workloads—or particular queries—or, alternatively, some measure of the distribution of that variance, a 2nd level effect. Since this working definition includes wall clock times, it implicitly includes factors such as optimizer quality, since a robustness metric such as this must include statement execution times. However, while the sophistication of query optimization, and the quality of access plans, is a component of a robust database management system, it is not the only component that impacts any notion of robustness.

This working definition of robustness raised as many questions as answers, and many of these were still unresolved by the end of the workshop. Those questions included:

- Sources of variance in query optimization include the statistics model, the cardinality model, and the cost model, with the latter usually being less critical in practice than the former two. One measure of ‘robustness’ is to assess the accuracy between estimates and actuals. What level of importance should the ‘correctness’ of a query optimizer have on a metric of robustness?
- Which offers more opportunity for disaster—and disaster prevention: robust query optimization or robust query execution?
- Is robustness a global property? Does it make sense to measure robustness at the component level? If so, which components, and what weight should be attached to each?
- Several of the attendees at the 2010 Dagstuhl workshop advocated a two-dimensional tradeoff between actual performance and ‘consistency’. But what is ‘consistency’? Is it merely plan quality, or something more?
- Robustness for who? Expectations are different between product engineers and end users; one should not try to define robustness unless one addresses whose expectations you are trying to satisfy. Both rely on an idealized model of how a system should behave. Can we define that model? At the same time, what expectations can a user have of a really complex query?
- Is adaptivity the only way to achieve robustness?
- What would a benchmark for robustness attempt to measure?

During the workshop we analyzed these questions from various perspectives. Unfortunately we failed to reach consensus on a clear definition of robustness, how to measure it, and what sorts of tradeoffs to include. Our hope, in this, the second Dagstuhl workshop on Robust Query Processing, is to make additional progress towards clarifying the problems, and possibly make some progress towards defining some general—or specific—approaches to improve DBMS robustness.

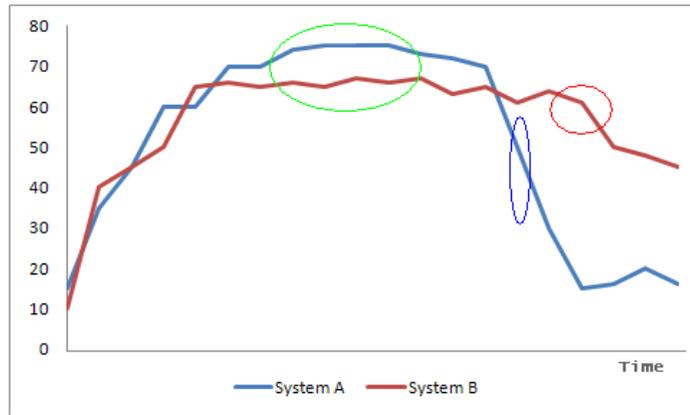


Fig. 4.7 Comparison of Systems A and B in response to increasing workloads over time



Fig. 4.8 Comparison of access plans for a single table range query

References

- 1 A. Aboulnaga and S. Chaudhuri. Self-tuning histograms: Building histograms without looking at data. In *ACM SIGMOD Int'l Conf. on Management of Data*, pp. 181–192, Philadelphia, Pennsylvania, May 1999.
- 2 M. Abouzour, K. Salem, and P. Bumbulis. Automatic tuning of the multiprogramming level in Sybase SQL Anywhere. In *ICDE Workshops*, pp. 99–104. IEEE, 2010.
- 3 M. Ahmad, A. Aboulnaga, S. Babu, and K. Munagala. QShuffler: Getting the query mix right. In *Proc. of the IEEE Int'l Conf. on Data Engineering*, pp. 1415–1417, 2008.
- 4 M. Ahmad, A. Aboulnaga, S. Babu, and K. Munagala. Interaction-aware scheduling of report-generation workloads. *The VLDB Journal*, 20:589–615, August 2011.
- 5 M. Ahmad, S. Duan, A. Aboulnaga, and S. Babu. Predicting completion times of batch query workloads using interaction-aware models and simulation. In *Proc. of the 14th Int'l Conf. on Extending Database Technology*, pp. 449–460, ACM, 2011.
- 6 R. Avnur and J.M. Hellerstein. Eddies: Continuously adaptive query processing. In *ACM SIGMOD Int'l Conf. on Management of Data*, pp. 261–272, 2000.
- 7 P. Bizarro, N. Bruno, and D.J. DeWitt. Progressive parametric query optimization. *IEEE Trans. on Knowledge and Data Engineering*, 21:582–594, 2009.
- 8 P.G. Bizarro. *Adaptive query processing: dealing with incomplete and uncertain statistics*. PhD thesis, University of Wisconsin at Madison, Madison, Wisconsin, 2006.
- 9 Surajit Chaudhuri, Hongrae Lee, and Vivek R. Narasayya. Variance aware optimization of parameterized queries. In *ACM SIGMOD International Conference on Management of Data*, pages 531–542, 2010.
- 10 Surajit Chaudhuri, Vivek R. Narasayya, and Ravishankar Ramamurthy. A pay-as-you-go framework for query execution feedback. *Proceedings of the VLDB Endowment*, 1(1):1141–1152, 2008.
- 11 Mitch Cherniack. *Building Query Optimizers with Combinators*. PhD thesis, Brown University, Providence, Rhode Island, May 1999.
- 12 Richard Cole, Florian Funke, Leo Giakoumakis, Wey Guy, Alfons Kemper, Stefan Krompass, Harumi Kuno, Raghunath Nambiar, Thomas Neumann, Meikel Poess, Kai-Uwe Sattler, Michael Seibold, Eric Simon, and Florian Waas. The mixed workload CH-benCHmark. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, New York, NY, USA, 2011. ACM.
- 13 Amol Deshpande, Zachary G. Ives, and Vijayshankar Raman. Adaptive query processing. *Foundations and Trends in Databases*, 1(1):1–140, 2007.
- 14 Kwanchai Eurviriyankul, Norman W. Paton, Alvaro A. A. Fernandes, and Steven J. Lynden. Adaptive join processing in pipelined plans. In *13th International Conference on Extending Database Technology (EDBT)*, pages 183–194, 2010.
- 15 Parke Godfrey, Jarek Gryz, and Calisto Zuzarte. Exploiting constraint-like data characterizations in query optimization. In *ACM SIGMOD International Conference on Management of Data*, pages 582–592, Santa Barbara, California, May 2001. Association for Computing Machinery.
- 16 Goetz Graefe, Arnd Christian König, Harumi Kuno, Volker Markl, and Kai-Uwe Sattler. Robust query processing. Dagstuhl Workshop Summary 10381, Leibniz-Zentrum für Informatik, Wadern, Germany, September 2010.
- 17 Michael Greenwald. Practical algorithms for self-scaling histograms or better than average data collection. *Performance Evaluation*, 20(2):19–40, June 1996.
- 18 Jarek Gryz, Berni Schiefer, Jian Zheng, and Calisto Zuzarte. Discovery and application of check constraints in DB2. In *Proceedings, Seventeenth IEEE International Conference on Data Engineering*, pages 551–556, Heidelberg, Germany, April 2001. IEEE Computer Society Press.
- 19 Waqar Hasan and Hamid Pirahesh. Query rewrite optimization in STARBURST. Research Report RJ6367, IBM Corporation, Research Division, San Jose, California, August 1988.
- 20 Stratos Idreos, Martin L. Kersten, and Stefan Mane-gold. Database cracking. In *CIDR*, pages 68–78. www.cidrdb.org, 2007.
- 21 Ihab F. Ilyas, Volker Markl, Peter J. Haas, Paul Brown, and Ashraf Aboulnaga. CORDS: Automatic discovery of correlations and soft functional dependencies. In *ACM SIGMOD International Conference on Management of Data*, pages 647–658, Paris, France, June 2004.
- 22 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisora Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, pages 7:1–7:6, New York, NY, USA, 2011. ACM.
- 23 Jonathan J. King. QUIST—A system for semantic query optimization in relational databases. In *Proceedings of the 7th International Conference on Very Large Data Bases*, pages 510–517, Cannes, France, September 1981. IEEE Computer Society Press.
- 24 Volker Markl, Vijayshankar Raman, David E. Simmen, Guy M. Lohman, and Hamid Pirahesh. Robust query processing through progressive optimization. In *ACM SIGMOD International Conference on Management of Data*, pages 659–670, 2004.
- 25 Rimma V. Nehme, Elke A. Rundensteiner, and Elisa Bertino. Self-tuning query mesh for adaptive multi-route query processing. In *12th International Conference on Extending Database Technology (EDBT)*, pages 803–814, 2009.
- 26 G. N. Paulley and Per-Åke Larson. Exploiting uniqueness in query optimization. In *Proceedings, Tenth IEEE International Conference on Data Engineering*, pages 68–79, Houston, Texas, February 1994. IEEE Computer Society Press.
- 27 Hamid Pirahesh, Joseph M. Hellerstein, and Waqar Hasan. Extensible/rule based query rewrite optimization in STARBURST. In *ACM SIGMOD International Conference on Management of Data*, pages 39–48, San Diego, California, June 1992. Association for Computing Machinery.
- 28 H. J. A. van Kuijk. The application of constraints in query optimization. Memoranda Informatica 88–55, Universiteit Twente, Enschede, The Netherlands, 1988.

4.32 Mobility Data Mining and Privacy

Organizers: Christopher W. Clifton, Bart Kuijpers, Katharina Morik, and Yucel Saygin
Seminar No. 12331

Date: 12.–17. August, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.8.16

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Chris Clifton and Bart Kuijpers



Participants: Gennady Andrienko, Nilgün Basalp, Joachim Biskup, Erik Buchmann, Christopher W. Clifton, Maria Luisa Damiani, Glenn Geers, Aris Gkoulalas-Divanis, Marco Gruteser, Christine Körner, Bart Kuijpers, Thomas Liebig, Bradley Malin, Michael Marhöfer, Walied Othman, Klaus Rechert, Claudia Santa, ErKay Savas, Florian Schaub, Olaf Spinczyk, Christian Wietfeld, Ouri E. Wolfson

Mobility Data Mining and Privacy aimed to stimulate the emergence of a new research community to address mobility data mining together with privacy issues. Mobility data mining aims to extract knowledge from movement behaviour of people. This is an interdisciplinary research area combining a variety of disciplines such as data mining, geography, visualization, data/knowledge representation, and transforming them into a new context of mobility while considering privacy which is the social aspect of this area. The high societal impact of this topic is mainly due to the two related facets of its area of interest, i.e., people's movement behaviour, and the associated privacy implications. Privacy is often associated with the negative impact of technology, especially with recent scandals in the US such as AOL's data release which had a lot of media coverage. The contribution of *Mobility Data Mining and Privacy* is to turn this negative impact into positive impact by investigating how privacy technology can be integrated into mobility data mining. This is a challenging task which also imposes a high risk, since nobody knows what kinds of privacy threats exist due to mobility data and how such data can be linked to other data sources.

The seminar looked closely at two application areas: Vehicular data and cellular data. Further discussions covered two specific new general approaches to protecting location privacy: context-dependent privacy, and location uncertainty as a means to protect privacy. In each of these areas, new ideas were developed; further information is given in the working group reports.

The seminar emphasized discussion of issues and collaborative development of solutions – the majority of the time was divided between working group breakout sessions followed by report-back and general discussion sessions. While the working group reports were written by subgroups, the contents reflect discussions involving all 22 participants of the seminar.

The seminar concluded that there are numerous challenges to be addressed in mobility data mining and privacy. These challenges require investigation on both technical and policy levels. Of particu-

lar importance is educating stakeholders from various communities on the issues and potential solutions.



Fig. 4.9
Steven Hautemaniere – Figures Libres. Part of the Dagstuhl art collection and donated by Angelika Mueller-von Brochowski and participants of the Dagstuhl Seminars 05311 and 05421.

4.33 Verifying Reliability

Organizers: Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda
Seminar No. 12341

Date: 19.–24. August, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.8.54

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Görschwin Fey, Masahiro Fujita, Natasa Miskov-Zivanov, Kaushik Roy, and Matteo Sonza Reorda



Participants: Jacob A. Abraham, Robert Aitken, Eli Arbel, Bernd Becker, Shawn Blanton, Cecile Braunstein, Mehdi Dehbashi, Giuseppe Di Guglielmo, Rolf Drechsler, Görschwin Fey, Masahiro Fujita, Carsten Gebauer, Jie Han, John P. Hayes, Bodo Hoppe, Ravishankar K. Iyer, Seiji Kajihara, Yoshiki Kinoshita, Wenchao Li, Igor L. Markov, Yusuke Matsunaga, Subhasish Mitra, Michael Orshansky, Laurence Pierre, Iliia Polian, Anand Raghunathan, Sudhakar M. Reddy, Kaushik Roy, Sachin Sapatnekar, Yasuo Sato, Matthias Sauer, Ulf Schlichtmann, Marcela Simková, Adit Singh, Matteo Sonza Reorda, Mehdi B. Tahoori, Massimo Violante, Tomohiro Yoneda

■ Introduction

Moore's law predicted the ever increasing computing power of the past decades from an economic perspective based on doubling the number of elements in a circuit about every two years. Moreover, Moore's law is expected to continue for another 10–20 years. On the physical level this integration is enabled by continuously shrinking feature sizes of basic components. But for future technology nodes reliability problems triggered by different sources are expected to increase rapidly. Process variations in the production process are one issue. While production processes become more accurate considering absolute measures, the relative inaccuracy compared to the component's size is increasing. One consequence are transistors with a wide range of threshold levels resulting in slightly faster or slower operating logic circuitry (both die-to die and within die). This may result for example in delay errors under certain operating conditions of a device. Increasing sensitivity to the omnipresent environmental radiation is another issue. In the past some errors induced by radiation have been observed infrequently while systems in space missions are already specified to be radiation resistant. Shrinking feature sizes result in sensitivity to radiation with lower energy causing more radiation induced events like Single Event Upsets (SEUs) even on sea level. Such effects are summarized as transient faults resulting in soft errors (as opposed to permanent faults resulting in a change of the functionality due to a modification of the physical structure). Consequently, approaches to design reliable circuits tolerating such transient faults without causing soft errors have been proposed. These design approaches to mitigate soft errors comprise all levels of design abstraction from the system specification down to the layout. Examples for these approaches are, e.g., fault tolerant algorithms and operating systems, fault tolerant processors, self-calibrating architectures, block level redundancy and error checking, synthesis approaches on the gate level, or hardening techniques on the layout level. In practical systems typically multiple mitigation techniques are implemented to guarantee reliability

across the full system stack. Functional verification has been and still is a challenge in current designs containing up to hundreds millions of transistors. Mature techniques for the formal verification and the dynamic verification of large systems exist. Research in verification is ongoing to match the rapid increase of the size of the systems. The verification of reliability is an interdisciplinary topic involving at least testing technology, verification methodology, and design experience. This makes the verification of reliable implementations an even harder problem. The testing community provides underlying models for transient faults to understand the effects at the functional and eventually at the system level. Using these models, the verification community designs efficient analysis tools and verification techniques to handle large systems. As in standard verification of large circuits a concerted action of formal methods, semi-formal techniques and simulation-based validation will be required. Still knowledge from the design community is required, to further speed up the verification task. Understanding the implemented approach to reliability on the application level and the system level is required to achieve a high degree of automation in the verification task.

■ Organization

The seminar was organized in short slots for talks followed by extensive discussions. A panel discussion in the afternoon summarized each day and focused on further questions. Each day was devoted to a special topic:

- Design – Techniques to ensure reliability by design.
- Fault models – Different types of fault models are required depending on the abstraction level and the type of design considered.
- Metrics – Measuring reliability requires some kinds of metrics. These metrics can be defined with respect to the fault models. But they should also reflect potential inaccuracies.

- Engines – Different types of engines are used in Electronic Design Automation (EDA) for circuits and systems.

■ Results

Documenting the results of intensive discussions in a compact manner is difficult. However, some results can be formulated in crisp statements. Approximate computing is a powerful technique for reliable design where the applications permit inaccuracy of operations up to a certain extent. Computing considering statistical nature of devices may be able to produce very accurate results, but providing compatible computing fabric at acceptable costs is a challenge. No single fault model will cover all aspects of reliability. In particular, fault models must be adapted to the application domain, the level of criticality and the step in the design process that is being considered. Appropriate metrics will then be applied to bridge gaps, e.g., between different levels of abstraction. An orchestration of reasoning engines ranging from formal techniques to simulation and emulation will always be required to gather data required for the different metrics. Design for Reliability will always affect all levels of abstraction. Only by concerted effort the same performance gains can be expected that we have seen in the past 50 years.

As a follow-up of the Dagstuhl Seminar, an Embedded Tutorial was successfully proposed for the DATE conference 2013. The Embedded Tutorial's title is "Reliability Analysis Reloaded: How Will We Survive?" and will include two presentations given by participants of the seminar or colleagues belonging to the research group of a participant.

4.34 Engineering Multiagent Systems

Organizers: Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke
Seminar No. 12342

Date: 19.–24. August, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.8.74

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Jürgen Dix, Koen V. Hindriks, Brian Logan, and Wayne Wobcke



Participants: Natasha Alechina, Jeremy Baxter, Michal Bida, Olivier Boissier, Rafael H. Bordini, Lars Braubach, Paolo Busetta, Rem Collier, Stephen Cranefield, Mehdi Dastani, Louise Dennis, Virginia Dignum, Jürgen Dix, Jorge J. Gomez-Sanz, Christian Guttman, Axel Heßler, Koen V. Hindriks, Tom Holvoet, Jomi Hübner, Yves Lespérance, Brian Logan, John-Jules Ch. Meyer, Berndt Müller, Jörg P. Müller, Alexander Pokahr, Alessandro Ricci, Andrea Santi, Federico Schlesinger, Amal El Fallah Seghrouchni, Maarten Sierhuis, Marija Slavkovic, Bas J. G. Testerink, Birna van Riemsdijk, Jørgen Villadsen, Michael Winikoff, Cees Witteveen, Wayne Wobcke

In 1993, Yoav Shoham's paper on agent-oriented programming was published in the *Artificial Intelligence Journal*. Shoham's ideas, and the work on agent-oriented programming it inspired, has had a profound impact on the field of multiagent systems, as evidenced by Shoham's paper receiving a 2011 IFAAMAS Influential Paper Award recognising seminal work in the field. Agent-oriented programming offers a natural approach to the development of complex systems in dynamic environments, and technology to support the development of agents and multiagent systems is beginning to play a more important role in today's software development at an industrial level.

Since Shoham's initial work, a range of platforms that support agent orientation have become available, and considerable experience has been gained with these platforms. Some key issues have also emerged from this work, however. First, given the plethora of systems and approaches that have become available in the field for developing multiagent systems, it is no longer clear which of these technologies is most appropriate for developing a particular application or what the distinctive benefits of various approaches are. It is especially important for practitioners to understand the benefits resulting from a particular choice of technology, when and how to apply it, and to develop standards that support the application of agent technology. Secondly, the very different style of agent-oriented programming potentially hampers the uptake of agent development tools and methods. To successfully apply the agent-oriented paradigm and to support the implementation and testing phases of agent-oriented development it is therefore very important to establish best practices and evaluate lessons learned from applying the technology in practice.

The aim of this seminar was to bring together researchers from both academia and industry to identify the potential for and facilitate convergence towards standards for agent technology. The seminar was very relevant for industrial research. The seminar meetings were meant to enable interaction, cross-fertilisation, and mutual feed-

back among researchers and practitioners from the different, but related areas, and provide the opportunity to discuss diverse views and research findings. The interaction in a Dagstuhl seminar was considered to be ideal for establishing common ground for defining standards, identifying best practices, and developing approaches to applying agent technology to the large scale, realistic scenarios found in industry. The aim of the discussions that were planned was therefore to establish a future research agenda, i.e. a *roadmap*, based on an evaluation of current state-of-the-art of agent-oriented programming languages, tools and techniques that are particularly important for large scale industrial applications.

The seminar took place August 19–24, 2012, with 37 participants from 15 countries. The programme included presentations by the participants and group discussions. Presentations were about 30 minutes long, including questions. We specifically asked participants not to present current research (their next conference paper), but rather asked for what should be considered the next step in their research area.

Participants were encouraged to use their presentations to provide input for discussion about the roadmap. They should show their perspectives and discuss what they think should be on the research agenda, try to explain why, and what it is they think this community should be aiming for. The group discussions took place in the afternoon, after the coffee break until 6pm. We put together four groups of 8–10 members, each headed by one discussion leader. The results of each working group were then presented to all participants before dinner. The seminar concluded with a general discussion on Friday morning and a wrap-up.

We identified the following important outcomes of the seminar.

MAS: Understanding of the uptake of multiagent systems technology in industry is seriously hampered by the current situation concerning paper acceptance at scientific conferences and workshops: While new theoretical approaches easily find their way into these events, papers about serious implementations

that scale up and put theoretical concepts to work are often considered not innovative enough and are thus not considered appropriate as scientific papers. We need a forum to publish such papers in order to generate research on the transfer of agent technology to industry.

Merger: During the seminar, eight out of 12 steering committee members of three important workshops in the area of agent systems development (ProMAS, DALT, and AOSE) met to discuss the possibility of merging the workshops. Based on the discussions at the seminar, it was generally agreed that greater focus is needed, and a single venue to present work in the field would be desirable. The workshop steering committees therefore decided (during the seminar) to merge ProMAS, DALT and AOSE to form a new workshop *Engineering Multiagent Systems*. 2012 will therefore be the last year in which the workshops will be held separately: They will be replaced by the new EMAS workshop at next year's AAMAS.

Roadmap: The organisers agreed to start a draft on the roadmap, based on the results of the group sessions. We plan to include the group leaders to produce a first draft, discuss it with the participants and afterwards, to finalise it.

4.35 Interaction Beyond the Desktop

Organizers: Alan Dix, James D. Hollan, Albrecht Schmidt, and Jürgen Steimle
Seminar No. 12351

Date: 26.–31. August, 2012 | Dagstuhl Seminar



Participants: Michel Beaudouin-Lafon, Susanne Boll, Alvaro Cassinelli, Alan Dix, Tanja Döring, Geraldine Fitzpatrick, Steve Gill, Jonna Häkkinen, Niels Henze, Uta Hinrichs, James D. Hollan, Clare Hooper, Eva Hornecker, Jochen Huber, Johann Habakuk Israel, Hans-Christian Jetter, Thorsten Karrer, Gudrun Klinker, Antonio Krüger, Wendy E. Mackay, Rainer Malaka, Ann Morrison, Max Mühlhäuser, Jörg Müller, Simon Olberding, Joseph A. Paradiso, Bastian Pfleging, Aaron Quigley, Harald Reiterer, Bernard Robben, Michael Rohs, Enrico Rukzio, Alireza Sahami, Christian Sandor, Arvind Satyanarayan, Bernt Schiele, Albrecht Schmidt, Johannes Schöning, Jürgen Steimle, Aurélien Tabard, Kaisa Väänänen-Vainio-Mattila, Volker Wulf

(No documentation available)

4.36 Information Flow and Its Applications

Organizers: Samson Abramsky, Jean Krivine, and Michael W. Mislove
Seminar No. 12352

Date: 26.–31. August, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.8.99

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Samson Abramsky, Jean Krivine, and Michael W. Mislove

Participants: Samson Abramsky, Marcus Appleby, David Balduzzi, Peter Bierhorst, Robert J. Bonneau, Gunnar Carlson, Eric Deeds, Ross Duncan, Jerome Feret, Tobias Fritz, Richard Gill, Jonathan Hayman, Peter Hines, Ricardo Honorato-Zimmer, Jean Krivine, Shane Mansfield, Michael W. Mislove, Catuscia Palamidessi, Prakash Panangaden, Mehrnoosh Sadrzadeh, Sandro Stucki, Baltasar Trancón y Widemann, Viktor Winschel, Glynn Winskel



The seminar “Information Flow and Its Applications” that took place in Schloss Dagstuhl in August 2012, has been the latest in a series of meetings concerning information flow that began with the 2008 Clifford Lectures by Samson Abramsky at Tulane University, and continued with two further meetings on informatic phenomena at Tulane, as well as a previous Dagstuhl seminar on “The Semantics of Information”³. The seminar “Information Flow and Its Applications” brought together mathematicians, computer scientists, physicists and researchers from related disciplines such as computational biology who are working on problems concerning information and information flow.

The seminar gathered 21 participants in addition to the 3 organizers, in the studios but cosy atmosphere of Schloss Dagstuhl. Armed with slides and chalks, each speaker described in terms as simple as can be, the questions and problems they were trying to solve, which, as the title of the seminar suggests, had all in common the issue of the representation and analysis of information flows.

The hypothesis underlying the organization of the seminar was the following: information flows leave on substrates which transport and transform data along time and space. From the modeling, analysis or simulation of these substrates will emerge unifying techniques or concepts. It is understood that such substrate can be *artificial*, for instance in the case of an electrical circuit, or *natural*, as in the complex signaling pathways that govern cellular fate. Moreover, information may be treated by systems in a *designed* manner, for instance a computer that processes its inputs according to a determined program, or be the result of *evolution*, like the internet which is a perfect example of a system that carries and processes information in spite of the absence of a pre-existing specification.

Although traditionally information processing is studied by distinct communities, scattered along the *Artificial-Natural* and *Desi-*

gned-Evolved axes, it is noteworthy that this separation is, to some extent, a historical artifact in the sense that artificial systems may be the fruit of evolution (as the internet) while natural ones may be used in a purely specified manner (as in synthetic biology). It is therefore natural to expect that tools and techniques developed in one field may be also relevant to others.

Another unifying scheme of the seminar was the emphasis on the use of formal languages in the representation of information flows. Indeed once a “real world” computing system, such as the cell or a quantum circuit, is abstracted as a formal programming language, one may then start to apply techniques imported from theoretical computer science. In the study of evolved systems, these techniques may be used to *extract* a specification of what is being observed, while in the context of systems where a specification is *a priori* at disposal, one may use these techniques to *verify* that the way information is processed conforms to the expectation.

Over the 4 days of talks, which gave rise to feedback that went beyond the expectation of the organizers, the participants of the seminar “Information Flow and its Applications” have had the opportunity to listen to talks ranging from Systems Biology to Theoretical Physics, from Quantum Computing to the study of Ecological systems. As organizers, we believe that the original guess that Information Flow should be a topic of its own was largely a good one.

³ <http://www.dagstuhl.de/10232>

4.37 Information-centric networking – Ready for the real world?

Organizers: Ali Ghodsi, Börje Ohlman, and Ignacio Solis

Seminar No. 12361

Date: 02.–09. August, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.1

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Ali Ghodsi, Börje Ohlman, Jörg Ott, Ignacio Solis, and Matthias Wählisch



Participants: Bengt Ahlgren, Somaya Arianfar, Carsten Bormann, Marcus Brunner, Ken Calvert, Giovanna Carofiglio, Antonio Carzaniga, Patrick Crowley, Elwyn Brian Davies, Lars Eggert, Anders Eriksson, Christian Esteve Rothenberg, Kevin Fall, Xiaoming Fu, Volker Hilt, Pan Hui, Van Jacobson, Myeong-Wuk Jang, Holger Karl, Gunnar Karlsson, Vikas Kawadia, Dirk Kutscher, Anders Lindgren, Luca Muscariello, Ashok Narayanan, Börje Ohlman, George Parisis, Pasi Sarolahti, Thomas C. Schmidt, Ignacio Solis, Matthias Wählisch, George Xylomenos, Eiko Yoneki, Lixia Zhang

Information-centric networking (ICN) defines a communication paradigm that recognizes the dominant usage of the Internet as a substrate to disseminate and access content of all kinds: from traditional web pages to online social networks to file distribution to live and on-demand media feeds. With ICN, the focus shifts from the location at which a content object is stored (typically some server) to the object itself so that scale, efficiency, and robustness of content publication and retrieval can be improved beyond what current Content Distribution Networks (CDNs) can deliver.

Diverse instances of ICN networking architectures were developed, including CCN/NDN, NetInf, DONA, and LIPSIN, among others, and see experimentation at different scale in both academia and industry. The fundamental concepts of ICN have gained popularity in the research community and have been taken up by several research activities that are addressing the topic from different angles.

Numerous research problems remain open, some of which (such as naming content) may find different (optimal) solutions in different deployments while others are more fundamental in nature and could affect the performance of all deployments. The latter include the performance benefits achievable through (cooperative) caching and caching at different points in the network, parallel content retrieval from multiple sources, and tradeoffs between native network layer and overlay-based ICNs. This second Dagstuhl Seminar on information-centric networking is intended to operate as a catalyst for these activities and provide a forum for discussing a selected subset of important research topics that have been identified so far. It will bring together researchers from different ICN backgrounds to discuss fundamentals that matter across the various platforms with the meta goal of identifying obstacles to be overcome, solutions, and paths towards real-world deployments.

In this seminar, we discussed the following core topics: (1) ICN applications and services, (2) ICN performance and comparison of alternative technologies, (3) business, legal, and deployment aspects.



Fig. 4.10
Hans Ortner – New York 1. Part of the Dagstuhl art collection and donated by Günther Hotz.

4.38 The Multilingual Semantic Web

Organizers: Paul Buitelaar, Key-Sun Choi, Philipp Cimiano, and Eduard H. Hovy
Seminar No. 12362

Date: 02.–09. September, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.15

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Paul Buitelaar, Key-Sun Choi, Philipp Cimiano, and Eduard H. Hovy



Participants: Guadalupe Aguado de Cea, Dimitra Anastasiou, Pushpak Bhattacharyya, Gerhard Budin, Paul Buitelaar, Nicoletta Calzolari Zamorani, Manuel Tomas Carrasco Benitez, Christian Chiarcos, Key-Sun Choi, Philipp Cimiano, Ernesto William De Luca, Gerard de Melo, Thierry Declerck, Bo Fu, Asuncion Gomez-Perez, Jorge Gracia, Marko Grobelnik, Iryna Gurevych, Sebastian Hellmann, Graeme Hirst, Chu-Ren Huang, Nancy Ide, Antoine Isaac, Christian Lieske, John McCrae, Elena Montiel-Ponsoda, Roberto Navigli, Sergei Nirenburg, Laurette Pretorius, Aarne Ranta, Kimmo Rossi, Felix Sasaki, Gabriele Sauberer, Hans Uszkoreit, Josef van Genabith, Jeroen van Grondelle, Daniel Vila-Suero, Martin Volk

The amount of Internet users speaking native languages other than English has seen a substantial growth in recent years. Statistics from 2010 in fact show that the number of non-English Internet users is almost three times the number of English-speaking users (1430 million vs. 536 million users). As a consequence, the Web is turning more and more into a truly multilingual platform in which speakers and organizations from different languages and cultural backgrounds collaborate, consuming and producing information at a scale without precedent. Originally conceived by Tim Berners-Lee et al. [1] as “an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation”, the Semantic Web has seen an impressive growth in recent years in terms of the amount of data published on the Web using the RDF and OWL data models. The kind of data published nowadays on the Semantic Web or Linked Open Data (LOD) cloud is mainly of a factual nature and thus represents a basic body of knowledge that is accessible to mankind as a basis for informed decision-making. The creation of a level playing field in which citizens from all countries have access to the same information and have comparable opportunities to contribute to that information is a crucial goal to achieve. Such a level playing field will also reduce information hegemonies and biases, increasing diversity of opinion. However, the semantic vocabularies used to publish factual data in the Semantic Web are mainly English, which creates a strong bias towards the English language and culture. As in the traditional Web, language represents an important barrier for information access as it is not straightforward to access information produced in a foreign language. A big challenge for the Semantic Web therefore is to develop architectures, frameworks and systems that can help in overcoming language and national barriers, facilitating the access to information originally produced for a different culture and language. An additional problem is that most of the information on the Web stems from a small set of countries where majority languages are spoken. This leads to a situation in which

the public discourse is mainly driven and shaped by contributions from those countries where these majority languages are spoken. The Semantic Web vision bears an excellent potential to create a level playing field for users with different cultural backgrounds, native languages and originating from different geo-political environments. The reason is that the information available on the Semantic Web is expressed in a language-independent fashion and thus bears the potential to be accessible to speakers of different languages if the right mediation mechanisms are in place. However, so far the relation between multilingualism and the Semantic Web has not received enough attention in the research community. Exploring and advancing the state-of-the-art in information access to the Semantic Web across languages is the goal of the seminar proposed here. A Semantic Web in which information can be accessed across language and national barriers has important social, political and economic implications:

- it would enable access to data in other languages and thus provide support for direct comparisons (e.g. of public spending), thus creating an environment where citizens feel well-informed and contributing to increasing their trust and participation in democratic processes as well as strengthening democracy and trust in government and public administration
- it would facilitate the synchronization and comparison of information and views expressed in different languages, thus contributing to opinion forming processes free of any biases or mainstream effects
- it would foster higher information transparency; the exchange of many data items is limited due to national boundaries and national idiosyncrasies, as it is e.g. the case with financial data, the exchange of which is limited due to the availability of very different accounting procedures and reporting standards. Creating an ecosystem in which financial information can be integrated across countries can contribute to a higher transparency of financial information, global cash flow and investments.

Vision, Goals and Topic: The vision underlying the proposed workshop is the creation of a Semantic Web in which all languages have the same status, every user can perform searches in their own language, and information can be contrasted, compared and integrated across languages. As a main topic for the seminar, we intend to discuss in how far the Semantic Web can be extended – from an infrastructural and conceptual point of view – in order to support access across languages. This will lead us to the discussion of two main questions:

- Ontological vocabularies that are available and used in the Semantic web cover a broad number of domains and topics to varying degrees of detail and granularity. For one thing we will discuss in how far these vocabularies can indeed be seen as an interlingua (language-independent) representation. This includes the question how, building on such an interlingual representation, the Semantic Web can indeed support access to semantic data across languages. This discussion will extend to the question which approaches are suitable to translate the user's information needs, expressed in natural language, into such a language-independent representation.
- For another thing, we will discuss how the multilingual Semantic Web can be constructed by publication and linking of available multilingual lexical resources following the Linked Data paradigms. In this context, we will also discuss how natural language processing tools can benefit from such a linked ecosystem of lexico-semantic background knowledge.

Other topics that we anticipated would be discussed at the seminar include the following:

- models for the integration of linguistic information with ontologies, i.e., models for multilingualism in knowledge representation, in particular OWL and RDF(S)
- collaborative design of ontologies across languages and cultures
- multilingual ontology alignment
- multilingual and cross-lingual aspects of semantic search and querying of knowledge repositories
- cross-lingual question answering over Linked Data
- architectures and infrastructure for a truly Multilingual Semantic Web
- localization of ontologies to multiple languages
- automatic integration and adaptation of (multilingual) lexicons with ontologies
- multi- and cross-lingual ontology-based information extraction and ontology population
- multilingualism and linked data (generation, querying, browsing, visualization and presentation)
- multilingual aspects of ontology verbalization
- ontology learning across languages
- NLP methods to construct the multilingual Semantic Web

■ Organization & Structure

The organizers asked participants to submit an abstract and to prepare a short presentation of about 10 minutes for the seminar.

The first day started with an introduction by the organizers, gi-

ving an overview of the main topics and goals of the seminar. Some guiding questions for the seminar as proposed by the organizers were the following:

- Can we exploit the LOD for NLP?
- Can we allow for multilingual access to the knowledge in the LOD?
- Can we regard the LOD as an interlingua?
- Can we apply Linked Data principles to the modelling of linguistic/lexical resources?
- How can we facilitate the localization of (semantic) web sites to multiple languages?

As technical and research challenges for the field in the next years, the organizers highlighted the following:

- Aggregating and summarizing content across languages
- Repurposing and verbalizing content in multiple languages
- Linking of information across languages
- Detection of inconsistent views across languages
- Translation of “objects” that have a context and are produced within some workflow
- Large-scale and robust text analysis in multiple languages
- Personalized and contextualized Interpretation of NL [2]
- Cross-lingual/cultural reconciliation of conceptualizations

Every day, between 10:30 and 12:00, a panel took place in which attendees of the seminar had 10 minutes to present their view on the main challenges in the field, answering to the following questions in particular:

1. What are in your view the most important challenges/ barriers/ problems and pressing needs with respect to the multilingual access to the Semantic Web?
2. Why does the problem matter in practice? Which industry sectors or domains are concerned with the problem?
3. Which figures are suited to quantify the magnitude or severity of the problem?
4. Why do current solutions fail short?
5. What insights do we need in order to reach a principled solution? What could a principled solution look like?
6. How can standardization (e.g. by the W3C) contribute?

After each panel the organizers attempted to group participants into teams around a certain topic. The groups worked together on the topic in the afternoons between 13:30 and 15:30. They were supposed to wrap-up their discussion and come up with a summary of their discussion until 17:00. These summaries were then presented in a plenary session to all the participants from Tuesday to Friday between 9:00 and 10:30.

Every day between 17:00 and 18:00 (just before dinner), we had an invited talk or special activity. On the first day, Kimmo Rossi from the European Commission shared his perspective on the challenges in our field. On the second day, there was a non-academic slot: First Jeroen van Grondelle showcased an industrial application of semantic, multilingual technologies; next, Christian Lieske and Felix Sasaki discussed perception and reality of the multilingual Semantic Web. On the fourth day we organized a demo session, giving participants the opportunities to give a hands-on look at their tools.

■ References

- 1 T. Berners-Lee, J. Hendler, O. Lassila, et al. The Semantic Web. *Scientific American*, 284(5):28–37, 2001.
- 2 G. Hirst. The future of text-meaning in computational linguistics. In *Proc. of 11th Int'l Conf. on Text, Speech and Dialogue*, pp. 3–11, Springer, 2008.

4.39 Software Defined Networking

Organizers: Pan Hui and Teemu Koponen

Seminar No. 12363

Date: 05.– 08. September, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.95

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Pan Hui and Teemu Koponen



Participants: Bengt Ahlgren, Marcus Brunner, Frank Dürr, Lars Eggert, Christian Esteve Rothenberg, Peter Feil, Anja Feldmann, Nate Foster, Howard Green, Pan Hui, Raimo Kantola, James Kempf, Teemu Koponen, Dirk Kutscher, Daniel Levin, Anders Lindgren, David Meyer, Toby Moncaster, Andrew W. Moore, Gerd Pflueger, Jarno Rajahalme, Wolfgang Riedel, Sasu Tarkoma, Fernando Manuel Valente Ramos, Cedric Westphal

Software Defined Networks (SDN) is seen as the most promising solution to resolve the challenges in realizing sophisticated network control. SDN builds its promise on the separation of the network control functions from the network switching elements. By moving the control plane out from the network elements into stand-alone servers, the switching elements can remain simple, general-purpose, and cost-effective and at the same time the control plane can rely on design principles of distributed systems in its implementation instead of being confined to distributed routing protocols.

The purpose of the seminar was to look at the current developments in this quickly evolving problem domain and identify future research challenges. The seminar brought together researchers with different domains and backgrounds. Given the high level of interest in SDN from industry, the organizers also invited many participants from companies working with SDN related networking products and services. This mix of people resulted in fruitful discussions and interesting information exchange. The structure of the seminar took advantage of these different backgrounds by focusing on themed talks and group discussions.

■ Organization of the Seminar

Software-Defined Networking (SDN) continues to remain relevant both for the industry and academia and indeed this was very much reflected in the backgrounds of the seminar participants; the seminar had a balanced mix of representatives both from industry and academia.

These two very active communities, industry and academia, are pursuing SDN with different mind-sets, different solutions and different implications in mind, however. The organizers felt that the interactions had been clearly insufficient in the past: practical challenges in SDN continue to remain little known in the academia whereas the industry often remains unaware of the recent useful developments in research. To this end, the two and half day seminar was explicitly structured around this observation; the goal was to allow

for fruitful interactions between the industry and academia to maximize the exchange of ideas, challenges and lessons learnt between these two communities.

The seminar discussions and talks were structured around three themes:

1. Status updates. From the very definition to the ongoing standardization work, SDN is still evolving. In these talks and discussions, we dived into the ongoing work at ONF as well as the perceived hard problems to be solved.
2. Industry use cases. In this theme the focus was on exposing the academia to the practical use cases on which industry is working.
3. Implementation. The third theme dived into the details and exposed the seminar participants to both the practical implementation issues faced as well as more theoretical observations about the system design.

For the status updates the seminar had the following talks at the first day. The talks were fairly short so enough discussion could be had between the talks:

- Teemu Koponen: Evolving SDN
- Peter Feil: ONF update
- David Meyer: Hard problems in OF/SDN
- Dirk Kutscher: Northbound interfaces

The discussions after (and during) the talks also bootstrapped the evening and its group discussions about the definition of the SDN and its use cases.

The second day started with the industry use cases.

- Peter Feil: Deutsche Telekom and SDN
- James Kempf: SDN: Definition and Use Cases
- Teemu Koponen: Network virtualization
- Cedric Westphal: SDN for content management/network-based CDN emulation/transparent caching

The rest of the day was dedicated for the implementation theme and

a set of short talks were given again to spark the discussion later in the evening about the implementation aspects.

- Dan Levin: State distribution trade-offs in SDN
- Nate Foster: Frenetic
- Toby Moncaster: SDN, can we (IP)FIX it?
- Andrew Moore: S/FPGA/NetFPGA
- Jarno Rajahalme: Issues in routing and tunneling in OF and OVS
- Wolfgang Riedel: Alignment of Storage, Compute and Networking
- Anders Lindgren: Use cases of SDN in information centric mobile networks

The third day was again about the use cases but this time from the academic participants. The following short talks were given with discussions between the talks:

- Christian Rothenberg: RouteFlow
- Fernando Ramos: Secure, trustworthy, resilient SDNs
- Raimo Kantola: Customer Edge Switching
- Frank Dürr: Supporting Communication Middleware with Software-Defined Networking

■ Outcome of the Seminar

The seminar was well received by the participants. Among the participants there were also organizers of future SDN workshops (IRTF SDN and DIMACS SDN) who signaled the intent of building their workshops around the similar discussion-oriented structure preferred at Dagstuhl.

4.40 Machine Learning Methods for Computer Security

Organizers: Anthony D. Joseph, Pavel Laskov, Fabio Roli, J. Doug Tygar, and Blaine Nelson
Seminar No. 12371

Date: 09.–14. September, 2012 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.2.9.109

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Anthony D. Joseph, Pavel Laskov, Blaine Nelson, Fabio Roli, and J. Doug Tygar



Participants: Battista Biggio, Christian Bockermann, Michael Brückner, Alvaro Cárdenas Mora, Christos Dimitrakakis, Felix C. Freiling, Giorgio Fumera, Giorgio Giacinto, Rachel Greenstadt, Anthony D. Joseph, Robert Krawczyk, Pavel Laskov, Richard P. Lippmann, Daniel Lowd, Aikaterini Mitrokotsa, Sasa Mrdovic, Blaine Nelson, Patrick Pak Kei Chan, Massimiliano Raciti, Nathan Ratliff, Konrad Rieck, Fabio Roli, Benjamin I. P. Rubinstein, Tobias Scheffer, Galina Schwartz, Nedim Srndic, Radu State, Doug Tygar, Viviane Zwanger

Arising organically from a variety of independent research projects in both computer security and machine learning, the topic of machine learning methods for computer security is emerging as a major direction of research that offers new challenges to both communities. Learning approaches are particularly advantageous for security applications designed to counter sophisticated and evolving adversaries because they are designed to cope with large data tasks that are too complex for hand-crafted solutions or need to dynamically evolve. However, in adversarial settings, the assets of learning can potentially be subverted by malicious manipulation of the learner's environment. This exposes applications that use learning techniques to a new type of security vulnerability in which an adversary can adapt to counter learning-based methods. Thus, unlike most application domains, computer security applications present a unique data domain that requires careful consideration of its adversarial nature to provide adequate learning-based solutions—a challenge requiring novel learning methods and domain-specific application design and analysis. The Perspectives Workshop, “Machine Learning Methods for Computer Security”, brought together prominent researchers from the computer security and machine learning communities interested in furthering the state-of-the-art for this fusion research to discuss open problems, foster new research directions, and promote further collaboration between the two communities.

This workshop focused on tasks in three main topics: the role of learning in computer security applications, the paradigm of secure learning, and the future applications for secure learning. In the first group, participants discussed the current usage of learning approaches by security practitioners. The second group focused on the current approaches and challenges for learning in security-sensitive adversarial domains. Finally, the third group sought to identify future application domains, which would benefit from secure learning technologies.

Within this emerging field several recurrent themes arose throughout the workshop. A major concern that was discussed

throughout the workshop was an uneasiness with machine learning and a reluctance to use learning within security applications and, to address this problem, participants identified the need for learning methods to provide better transparency, interpretability, and trust. Further, many workshop attendees raised the issue of how human operators could be incorporated into the learning process to guide it, interpret its results, and prevent unintended consequences, thus reinforcing the need for transparency and interpretability of these methods. On the learning side, researchers discussed how an adversary should be properly incorporated into a learning framework and how the algorithms can be designed in a game-theoretic manner to provide security guarantees. Finally, participants also identified the need for a proper characterization of a security objective for learning and for benchmarks for assessing an algorithm's security.



Fig. 4.11

Boy Müller – ohne Titel. Part of the Dagstuhl art collection and donated by participants of the Dagstuhl Seminars 10181 and 10421, and further people.

4.41 Biological Data Visualization

Organizers: Carsten Görg, Lawrence Hunter, Jessie Kennedy, Seán O'Donoghue, and Jarke J. van Wijk

Seminar No. 12372

Date: 09.–14. September, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.131

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Carsten Görg, Lawrence Hunter, Jessie Kennedy, Seán O'Donoghue, and Jarke J. van Wijk



Participants: Jan Aerts, Sheelagh Carpendale, Min Chen, Kasper Dinkla, Daniel Evanko, Jean-Fred Fontaine, Tom Freeman, Nils Gehlenborg, Carsten Görg, Martin Graham, Julian Heinrich, Lawrence Hunter, Graham Johnson, Jessie Kennedy, Andreas Kerren, Karsten Klein, Oliver Kohlbacher, Martin Krzywinski, Michael Lappe, Raghu Machiraju, Miriah Meyer, Torsten Möller, Cydney Nielsen, Kay Nieselt, Sean O'Donoghue, Arthur J. Olson, James Procter, Mark Ragan, Jos B.T.M. Roerdink, Francis Rowland, Andrea Schafferhans, Falk Schreiber, Svenja Simon, Christian Stolte, Margaret-Anne Storey, Hendrik Strobel, Jarke J. Van Wijk, Corinna Vehlow, Matthew O. Ward, Michel A. Westenberg, Bang Wong

■ Introduction and Motivation

Biology is rapidly evolving into a 'big data' science, and as a consequence there is an urgent and growing need to improve the methods and tools used for gaining insight and understanding from biological data. Over the last two decades, the emerging fields of computational biology and bioinformatics have led to significant advances primarily in automated data analysis. Today, however, biologists increasingly deal with large, complex datasets (e.g., 'omics' data) where it is not known in advance what they are looking for and thus, automated analyses alone cannot solve their problems. Interactive visualizations that can facilitate exploratory data analysis and support biologists in creating new hypotheses lend themselves to complement automated analyses. Bioinformaticians already have built a variety of tools for visualizing different types of biological data and those tools are widely used in the community. So far, most bio-related visualization research has been conducted by people outside of the visualization community, people who have learned about visualization but are often not aware of research in the visualization community. Consequently, the current tools do not embody the latest advancements in design, usability, visualization principles, and evaluation.

One main goal of this first Dagstuhl Seminar on Biological Data Visualization was to bring together the users (biologists), current visualization tool builders (bioinformaticians), and visualization researchers to survey the state-of-the-art of the current tools and define a research agenda for systematically developing the next generation of tools for visualizing biological data. Only a close collaboration of the researchers from all three communities can create the synergies necessary to address the challenges in analyzing and visualizing large and complex biological datasets.

Topics discussed during the seminar included:

- Challenges in visualizing biological data. Biological data is very heterogeneous. It contains spatial data, graphs, tabular data, and textual data. Challenges are wide spread: open-ended da-

ta quantity, open-ended exploratory tasks, long-term analyses, rich analytics, heterogeneous data, usability and evaluation of tools.

- Design and visualization principles, research in human-centered design, usability, and evaluation of interactive data-analysis tools.
- Creating a common research agenda and a common understanding of the problem field of biological data visualization.
- Integration of multiple visualizations for different data types and tasks into one tool to support more complex analysis scenarios.
- Designing an infrastructure for next generation visualization tools.
- Establishing collaborations between computer scientists and biologists.

■ Participants and Program

41 researchers from 9 countries participated in this seminar. Many participants came from the US and from Germany, others came from Canada, Australia, and a number of other European countries. There was a good mix of researchers from the visualization, bioinformatics, and biology communities. About a third of the participants attended their first seminar at Dagstuhl.

The program was designed to facilitate in-depth discussions in small working groups. To get to know each other—the seminar brought together researchers from different communities—participants introduced themselves and their research interests with a 'personal ad' in the Monday morning session. This was a great way to set the tone for informal and engaging discussions during the seminar.

Previous to the seminar, the organizers collected interesting ideas and suggestions from the participants for possible topics for working groups. To allow participants to work on different topics and with different people, the topics and groups changed halfway

through the seminar. On Monday morning and Wednesday morning all participants discussed and refined the suggested topics and formed groups according to their interests. The groups (four on Monday/Tuesday and five on Wednesday/Thursday) worked in parallel on their topics and reported regularly on their progress. The work in the breakout groups was complemented by a discussion on the BioVis Community on Friday and a number of talks given throughout the seminar:

- *Seán I. O'Donoghue*: BioVis Introduction: A Practitioner's Viewpoint
- *Daniel Evanko*: Visualization on nature.com
- *Matt Ward*: Biovisualization Education: What Should Students Know?
- *Arthur J. Olson*: The Promise and Challenge of Tangible Molecular Interfaces
- *Martin Krzywinski*: visualization – communicating, clearly
- *Bang Wong*: Concepts gleaned from disparate communities

These talks, presented to all participants in the morning sessions and after the lunch breaks, intentionally touched on broad and high-level topics to make them more interesting to the diverse audience in the seminar.

■ Discussion and Outcome

Some of the working groups followed a classical design process [1, 2] to structure their collaborative work. They split their discussions into a *problem phase* and a *solution phase*. Both phases featured divergent and convergent stages: *discover* and *define* for the problem phase and *develop* and *deliver* for the solution phase. Francis Rowland, a seminar participant with expertise in user experience design, facilitated these discussions.

Figure 4.12 shows some artifacts produced by the Ontologies in Biological Data Visualization working group that followed this design process. The *Four C's* approach (left) is an example for the discover and design stages. The group broke down their topic into four aspects: *Components* (parts), *Characters* (people involved), *Challenges*, and *Characteristics* (features and behavior). The Four C's approach helped the group to provide a holistic view on the design problem and to better define the topic. The *Draw the Box* approach (right) is an example for the develop and deliver stages. Members of the group collaboratively imagined an end product of their work that would be sold in a box on a shelf and designed its package. This approach helped the group members to gather ideas, visualize the outcome, and focus on the most important features of the product.

The diverse outcomes from the nine working groups are summarized below.

Comparative Analysis of Heterogeneous Networks: The analysis of the transcriptome produces a large number of putatively disrupted transcripts, and prioritizing which disruptions are most likely to be meaningful (causal or diagnostic) is a time-consuming process. To guide their interpretation researchers create heterogeneous networks by integrating information from a wide variety of annotation databases. The working group investigated how the analysis of the transcriptome can be facilitated by interactive visualizations of transcriptome assemblies and proposed a method to infer the functional consequence of a transcript's disruption based on the local structure of the annotation networks. A tight coupling of network analysis algorithms and interactive visualizations, specifically designed to support these analysis tasks, could accelerate identification of important transcript alterations.

Sequence Data Visualization: Genome-associated data is growing at a fast rate and genome browsers are still the tool of choice for integrating and analyzing different types of data in one single representation. The working group analyzed the different challenges

of visualizing genome-associated data and separated them into two different dimensions: problems associated with rearrangements of the genomic coordinates and problems with the abundance of data at each genomic position. To address these problems, the group discussed and developed a number of possible solutions, including the development of a reference-free gene-centric approach, compressing tracks by aggregation or summarization, and using meta-data or data itself as a novel way for selecting tracks. These approaches can lay the foundation for the development of new visualization tools.

Bridging Structural & Systems Biology via DataVis: There exist several gaps between the field of structural biology, which has yielded detailed insight into the molecular machines of life, and the field of systems biology, which has evolved more recently in the wake of the genomics revolution, but separately from the advances of the more structural view of biology. The integration of both fields and their visualization tools could create new tool sets to enhance the exploration and understanding of biological systems. The working group analyzed and described the existing gaps and proposed seven strategies to facilitate collaboration and professional advancement in structural biology, systems biology, and data visualization.

Ontologies in Biological Data Visualization: Ontologies are graph-based knowledge representations in which nodes represent concepts and edges represent relationships between concepts. They are widely used in biology and biomedical research, for the most part as computational models, in computational analyses, and for text mining approaches. The working group examined the potential impact of ontologies on biological data visualization. The group identified challenges and opportunities from the perspectives of three different stakeholders: ontologists (who create and maintain ontologies), data curators (who use ontologies for annotation purposes), and data analysts (who use ontologies through applications to analyze experimental data). Identified challenges include the dynamic nature of ontologies, scalability, how to utilize the complex set of relationships expressed in ontologies, and how to make ontologies more useful for data analysis. Identified research opportunities include the visualization of ontologies themselves, automated generation of visualization using ontologies, and the visualization of ontological context to support search. The group submitted a Viewpoints article on Ontologies in Biological Data Visualization to the IEEE Computer Graphics & Applications journal.

A Framework for Effective Visualization Design: Visualizations are not only an important aspect of how scientists make sense of their data, but also how they communicate their findings. The techniques and guidelines that govern how to design effective visualizations, however, can be quite different whether the goal is to explore or to explain. Unfortunately, scientists are often not aware of the spectrum of considerations when creating visualizations. To help clarify this problem, the working group has developed a framework to reason about the spectrum and considerations to help scientists better match their visualization goals with appropriate design considerations.

Uncertainty Visualization: Uncertainty is common in all areas of science, and it poses a difficult problem for visualization research. Visualization of uncertainty has received much attention in the areas of scientific visualization and geographic visualization; however, it appears much less common in information visualization and in biological data visualization. The working group analyzed and described the sources of uncertainty and types of uncertainty specific to biology. Uncertainty visualization in networks was identified as an open issue, including uncertainty in the network topology and uncertainty in attributes on nodes, edges, and their interdependencies. The group started a survey of the literature on uncertainty visualization for biological data and proposed to construct a taxonomy of uncertainty visualization approaches, and investigate how



Fig. 4.12
Examples of design processes: the *Four C's* approach (left) and the *Draw the Box* approach (right)

they could be employed in the context of a collection of biological problems.

Evaluation: The working group identified two central problems with respect to the evaluation of tools for visualizing biological data: (1) How to motivate biologists to participate in evaluations? and (2) How to evaluate the tools? The answer to the first question was (simply) that biologists have to benefit from the evaluation to be motivated to participate, e.g. they might get a tool they can use to solve their problems. The second question was more complex and the working group discerned a number of dimensions, centered around what, why, when, where, and how. The discussion of these dimensions lead to the insight that there is a strong difference between approaches taken by designers working at a bio-institute and approaches taken by infovis researchers. Both approaches have merit, the challenge is to close the gap and combine them.

Multiscale Visualization: Biology involves data and models at a wide range of scales and researchers routinely examine phenomena and explore data at multiple scales. Visual representations of multi-scale datasets are powerful tools that can support data analysis and exploration, however, visualizing multi-scale datasets is challenging and not many approaches exist. The working group identified four common dimensions of biological multi-scale datasets: 3D space, time, data complexity (modality), and data volume (size). The group produced a short video to introduce each dimension independently in order to provide a quick and understandable view on the nature of the different scales and how they apply to biological data and exploration. Additionally, the group discussed in more detail a number of biological multi-scale data and models that can be visualized across multiple dimensions and introduced case studies to highlight issues like navigation, interaction, and human-computer interfaces. Carsten Görg presented a talk on the results

from this working group at the 2012 Rocky Mountain Bioinformatics Conference.

Infrastructure: The working group discussed needs from both a technical and community standpoint regarding the challenges involved in the analysis of biomedical data and mechanisms to facilitate interactions between visualization communities in computer science and biology. Eight key criteria were identified: interoperability, reusability, compatibility, references & benchmarks, middleware, vertical integration, scalability, and sustainability. The group developed a model for a community-maintained, biological visualization resource that would enable biological questions, task descriptions, sample datasets and existing tools for the problems to be disseminated to the computational visualization and biological research communities. Additionally, the group developed a detailed use-case based on the data and analysis pipelines of the cancer genome atlas that will allow technical aspects of the eight key criteria to be explored and practical solutions proposed.

Finally, based on feedback from the participants (from the seminar questionnaire as well as from personal communication with the organizers) another important outcome of the seminar was to establish collaborations between computer scientists and biologists. The academic cultures in biology and computer science, including publication models, are quite different. In addition, biologists have a different mindset than computer scientists: biologists often work in a detail-oriented manner whereas computer scientists often seek to generalize. Understanding each other's culture is important for successful collaborations and the Dagstuhl seminar provided a unique setting to meet enthusiastic people from different communities, have long group discussions with a focus on problem solving, and form synergies with researchers that have a different outlook and expertise.

References

- 1 D. Gray. *Gamestorming: A Playbook for Innovators, Rulebreakers, and Changemakers*. O'Reilly, 2010.
- 2 L. Hohmann. *Innovation Games: Breakthrough Products Through Collaborative Play: Creating Breakthrough Products and Services*. Addison Wesley, 2006.



Fig. 4.13

Aloys Ohlmann – Frau mit Katze. Part of the Dagstuhl art collection and donated by Wolfgang Wahlster.

4.42 Privacy-Oriented Cryptography

Organizers: Jan Camenisch, Mark Manulis, Gene Tsudik, and Rebecca Wright
Seminar No. 12381

Date: 16.–21. September, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.165

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Jan Camenisch, Mark Manulis, Gene Tsudik, and Rebecca N. Wright



Participants: Giuseppe Ateniese, Johannes Blömer, Nikita Borisov, Xavier Boyen, Jan Camenisch, Claude Castelluccia, Bruno Crispo, George Danezis, Emiliano De Cristofaro, Claudia Diaz, Roger Dingledine, Yevgeniy Dodis, Maria Dubovitskaya, Marc Fischlin, Cédric Fournet, Ian Goldberg, Dennis Hofheinz, Jean Pierre Hubaux, Aaron M. Johnson, Stefan Katzenbeisser, Dogan Kesdogan, Markulf Kohlweiss, Anja Lehmann, Anna Lysyanskaya, Mark Manulis, Gregory Neven, Bertram Poettering, Bartosz Przydatek, Kai Rannenberg, Jean-Pierre Seifert, Vitaly Shmatikov, Radu Sion, Douglas Stebila, Gene Tsudik, Markus Ullmann, Susanne Wetzel, Douglas Wikström, Rebecca Wright

The constantly increasing volume of electronic interactions and sensitive information disseminated online raises privacy concerns and motivates the need for efficient privacy-oriented techniques. The aim of our “Privacy-Oriented Cryptography” seminar was to bring together (mainly, but not only) researchers working in different domains of cryptography and privacy. Although non-cryptographic measures can, at times, aid privacy (e.g., statistical or ad hoc obfuscation techniques) — cryptography, via its mathematical mechanisms and formal concepts, helps obtain novel and efficient privacy-enhancing solutions, achieving concrete and measurable privacy guarantees.

Since privacy is a very broad area, being explored not only by security and cryptography experts, this seminar focused on two domains: *user privacy* and *data privacy*, for which the benefit from using cryptographic techniques is especially significant. Seminar participants presented and discussed many novel privacy-oriented cryptographic algorithms and protocols that admit various fields of deployment for protecting privacy in a broad range of applications, involving possibly huge amounts of data (e.g., cloud computing) and many different users (e.g. online communities). The seminar further addressed the emerging research direction of *provable privacy*, by discussing various mechanisms and techniques for defining concrete privacy goals and enabling their formal analysis.

The seminar brought together 39 of the leading scientists in the areas of (applied) cryptography and privacy. The participants came from all over the world, including the US (13 participants), Germany (8), Switzerland (6), Great Britain (5), Australia (1), Belgium (1), Canada (1), France (1), Italy (1), and Sweden (1).

The program contained 26 interactive presentations, each about 35–40 minutes and two panel discussions, with a free afternoon on Wednesday to offer time for social activities or for conducting collaborative research in smaller groups. The seminar ended on Friday after lunch to enable time for traveling. We asked participants prior to the seminar to suggest talks based on their most recent results.

Most presentations followed this suggestion and introduced new, sometimes even not yet submitted or still work-in-progress results. The first panel — “Privacy Models: UC or Not UC?” — discussed the advantages and disadvantages of existing cryptographic methods for formal specification and analysis of security and privacy guarantees. The second panel — “Privacy-Oriented Cryptography: Why is it not adopted more in practice?” — discussed challenges that arise in the practical deployment of existing privacy-oriented cryptographic solutions but also considered some success stories like Tor, a popular anonymous communications service, which is widely used in different parts of the world.

The nature of the seminar allowed experts and practitioners to air ideas and discuss preliminary concepts and work-in-progress results. This might have led to the exposure and subsequent exploration of new research directions that may offer both practical significance and intellectual challenge.

The organizers would like to thank all participants for accepting our invitations and attending the seminar, and for sharing their ideas and contributing to the interesting seminar program. We hope that discussions were fruitful and the opportunity to work face-to-face during the seminar helped to create impulses for exciting new research projects, paving the way for further progress and new discoveries in Privacy-Oriented Cryptography.

Finally, the organizers, also on behalf of the participants, would like to thank the staff and the management of Schloss Dagstuhl for their support throughout the 1,5 years of preparations of this very pleasant and successful event.

4.43 Computation and Palaeography: Potentials and Limits

Organizers: Tal Hassner, Malte Rehbein, Peter A. Stokes, and Lior Wolf
Seminar No. 12382

Date: 18.–21. September, 2012 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.2.9.184

© Creative Commons BY-NC-ND 3.0 Unported license

© Tal Hassner, Malte Rehbein, Peter A. Stokes, and Lior Wolf

Participants: Dimitris Arabadjis, Nachum Dershowitz, Matthieu Exbrayat, Shira Faigenbaum, Melanie Gau, Tal Hassner, R. Manmatha, Ophir Münz-Manor, Eyal Ofek, Micalis Panagopoulos, Robert Sablatnig, Wendy Scase, Timothy Stinson, Peter A. Stokes, Dominique Stutzmann, Ségolène Tarte, Lior Wolf



The Schloss Dagstuhl Perspectives Workshop on “Computation and Palaeography: Potentials and Limits” focused on the interaction of palaeography, the study of ancient and medieval documents, and computerized tools developed for analysis of digital images in computer vision. During the workshop, the interaction between domain experts from palaeography and computer scientists with computer vision backgrounds has yielded several very clear themes for the future of computerized tools in palaeographic research. Namely,

- Difficulties in communication between palaeographers and computer scientists is a prevailing problem. This is often reflected not only in computerized tools failing to meet the requirements of palaeography practitioners but also in the terminology used by the two disciplines. Better communication should be fostered by joint events and long-term collaborations.
- Computerized palaeographic tools are often black boxes which put the palaeography scholar on one end of the system, only receiving a systems output, with little opportunity to directly influence how the system performs or to communicate with it using natural palaeographic terminology. The long-term desire is to have the scholar at the center of the computerized system, allowing interaction and feedback in order to both fine-tune performance and better interpret and communicate results. This is crucial if palaeography is to become a truly evidence-based discipline. To this end the use of high-level terminology, natural to palaeography, should be integrated into computerized palaeographic systems.
- Palaeographic data, scarce to begin with, is even more restricted by accessibility and indexing problems, non-standard benchmarking techniques and the lack of accurate meta-data and ground truth information. Multiple opportunities were identified for acquiring data and disseminating it both in the palaeographic research community and outside to the general public.
- Palaeographic research is largely restricted to the domain of experts. Making palaeography accessible to non-experts by using

computerized tools has been identified as an effective means of disseminating valuable cultural heritage information while at the same time potentially giving rise to crowdsourcing opportunities, such as those proved successful in other domains.

The manifesto which resulted from this work elaborates on the existing challenges and limitations of the field and details the long-term recommendations that have emerged from the workshop.

4.44 Algorithms and Complexity for Continuous Problems

Organizers: Alexander Keller, Frances Kuo, Andreas Neuenkirch, and Joseph F. Traub
Seminar No. 12391

Date: 23.–28. September, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.9.200

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Alexander Keller, Frances Kuo, Andreas Neuenkirch, and Joseph F. Traub



Participants: Christoph Aistleitner, Martin Altmayer, James M. Calvin, Ronald Cools, Stephan Dahlke, Thomas Daun, Steffen Dereich, Nicolas Döhring, Massimo Fornasier, Stefan Geiss, Michael Gnewuch, Mario Hefter, Stefan Heinrich, Aicke Hinrichs, Peter Kritzer, Thomas Kühn, Frances Kuo, Gunther Leobacher, Peter Mathé, Klaus Meer, Thomas Müller-Gronbach, Valeriya Naumova, James Nichols, Erich Novak, Dirk Nuyens, Jens Oettershagen, Anargyros Papageorgiou, Sergei V. Pereverzev, Iasonas Petras, Friedrich Pillichshammer, Leszek Plaskota, Pawel Przybylowicz, Klaus Ritter, Daniel Rudolf, Otmar Scherzer, Wolfgang Ch. Schmid, Reinhold Schneider, Christoph Schwab, Winfried Sickel, Ian H. Sloan, Rob Stevenson, Shu Tezuka, Joseph F. Traub, Mario Ullrich, Tino Ullrich, Jan Vybiral, Grzegorz Wasilkowski, Markus Weimar, Henryk Wozniakowski, Larisa Yaroslavtseva, Marguerite Zani

This was already the 11th Dagstuhl Seminar on Algorithms and Complexity for Continuous Problems over a period of 21 years. It brought together researchers from different communities working on computational aspects of continuous problems, including computer scientists, numerical analysts, applied and pure mathematicians. Although the seminar title has remained the same many of the topics and participants change with each seminar and each seminar in this series is of a very interdisciplinary nature.

Continuous computational problems arise in diverse areas of science and engineering. Examples include path and multivariate integration, approximation, optimization, as well as operator equations. Typically, only partial and/or noisy information is available, and the aim is to solve the problem within a given error tolerance using the minimal amount of computational resources. For example, in high-dimensional integration one wants to compute an ε -approximation to the integral with the minimal number of function evaluations. Here it is crucial to identify first the relevant variables of the function. Understanding the complexity of such problems and construction of efficient algorithms is both important and challenging.

The current seminar attracted 51 participants from more than 10 different countries all over the world. About 30% of them were young researchers including PhD students. There were 40 presentations covering in particular the following topics:

- Biomedical learning problems
- Random media
- Computational finance
- Noisy data
- Tractability
- Quantum computation
- Computational stochastic processes
- High-dimensional problems

The work of the attendants was supported by a variety of funding agencies. This includes the Deutsche Forschungsgemeinschaft, the

Austrian Science Fund, the National Science Foundation (USA), and the Australian Research Council. Many of the attendants from Germany were supported within the DFG priority program SPP 1324 on “Extraction of Quantifiable Information from Complex Systems”, which is strongly connected to the topics of the seminar.

As always, the excellent working conditions and friendly atmosphere provided by the Dagstuhl team have led to a rich exchange of ideas as well as a number of new collaborations. Selected papers related to this seminar will be published in a special issue of the *Journal of Complexity*.



Fig. 4.14

Uwe Loebens – Lagune I. Part of the Dagstuhl art collection and donated by SaarToto, Reinhard Wilhelm, friends and relatives of Hannah and Sebastian Fischer, Roland and Ute Vollmar, and participants of the Dagstuhl Seminars 06121, 06081, and 06281.

4.45 Web Application Security

Organizers: Lieven Desmet, Martin Johns, Benjamin Livshits, and Andrei Sabelfeld
Seminar No. 12401

Date: 30. September – 05. October, 2012 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.2.10.1

© © ⊖ Creative Commons BY-NC-ND 3.0 Unported license

© Lieven Desmet, Martin Johns, Benjamin Livshits, and Andrei Sabelfeld



Participants: Marco Balduzzi, Nataliia Bielova, Arnar Birgisson, Egon Börger, Bastian Braun, Juan Chen, Ravi Chugh, Jorge Cuellar, Valentin Dallmeier, Philippe De Ryck, Lieven Desmet, Akhawe Devdatta, Daniele Filaretti, Cormac Flanagan, Cédric Fournet, Michael Franz, Dieter Gollmann, Arjun Guha, Daniel Hedin, Mario Heiderich, Boris Hemkemeier, Michael Hicks, Thorsten Holz, Thomas Jensen, Ranjit Jhala, Martin Johns, Shriram Krishnamurthi, Benjamin Livshits, Sergio Maffei, Fabio Massacci, John C. Mitchell, Nick Nikiforakis, Martin Ochoa, Frank Piessens, Joe Gibbs Politz, Joachim Posegga, Tamara Rezk, Eric Rothstein, Andrei Sabelfeld, Sebastian Schinzel, Juraj Somorovsky, Nikhil Swamy, Steven Van Acker, John Wilander

The Dagstuhl seminar on *Web Application Security* aimed to bring researchers together in the field of web security, both from academia and industry. The seminar is a follow-up of the Dagstuhl Seminar 09141 on Web Application Security in 2009 [1, 2].

■ Research context

Since its birth in 1990, the web has evolved from a simple, stateless delivery mechanism for static hypertext documents to a fully-fledged run-time environment for distributed multi-party applications. Recently, the web technologies have gradually shifted from a central server technology towards a rich/stateful client paradigm and lively interaction models. The wave of popular peer-to-peer web applications and web mashup applications confirm this emerging trend. But the shift from the server-centered paradigm poses a significant challenge of securing web applications in the presence of multiple stakeholders, including security-ignorant end-users. This motivates the need for solid *web application security*.

The seminar aimed to address the open question of how to protect against the pervasive threats to web applications. Some of the key objectives put forward are (i) over-viewing the state of the art to consolidate and structure it, (ii) identifying key challenges, and (iii) brainstorming on new ideas and approaches towards resolving these challenges.

The inception of this Dagstuhl seminar was strongly inspired by the following emerging trends and challenges in the web security landscape:

Fine-grained access control. Fine-grained access control policies define how the application authenticates and authorizes end users, from which application contexts the application can be consulted, and which interaction sequences maintain the application's integrity (i.e. control-flow integrity). Our objective was to address a range of questions from formal foundation of

authentication policies and protocols to the practicalities of authentication such as secure session management.

Information-flow control. Information-flow control specifies how sensitive data, possibly originating from multiple content providers in multiple trust domains, can be used in data aggregations, and client-side and server-side processing as is typically done in mashups. Challenges here include reconciling information-flow policies from several involved parties, with possibly conflicting goals. Moreover, tracking end-to-end information flow in web applications remains an open question. Our objective was to establish an enhanced understanding of how to make information-flow control policies and mechanisms practical in a web setting.

Secure composition. Secure composition policies specify how active third-party components, for instance written in JavaScript, can be securely integrated into applications via client-side and server-side mashups. By nature, web mashups heavily depend on interaction and communication across different origins, but contradictory, mashup security relies on separation techniques for protecting both code and data. As a result, traditional HTML techniques (mainly based on the same-origin policies) fail to address both the interaction and separation needs. We wanted to explore principled approaches to achieve the delicate balance between interaction and separation in security composition.

Cross-domain interaction. One of the original and still unresolved problems of the web is the inherent incompatibility between the cross-domain nature of the hyperlink and the same-origin security policy of its active content. In the recent past the situation has become even more complex with the introduction of client-side primitives for cross-domain interaction, such as CORS. Our objective was to assess the impact of current developments and identify promising directions for solutions.

Recent advances in JavaScript and HTML5. There are several

technological advances in the latest versions of JavaScript (such as strict mode, frozen objects, proxies and SES), that might contribute to the security of web applications. In addition, the research community did make important steps forward in understanding and improving the language by formalizing its semantics. At the same time, web specification (including HTML5 and CSP) are adding tons of new features as well as security measures as part of the browsing environment. Our objective was to have an enhanced understanding of the latest trends and research advances in JavaScript and HTML5 with respect to security.

■ Seminar program

The seminar attracted 44 participants, coming from companies and research institutions across Europe and the US. The group represented a nice mix of participants of academia and industry (including researchers of Siemens, SAP, Trend Micro, and Microsoft as well as two banks) and a good balance between junior and senior researchers.

The seminar had a well-filled program, with 3 keynotes, 28 research talks, 3 break-out sessions and 15 5-minute talks. The organizers aimed at keeping enough time during the breaks and in the evening for informal discussion. In addition, the participants went on a hike to the lake on Wednesday afternoon, as part of the social program.

■ Keynotes

The first three days, the floor was opened by keynotes to set the scene and inspire the discussions. The organizers invited the following three keynote speakers.

- Martin Johns (SAP Research – Karlsruhe, DE) – *Web Application Security: Are we there yet?*
- Shriram Krishnamurthi (Brown University – Providence, US) – *Browser Extension Analysis and Other JavaScript Adventures*
- John C. Mitchell (Stanford University, US) – *Science of Web Security and third-party tracking*

Martin Johns opened the Dagstuhl seminar on Monday by assessing the current state-of-practice in web security, 3 years after the previous Dagstuhl Seminar on Web application security. He sketched the evolving web landscape, and surveyed to what extent the results achieved so far suffice, and what is still missing. In particular, Martin gave a heads-up on client-side complexity and server-driven security, as being developed in the EU-FP7 project WebSand⁴.

Shriram Krishnamurthi discusses techniques based on typing to verify web applications, and demonstrated how these techniques can also be used to verify browser extensions. Such a verification can for instance assure that no unsafe functions are called within an extension, while operating in *private browsing mode*. As part of the underlying toolkit, Shriram presented core semantics of JavaScript in λ_{JS} (lambdaJS), and showed how a JavaScript program can be desugared in λ_{JS} [3].

John C. Mitchell focused on the *science of security and principles*, and demonstrated this by means of relevant web security examples. He emphasized the importance of defining system models, adversary models as well as desired properties of system, and argued that it seems feasible to verify web security properties. An interesting research question to be answered by such a *scientific* approach would be “Does CSP prevent XSS?”, and John challenged the audience to tackle this challenge. In addition, he discussed the

importance of experimental studies and gave some highlights on recent research results on web tracking.

■ Research talks

The organizers invited all the participants to take the floor during the seminar, and encouraged the presenters to step away from typical conference presentations, but rather strive for interaction with the audience and engage discussions.

Web security is a broad research domain, and the seminar was able to attract web security researchers with various backgrounds. As a result, a diverse set of recent research results was presented during the seminar, and these can be grouped in 5 topical clusters:

1. Web security vulnerability landscape
2. Information-flow control
3. JavaScript formalization
4. JavaScript confinement
5. Infrastructure and server hardening

■ Break-out sessions

To complement the keynotes and the research talks, the organizers opted to have three parallel break-out sessions as part of the seminar program. The break-out sessions enabled participants to discuss selected topics in web security research in an informal setting and in smaller teams. The three topics of the break-out sessions were:

- Cross-Site Scripting (XSS)
- JavaScript
- Information-Flow

The main purpose of the break-out sessions was to informally discuss the most important state-of-the-art and research challenges. As part of the break-out sessions, the teams identified and enlisted in a bottom-up way the most relevant state-of-the-art work, as well as the set of main challenges and research directions for the specific web security research area. The break-out sessions consisted of three slots of 70 minutes on Monday, Tuesday and Thursday. Participants joined the break-out sessions of their choice on Monday and Tuesday, and were encouraged to take part of two different sessions. The session on Thursday was used to report back the results of the three break-out sessions to the full group by means of a small presentation.

■ 5-minute talks

Finally, to encourage participants to pitch new research ideas, or highlight some relevant results, we also had two sessions specifically targeted at **5-minute talks**.

■ Conclusion

The Dagstuhl seminar on Web Application Security was a timely follow-up of the previous Dagstuhl seminar on this topic in 2009. The research domain has been maturing over the last five years, and new challenges have emerged such as the client-side complexity, the need of information-flow control enforcement, and hardening of JavaScript code.

The seminar brought 44 web security researchers together, coming from companies and research institutions across Europe and the US. The seminar had a well-filled program, with 3 keynotes, 28 research talks, and 15 5-minute talks. As web application security is a broad research domain, a diverse set of recent research results was presented during the talks, covering the web security vulnerabi-

⁴ EU-FP7 STREP WebSand, <https://www.websand.eu/>

lity landscape, information-flow control, JavaScript formalization, JavaScript confinement, and infrastructure and server hardening.

In addition to the plenary program, the seminar also featured three parallel break-out sessions on Cross-Site Scripting (XSS), JavaScript and Information-flow control. The main goal of the break-out sessions was to informally discuss the most important state-of-the-art work, as well as to identify the main challenges and research directions for future research, as documented in this report.

References

- 1 D. Boneh, U. Erlingsson, M. Johns, and B. Livshits. 09141 abstracts collection – web application security. In *Proc. of Web Application Security*, number 09141 in Dagstuhl Seminar Proceedings, 2010. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Germany.
- 2 D. Boneh, U. Erlingsson, M. Johns, and B. Livshits. 09141 executive summary – web application security. In *Proc. of Web Application Security*, number 09141 in Dagstuhl Seminar Proceedings, 2010. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, Germany.
- 3 A. Guha, C. Saftoiu, and S. Krishnamurthi. The essence of javascript. In *Proc. of the 24th European Conf. on Object-oriented programming*, ECOOP'10, pp. 126–150. Springer, 2010.
- 4 N. Nikiforakis, L. Invernizzi, A. Kapravelos, S. Van Acker, W. Joosen, C. Kruegel, F. Piessens, and G. Vigna. You are what you include: large-scale evaluation of remote javascript inclusions. In *Proc. of the 2012 ACM Conf. on Computer and communications security*, CCS '12, pp. 736–747, ACM, 2012.
- 5 S. Van Acker, N. Nikiforakis, L. Desmet, W. Joosen, and F. Piessens. FlashOver: Automated discovery of cross-site scripting vulnerabilities in rich internet applications. In *AsiaCCS*, May 2012.
- 6 M. Heiderich, M. Niemietz, F. Schuster, T. Holz, and J. Schwenk. Scriptless attacks: stealing the pie without touching the sill. In *Proc. of the 2012 ACM Conf. on Computer and communications security*, CCS '12, pp. 760–771, ACM, 2012.
- 7 W. De Groef, D. Devriese, N. Nikiforakis, and F. Piessens. Flowfox: a web browser with flexible and precise information flow control. In *Proc. of the 2012 ACM Conf. on Computer and communications security*, CCS '12, pp. 748–759, ACM, 2012.
- 8 T. H. Austin and C. Flanagan. Multiple facets for dynamic information flow. pp. 165–178, 2012.
- 9 D. Hedin and A. Sabelfeld. Information-flow security for a core of javascript. In *Proc. of the 25th IEEE Computer Security Foundations Symp.*, CSF '12, pp. 3–18, IEEE, 2012.
- 10 A. Birgisson, D. Hedin, and A. Sabelfeld. Boosting the permissiveness of dynamic information-flow tracking by testing. In *Proc. of Computer Security*, ESORICS'12, vol. 7459 of LNCS, pp. 55–72. Springer, 2012.
- 11 P. Mardziel, S. Magill, M. Hicks, and M. Srivatsa. Dynamic enforcement of knowledge-based security policies. In *Proc. of the 2012 IEEE 25th Computer Security Foundations Symp.*, CSF'11, 0:114–128, IEEE, 2011.
- 12 B. Köpf, L. Mauborgne, and M. Ochoa. Automatic quantification of cache side-channels. In *Computer Aided Verification*, vol. 7358 of LNCS, pp. 564–580. Springer, 2012.
- 13 F. Freiling and S. Schinzel. Detecting hidden storage side channel vulnerabilities in networked applications. In *Proc. of Future Challenges in Security and Privacy for Academia and Industry*, vol. 354 of *IFIP Advances in Information and Communication Technology*, pp. 41–55, Springer, 2011.
- 14 S. Schinzel. An efficient mitigation method for timing side channels on the web. In *2nd International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE)*, 2011.
- 15 R. Chugh, D. Herman, and R. Jhala. Dependent types for javascript. In *Proc. of the ACM Int'l Conf. on Object oriented programming systems languages and applications*, OOPSLA '12, pp. 587–606, ACM, 2012.
- 16 C. Fournet, N. Swamy, J. Chen, P.-E. Dagand, P.-Y. Strub, and B. Livshits. Fully abstract compilation to javascript. In *Proc. of the Symp. on Principles of Programming Languages (POPL)*, January 2013.
- 17 J. Gibbs Politz, M.J. Carroll, B.S. Lerner, J.n Pombrio, and S. Krishnamurthi. A tested semantics for getters, setters, and eval in javascript. In *Proc. of the 8th Symp. on Dynamic languages*, DLS'12, pp. 1–16, ACM, 2012.
- 18 D. Akhawe, P. Saxena, and D. Song. Privilege separation in html5 applications. pp. 23–23, 2012.
- 19 P. Agten, S. Van Acker, Y. Brondsema, P. H. Phung, L. Desmet, and F. Piessens. Jsand: complete client-side sandboxing of third-party javascript without browser modifications. In *Proc. of the 28th Annual Computer Security Applications Conf.*, ACSAC '12, pp. 1–10, ACM, 2012.
- 20 K. Bhargavan, C. Fournet, R. Corin, and E. Zalinescu. Verified cryptographic implementations for tls. *ACM Trans. Inf. Syst. Secur.*, 15(1):3, 2012.
- 21 T. Jager, S. Schinzel, and J. Somorovsky. Bleichenbacher's attack strikes again: Breaking pkcs#1 v1.5 in xml encryption. In *Proc. of Computer Security*, ESORICS 2012, vol. 7459 of LNCS, pp. 752–769, Springer, 2012.
- 22 T. Jager and J. Somorovsky. How to break xml encryption. In *Proc. of the 18th ACM Conf. on Computer and communications security*, CCS '11, pp. 413–422, ACM, 2011.
- 23 B. Braun, P. Gemein, H. P. Reiser, and J. Posegga. Control-flow integrity in web applications. In *International Symp. on Engineering Secure Software and Systems*, ES-SoS'13, February 2013. [to appear].
- 24 C. Bansal, K. Bhargavan, and S. Maffei. Discovering concrete attacks on website authorization by formal analysis. In *Proc. of 25th IEEE Computer Security Foundations Symp.*, CSF'12, pp. 247–262, IEEE, 2012.



Fig. 4.15

Natascha Popp – ohne Titel. Part of the Dagstuhl art collection and donated by Dominik Heckmann and Ute and Roland Vollmar.

4.46 Coalgebraic Logics

Organizers: Ernst-Erich Doberkat, and Alexander Kurz
Seminar No. 12411

Date: 08.–12. October, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.10.38

© Creative Commons BY-NC-ND 3.0 Unported license

© Alexander Kurz and Ernst-Erich Doberkat



Participants: Samson Abramsky, Octavian Vladut Babus, Jort Bergfeld, Marta Bilkova, Marcello M. Bonsangue, Liang-Ting Chen, Vincenzo Ciancia, Corina Cirstea, Josée Desharnais, Ernst-Erich Doberkat, Matej Dostal, Norman Francis Ferns, H. Peter Gumm, Helle Hvid Hansen, Ichiro Hasuo, Dirk Hofmann, Bart Jacobs, Achim Jung, Klaus Keimel, Dexter Kozen, Clemens Kupke, Alexander Kurz, Paul Blain Levy, Tadeusz Litak, Stefan Milius, M. Andrew Moshier, Rob Myers, Alessandra Palmigiano, Prakash Panangaden, Dusko Pavlovic, Daniela Petrisan, Katsuhiko Sano, Lutz Schröder, Alexandra Silva, Ana Sokolova, Shashi M. Srivastava, Sam Staton, Pedro Sánchez Terraf, Baltasar Trancón y Widemann, Henning Urbat, Jiri Velebil, Yde Venema, Joost Winter, Chunlai Zhou

■ Background

Modal Logic is a field with roots in philosophical logic and mathematics. As applied to Computer Science it has become central in order to reason about the behavioural and temporal properties of computing and communicating systems, as well as to model properties of agents such as knowledge, obligations, and permissions. Two of the reasons for the success of Modal Logic are the following. First, many modal logics are—despite their remarkable expressive power—decidable and, therefore, amenable to automated reasoning and verification. Second, Kripke’s relational semantics of modal logic turned out to be amazingly flexible, both in terms of providing techniques to prove properties of modal logics and in terms of allowing the different applications of Modal Logic to Artificial Intelligence, Software Agents, etc.

Coalgebra is a more recent area. Following on from Aczel’s seminal work on non-well founded set theory, coalgebra has been developed into a general theory of systems. The basic idea is that coalgebras are given with respect to a parameter F . Technically, the parameter F is a *functor* on a *category* \mathcal{C} .

Different choices of F yield, for example, the Kripke frames and models of modal logic, the labelled transition systems of process algebra, the deterministic automata of formal language theory, or the Markov chains used in statistics. Rutten showed that, in analogy with Universal Algebra, a theory of systems, called Universal Coalgebra, can be built uniformly in the parameter F , simultaneously covering the above and other examples. Crucial notions such as behavioural equivalence (observational equivalence, bisimilarity), final semantics and coinduction find their natural place here.

Coalgebraic Logic combines Coalgebra and Modal Logic to study *logics of systems* uniformly in the parameter F . Given the plethora of different transition systems and their ad hoc logics, such a uniform theory is clearly desirable. *Uniformity* means that results on, for example, completeness, expressivity, finite model property and complexity of satisfiability can be established at once for all

functors (possibly satisfying some, usually mild, conditions). Additionally, there is also a concern for *modularity*: Typically, a parameter F is composed of basic features (such as input, output, non-determinism, probability). Modularity then means that the syntax/proof systems/algorithms for the logic of F are obtained compositionally from the syntax/proof systems/algorithms for the logics of the basic features.

What has been achieved: The power of uniformity and modularity Following on from Moss’ seminal paper, Coalgebraic Logic is now growing into a successful area. Conferences in this area now treat topics such as completeness, expressivity, compositionality, complexity, rule formats for process calculi, containing several hitherto unknown results on these classic topics.

The uniformity achieved in the above cited work is based on varying the type F for a given base category \mathcal{C} , usually the category of sets. But it is also of interest to vary \mathcal{C} .

Here probabilistic approaches deserve to be mentioned. In a number of papers Markov transition systems could be shown to interpret modal logics under different assumptions on the probabilistic structure. It was shown that general measurable spaces provide too general a structure, but that analytic spaces with Borel transition laws offer just the right blend of generality and measure theoretic accessibility. In this context, it was shown that logical equivalence, bisimilarity, and behavioral equivalence are equivalent concepts. Recent work shows that this can be extended to distributional aspects as well: instead of comparing states proper, one has a look at distributions over the states of a Kripke model. This approach was recently generalized from general modal logics to coalgebraic logics; these logics are interpreted through coalgebras in which the subprobability functor and the functor suggested by the phenomenon to be modelled form various syntactic alliances. This generalization brings stochastic coalgebraic logic into the mainstream of coalgebraic logics: the problems considered are similar, and one sees a convergence of methods.

Nevertheless it is to be mentioned that the probabilistic approach brings its own idiosyncratic touch due to measure theoretic problems. This entails among others that one sometimes has to work in a very specific topological context, for otherwise solutions are not available. On the other hand, leaving a topological context and working in general measurable spaces poses the question of the limits to the coalgebraic approach: What can be achieved in general measurable spaces, or in measurable spaces in which some of the properties are available (like Blackwell spaces, which are countably generated without being topological)?

Quantitative aspects are also considered when it comes to approximate Markov transition processes defined on uncountably infinite state spaces through finite processes. This is a classical problem that arises mostly in practical applications of Markov transition systems; it has to be investigated from a logical vantage point as well.

■ Structuring the Seminar

When we planned the seminar, we envisaged six broad topics. One of the outcomes of this seminar, as compared to the one of 2009, is that the different subcommunities in coalgebraic logic moved closer together, exchanging ideas, techniques, problems and also researchers. Consequently, it seems difficult, if not impossible, to divide up all the talks consistently among the distinct research topics. We will nevertheless try to describe some trends.

■ Probabilistic Transition Systems

The focus of Markov transition systems shifted from the consideration of specific problems (like interpreting a particular class of logics) to structural problems which are treated with the instruments provided by coalgebras. The talk presented Panangaden concentrated on the duality of Markov transition systems and various function spaces, most of them well known in functional analysis. The Radom-Nikodym Theorem provides a very sophisticated tool for switching between these representations. Doberkat's talk dealt with stochastic effectivity functions as an extension of Markov transition systems for the interpretation of more complicated logics like, e.g., Parikh's game logic. Urbat showed that both the Hausdorff and the Kantorovic functor, which are widely used to model probabilistic nondeterminism are finitary, improving some well known results; at the same time, this results raises some interesting topological questions.

Quite apart from structural problems, another approach has been presented by Srivastava; he gave a tutorial talk on deduction systems for probabilistic logics, based on the work by Goldblatt and by Zhou. The set theoretic problems which originate with bisimilarity were taken up by Terraf, who extended a well-known result from descriptive set theory on the structure of equivalence relations to bisimulations, hereby indicating some of the caveats one has to observe in classical set theory.

■ Coalgebras and automata theory

Whereas the final coalgebra describes all infinite behaviours, the theory of formal languages suggests that the regular or rational sets of behaviours should be of special interest. This is indeed the case and the talks of Milius, Myers, Sokolova and Winter presented some of the latest developments. More generally, this direction of generalising results from automata theory also saw talks of Hansen/Silva and of Venema.

■ Process algebra and operational semantics

Bonsangue presented a coalgebraic account of the 'bisimulation-up-to' proof technique and Staton had new results on finite power set functors. Another direction is concerned with applying

coalgebraic techniques to other process equivalences than bisimulation. In particular, Hasuo and Cirstea studied trace equivalence, whereas Levy's tutorial on relation liftings was concerned with various notions of simulation.

■ Coalgebraic logic beyond sets

After the successes of set-based coalgebra, quite some effort goes now into extending results to more general settings. Jacobs presented a novel framework uniformly covering the classical, probabilistic and quantum case. Pavlovic introduced his ideas about a monoidal computer to bridge the gap between high-level specification and low-level computational models such as Turing machines. Talks by Bilkova, Dostal, and Velebil explored how to harness enriched category theory whereas Moshier is extending coalgebraic logics from the discrete to the setting of compact Hausdorff topological spaces, a topic that also surfaced in Hofmann's contribution. Petrisan studied final coalgebra in nominal sets.

■ Extensions of coalgebraic logics

Litak led a discussion session about the directions of generalising coalgebraic modal logic to formalisms with explicit quantifiers. Palmigiano reported latest results on extensions with fixpoint operators and Venema discussed some of the challenges and open problems in this area. Sano showed how to extend coalgebraic logic by an actuality operator and whereas Schröder explored the border of decidability for coalgebraic hybrid logic.

■ Applications

One of the outcomes of the seminar was the excitement generated by the wide range of applications which are now coming into the scope of coalgebraic techniques. Examples include Abramsky's results on infinite economic non-cooperative games, Trancon y Wiedemann's contributions to a reformulation of the foundations of ecology, and Kozen's ideas of making coalgebraic techniques available to the working programmer and to the working mathematician.

4.47 Algebraic and Combinatorial Methods in Computational Complexity

Organizers: Manindra Agrawal, Thomas Thierauf, and Christopher Umans

Seminar No. 12421

Date: 14.–19. October, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.10.60

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Manindra Agrawal, Thomas Thierauf, and Christopher Umans



Participants: Farid Ablyev, Manindra Agrawal, Eric Allender, Vikraman Arvind, David A. Mix Barrington, Markus Bläser, Andrej Bogdanov, Harry Buhrman, Sourav Chakraborty, Arkadev Chattopadhyay, Samir Datta, Volker Diekert, Klim Efremenko, Stephen A. Fenner, Lance Fortnow, Anna Gál, William Gasarch, Parikshit Gopalan, Frederic Green, Joshua A. Grochow, Steve Homer, Valentine Kabanets, Neeraj Kayal, Pascal Koiran, Swastik Kopparty, Michal Koucký, Matthias Krause, Klaus-Jörn Lange, Sophie Laplante, Bruno Loff, Meena Mahajan, Pierre McKenzie, Or Meir, Jochen Messner, Natacha Portier, Atri Rudra, Chandan Saha, Rahul Santhanam, Shubhangi Saraf, Uwe Schöning, Rocco Servedio, Ronen Shaltiel, Simon Straub, Amnon Ta-Shma, Thomas Thierauf, Jacobo Toran, Christopher Umans, Virginia Vassilevska Williams, Nikolay K. Vereshchagin, Ryan Williams, Amir Yehudayoff, David Zuckerman

The seminar brought together more than 50 researchers covering a wide spectrum of complexity theory. The focus on algebraic and combinatorial methods showed the great importance of such techniques for theoretical computer science. We had 30 talks, most of them lasting about 40 minutes, leaving ample room for discussions. In the following we describe the major topics of discussion in more detail.

Circuit Complexity is an area of fundamental importance to Complexity, which has resisted all efforts to prove strong lower bounds. We had several talks on circuit upper and lower bounds.

Valentine Kabanets considered the following compression problem: Given a truth table of an n -variate Boolean function f , find a Boolean circuit C of non-trivial size (less than $2^n/n$) that computes f . The motivation comes from the desire to understand “easy” functions (what property of their truth tables makes such functions compressible), and “hard” functions (once we understand which functions are compressible, we may be able to prove certain functions to require high circuit complexity). As an example, he showed that the class of functions computable by polysize AC^0 -circuits, and linear-size de Morgan formulas are compressible.

The Shub and Smale’s “tau-conjecture” states that the number of integer roots of a univariate polynomial should be polynomially bounded in the size of the smallest straight-line program computing it. Pascal Koiran proposed a real version of the tau-conjecture in his talk. If this new conjecture is true, then the permanent polynomial cannot be computed by polynomial-size arithmetic circuits.

Fred Green showed that degree- d block-symmetric multivariate polynomials modulo any odd p correlate with parity exponentially better than degree- d symmetric polynomials, for certain values of d . The result is obtained through the development of a theory call spectral analysis of symmetric correlation, which originated in work of Cai, Green, and Thierauf.

Chaudhuri and Radhakrishnan used certifying polynomials

to show that Approximate Majority cannot be computed by AC^0 -circuits of size $n^{1+o(1)}$. In his talk, Swastik Kopparty extended their technique and showed that Approximate Majority cannot be computed by AC^0 [parity]-circuits of size $n^{1+o(1)}$. This implies a separation between the power of AC^0 [parity]-circuits of near-linear size and uniform AC^0 [parity]-circuits of polynomial size.

Neeraj Kayal talked on the problem of computing the smallest formula for a polynomial given as a blackbox. The complexity of this problem is still unclear. It is conjectured that it is NP-hard. Neeraj presented his very impressive result, a randomized algorithm that given blackbox access to the polynomial f computed by an unknown/hidden arithmetic formula reconstructs, on the average, an equivalent or smaller formula in time polynomial in the size of its output. This is the strongest model of arithmetic computation for which a reconstruction algorithm is presently known, albeit efficient in a distributional sense rather than in the worst case.

Coding Theory Error-correcting codes, particularly those constructed from polynomials, lie at the heart of many of the most significant results in Computational Complexity (e.g. interactive proofs, PCPS, hardness amplification, explicit constructions, derandomization, etc.) In many of these applications it is evident that the *local-testability/decodability* of the code is critical.

A q -query Locally Decodable Code (LDC) is an error-correcting code that allows to read any particular symbol of the message by reading only q symbols of the codeword. In a completely new approach, Klim Efremenko showed how to construct q -query LDCs from representation theory. Parikshit Gopalan showed an equivalence between locally testable codes and Cayley graphs with certain spectral properties. These Cayley graphs can be viewed as “derandomized hypercubes” which preserve several important properties of the Boolean hypercube such as small-set expansion, large threshold rank and hypercontractivity.

Shubhangi Saraf talked about the classical theorem of Sylves-

ter-Gallai, which says that, if for every two points there is a third point on the line through them, then all points are on the same line. In the stable versions of the theorem, it is only guaranteed that many triples of points are approximately collinear. Configurations with many approximately collinear q -tuples of points also arise naturally in stable versions of Locally Correctable Codes over the complex numbers. She showed that such stable codes with constant query complexity do not exist.

Explicit Constructions Until recently the best-known construction of extractors (and their cousins, *condensers*) was a primarily combinatorial construction of Lu, Reingold, Vadhan, and Wigderson. Then Guruswami, Umans and Vadhan gave an entirely algebraic construction, utilizing the new polynomial error-correcting codes of Parvaresh and Vardy. Amnon Ta-Shma presented a new construction of condensers based on Parvaresh-Vardy codes. Amnon's condensers have entropy rate $1 - \alpha$ for subconstant α (in contrast to GUV which required constant α) and suffer only sublinear entropy loss.

Ronen Shaltiel presented new constructions of zero-error seedless dispersers for bit-fixing sources and affine sources. Ronen used these dispersers to construct an algorithm for a problem related to the Write-Once-Memory (WOM) problem in which once we raise a storage cell from zero to one, it is stuck at this value. He gives the first explicit scheme with asymptotically optimal rate.

Anna Gál identified a new class of superconcentrator-like graphs with connectivity properties distinct from previously studied ones. Anna showed that any circuit computing good codes must satisfy such superconcentrator-like properties.

Probabilistic proof systems is a sub-field of complexity theory that investigates questions such as “how can we use randomness to prove and verify assertions?”, “what do we gain from using randomness in verification procedures?”, and “what assertions can be verified by probabilistic verification procedures?”. Research in this area began in the 1980, and has led to several of the most important achievements of complexity theory in those decades.

A line of research from the recent years is aimed at finding alternative “combinatorial” proofs for those key results, i.e., proofs that do not rely on algebra. This line of research is motivated by trying to gain more intuition of those results, as well as to understand the properties of polynomials that make them useful for such constructions. Or Meir gave a very interesting survey talk about this line of research.

Complexity In a much appreciated talk, Joshua Grochow gave a very interesting survey-type talk on the Geometric Complexity Theory (GCT) program, which was introduced by Mulmuley and Sohoni to attack fundamental lower bound problems in complexity – such as P vs NP – using algebraic geometry and representation theory. Joshua succeeded in explaining very nicely some of the intuition behind the use of algebraic geometry and representation theory in complexity.

Michal Koucký gave a very interesting overview talk on the on-line labeling problem, where one receives n integers from the set $\{1, \dots, r\}$ and has to store them in an array of size m . The integers are presented sequentially in an arbitrary order, and must be stored in the array in sorted order. The complexity measure is essentially the number of times an element has to be moved in to make space for a newly arrived item. Michal showed that various known algorithms in the literature solve the problem asymptotically optimal.

Perfect matching is in P but not known to be in NC. Counting the number of perfect matchings in a graph is #P-complete. In contrast, Vazirani showed that counting the number of perfect matchings in a planar graph is in NC. So in particular, the decision version of perfect matching in planar graphs is in NC. Hence one

way to get perfect matching in NC could be to reduce perfect matching to perfect matching in planar graphs. An obvious approach to construct such a reduction is to come up with a *planarizing gadget*. Jochen Messner proved in his talk unconditionally that such a reduction is not possible for the perfect matching problem.

Steve Fenner considered the following two-player game on a finite partially ordered set (poset) S : each player takes turns picking an element x of S and removes all $y \succ x$ from S . The first one to empty the poset wins. Recently, Daniel Grier, an undergrad at the University of South Carolina, has settled the problem and showed that determining the winner of a poset game is PSPACE-complete. The reduction shows, that the game is already PSPACE-complete when the poset has only 3 levels. The complexity of two-level poset games is still open. Steve presented a simple formula allowing one to compute the status for a large class of two-level poset game.

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

4.48 Time-of-Flight Imaging: Algorithms, Sensors and Applications

Organizers: James Davis, Bernd Jähne, Andreas Kolb, Ramesh Raskar, and Christian Theobalt
Seminar No. 12431

Date: 21.–26. October, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.10.79

© © © Creative Commons BY-NC-ND 3.0 Unported license

© James Davis, Bernd Jähne, Andreas Kolb, Ramesh Raskar, and Christian Theobalt



Participants: Michael Balda, Christopher Barsi, Erhardt Barth, Sebastian Bauer, James E. Davis, Adrian Dorrington, Martin Eisemann, Peter Eisert, Jürgen Gall, Marcin Grzegorzek, Diego Gutierrez, Uwe Hahne, Thomas Helten, Ivo Ihrke, Slobodan Ilic, Shahram Izadi, Bernd Jähne, Andreas Jordt, Thomas Kilgus, Kwang In Kim, Reinhard Klein, Andreas Kolb, Daniel Kondermann, Jens Kubacki, Cornelia Lanz, Seungkyu Lee, Damien Lefloch, Frank Lenzen, Oisín Mac Aodha, Aditi Majumder, Rahul Nair, P. J. Narayanan, Shohei Nobuhara, Martin Profittlich, Ramesh Raskar, Christian Schaller, Andreas Schilling, Alexander Seitel, Christian Theobalt, Julien Thollot, Gordon Wetzstein, Giora Yahav, Ruigang Yang

In recent years, Time-of-Flight (ToF) depth imaging technology has seen immense progress. Time-of-Flight imaging is based on measuring the time that light, emitted by an illumination unit, requires to travel to an object and back to a detector. From this time, scene depth and possibly additional information that can not be measured by traditional intensity imaging, is inferred. While early ToF cameras were merely lab prototypes to prove a concept, recent sensor designs are at the edge of becoming operative products for mass market applications. A wide range of research disciplines is able to benefit from reliable and fast depth imaging technology, such as computer vision, computer graphics, medical engineering, robotics and computational photography, to name a few. Easy availability of affordable depth cameras will open the door for many new applications. The commercial success of the Microsoft™ Kinect device – a depth sensor based on an alternative measurement principle – gives a first impression on this.

Currently, manufacturers of ToF systems mainly focus on sensor technology and on the design of cameras. Sensor design has seen great advancements, but the data delivered by the cameras remain challenging and are affected by many types of systematic distortions and difficult scene dependencies. ToF data are thus hardly usable out-of-the-box and it takes proper mathematical modeling and algorithmic processing to apply the data in practical imaging and reconstruction scenarios. Algorithm design for ToF imagers, however, is still in its early days and many challenges remain. In this seminar, we plan to discuss and extend the state of the art in ToF imaging algorithms and applications with leading researchers in the field.

Also, currently, there is little dialogue between researchers developing ToF algorithms and sensor designers. Therefore, the seminar also strongly supported the manufacturers in getting up to date with all relevant research results and, even more importantly, it offered the opportunity to establish long-term partnerships and research collaborations. We also believe that this stronger interaction

will lead to more advanced sensor designs, as well as more powerful algorithmic solutions at the same time.

■ Description of the Seminar: Topics, Goals and Achievements

■ General Motivation

Time-of-Flight technology is based on measuring the time that light, emitted by an illumination unit, requires to travel to an object and back to a detector. This allows to measure distances with high precision and high speed. Recently, this principle has been the basis for the development of new range-sensing devices realized in standard CMOS and CCD technology and which are called ToF cameras, as well as in the context of photogrammetry Range Imaging (RIM) sensors. Unlike other 3D systems, the ToF camera is a very compact device. It has the potential of being one of the first low-price, off-the-shelf system to provide full-range distance information in at video rate.

Today the community using Time-of-Flight technology is scattered over many research disciplines without intense communication across research areas. Such communication is necessary, however, to fuse results from sensor technology, low-level ToF data processing and high-level image processing. Each of the above research disciplines that employs time-of-flight imaging has to develop algorithmic solutions to these very same core problem areas. Additionally, there are new hot topics that currently do not make use of this new technology but might benefit from it in the future, which further underlines the importance of ToF algorithm design.

In this seminar, we exploited this multi-disciplinarity, and brought researchers from computer vision, computer graphics, computational photography, image processing and engineering disciplines together that work with ToF imagers. Together, we defined the current state of the art in core algorithmic questions that ToF ima-

ging researchers are confronted with (additionally to the seminar by an edited book on the main results). We also contributed advancing the field by identifying current limitations, important future research directions, and by enabling a closer dialogue between algorithm and hardware designers to discuss future sensor designs.

■ Topics

Time-of-Flight imaging devices can measure scene depth largely independently of scene appearance and are generally based on extensions of standard video intensity camera hardware. ToF sensors can thus be used for static and dynamic scene capture. However, the data of these sensors suffer from a variety of deficiencies, such as low resolution, strong random noise, and non-trivial systematic distortions. These challenges have to be algorithmically addressed before ToF cameras become mainstream in any field of application. The main topic of this seminar was the definition and extension of the state of the art in ToF imaging problems in three core areas of algorithm and technology development that are described in the following.

Low level data processing, calibration and characterization Researchers in computer vision, computer graphics and image processing only just started to mathematically model the measurement characteristics of ToF sensors [1]. This is a fundamental prerequisite for calibration [2], as well as for well-founded design of low level ToF data processing.

The phase-based Time-of-Flight technology suffers from some specific problems that cause systematic calibration errors and parameter correlation issues. Due to the physical realization of light modulation in the emitting LEDs, the ideal sine-waveform light emittance is approximated by a band-limited rectangular waveform. This causes nonlinear depth distortions, called *wiggling errors*. In addition, there are several non-linear effects depending on multi-path light propagation, for example in the optical system or due to multiple reflections in the scene. Some effects are well-understood, but there are still open issues in depth calibration [3]. In addition, the calibration of external camera parameters suffers from strong correlation, since typically the cameras have limited field of view and low image resolution. Solutions to this problem can be found if a synchronous multi-camera calibration with rigidly coupled color and range camera rigs are investigated [4]. Coupling of high-resolution color video cameras with ToF cameras is hence an issue of further investigation. Latest ideas on sensor calibration will be reviewed and augmented in the seminar.

The knowledge gained through sensor calibration can also be exploited to create sensor simulations of high fidelity in software. This will be an invaluable tool test new algorithms. Proper sensor modeling also enables detailed sensor comparison and evaluation, and eventually even certification. A couple of research initiatives are underway to build in-depth mathematical sensor model of ToF imagers which will be discussed at the seminar [5].

Low level sensor calibration and sensor modeling enables more efficient and effective design of algorithms for low-level TOF processing. For instance, first low-level ToF filtering [6, 7] and ToF 3D superresolution approaches have been proposed [8, 9]. Most of these approaches have already demonstrated that a proper sensor model can be exploited for higher quality processing. In the seminar, we reviewed latest low-level processing techniques, and evaluated how new and better filtering and data enhancement techniques can be developed, also for rarely considered depth camera artifacts, such as ToF motion blur. We also discussed how such techniques can be integrated on the sensor and how the gained understanding of sensor characteristics can benefit the design of future sensors.

High level data processing for 3D reconstruction, understanding and recognition Low-level ToF imaging builds the foundations for the higher-level processing tasks that researchers and practitioners from many disciplines are confronted with. In most cases, such higher level processing aims to recover high-quality 3D models of static and/or dynamic scenes that should be displayed, analyzed, interactively modified, or used for recognition and scene understanding.

One major field of research using higher-level ToF image processing is computer graphics. Here, efficient acquisition of geometric models of static and dynamic scenes is of tremendous importance, and has many applications in interactive 3D rendering, geometric modeling and product design, 3D human computer interaction, cultural heritage, as well as professional media and game productions. ToF sensors can be an important asset here in order to replace the costly, highly specialized, complex and often intrusive acquisition technology currently used for such tasks. Static scene acquisition is mostly performed based on active scanners, using structured light or laser-based triangulation. Dynamic scene capture can also be achieved with structured light devices, and specialized optical systems that track fiducial markers exist for capturing motion. Being able to solve similar reconstruction tasks with only ToF cameras would be a big step ahead and eventually make 3D acquisition technology available to a wider range of users.

For a long time computer vision researchers have successfully developed 3D reconstruction approaches from single or multiple cameras that exploit certain photometric or radiometric cues. Many of them have in common that they are computationally expensive and that they only succeed under certain scene conditions, such as if scenes are sufficiently textured. An enormous potential lies in the fusion of ToF sensors with standard sensors for computer vision and robotics problems. Most areas in computer vision benefit from depth or range information; however, due to the difficulty in reconstruction of robust depth maps in real-world environments — especially in real-time applications — most state of the art solutions in areas like object recognition, gesture and action recognition in man-machine communication, pedestrian detection, and low-level tasks like segmentation just rely on 2D intensity information. Available depth and shape cues in real-time together with intensity information will open new possibilities to improve quality and robustness of algorithms and systems in such areas [10–15]. In this context, there are several open problems, which were discussed during this seminar: do we need to define new features to be extracted from Time-of-Flight data and which feature will lead to a gain in quality compared to nowadays state of the art solutions? How can we deal with resolution and noise level of such cameras to complement normal 2D intensity information? Will we need to fuse Time-of-Flight information with 2D intensity data of standard CCD cameras, or are there applications, that can benefit from Time-of-Flight cameras by itself?

Another area in which ToF imaging will play a major role in future, is video processing, in particular 3D video and 3D TV. The analysis of dynamic 3D scenes from video requires the simultaneous processing of color video and range data. While traditional approaches using multi-view video are already quite successful, the advent of ToF range technology allows novel insights, novel applications and ease of acquisition. Traditional multi-view depth reconstruction requires sufficiently textured scenes, which might not be the case for arbitrary scenes, especially in indoor environments. This might lead to incomplete reconstruction results. ToF range acquisition has the potential for handling range data in dynamic video, but still many issues need to be solved and discussed by experts: in particular the challenging noise, uncertainty in the measurements, and low resolution of current ToF cameras represent a challenge. First applications handling video-rate HD-TV depth processing can

be found in systems for 3D-Television capture [16] or in general computer graphics applications [17].

In many other areas, for instance computational photography, computational videography and medical engineering, researchers are facing similar reconstruction problems and can benefit from ToF sensors. For instance, in medical engineering, ToF cameras have been used to detect patient position [18] and respiratory motion in radiotherapy [19, 20].

The above list of examples shows that the algorithmic problems to be solved for making ToF sensors usable for high level reconstruction in different areas are very similar. The main challenge will be to enable high quality reconstruction despite strongly distorted and low-resolution raw ToF sensor output. Several strategies have been explored to attack these problems: Sensor fusion approaches combine depth and intensity cameras, spatio-temporal reconstruction approaches recover higher detail by accumulating and aligning measurements over time, superresolution and alignment can be combined to enable high-quality 3D reconstruction. Given such better quality reconstructions, the captured data can be employed as scene models or further analyzed for capturing motion and gestures, for recognizing activities, for recognizing objects, or for analyzing the environment in a navigation scenario. The seminar therefore reviewed latest algorithms for static and spatio-temporal 3D reconstruction from ToF data. We have also discussed how they need to be tuned for specific applications, such as motion capture and recognition. Finally, we discussed ways to better integrate low-level and high-level processing.

Sensor technology and new depth sensor designs

While algorithm design for low-level and high-level TOF imaging were the main focus of this seminar, we also initiated to enable a dialogue between hardware manufacturers and algorithm designers. On the one hand this familiarized hardware designers with the state-of-the-art in ToF data processing, and sensibilized them for the existing challenges and specific application requirements. In return, algorithm designers deepened their knowledge about the fundamental physical principles of ToF imaging and gain a better understanding for the physical origins of sensor characteristics.

It is possible that relatively simple changes to the ToF hardware would result in the possibility of new sensor designs. ToF cameras make use of a CMOS sensor that is an enhanced version of a normal camera with extra circuitry at each pixel, and a structured IR illuminator. A great deal of prior research exists on using “normal” CMOS cameras together with triangulation based structured light to recover depth. The structured illuminator in these two research areas makes use of different principles, and the internal frame rate of the ToF camera is much higher, but the hardware components are broadly similar, suggesting that sharing of ideas might be fruitful.

Importantly, ToF and triangulation have complementary error characteristics, strengths, and weaknesses. For example, ToF sensors tend to perform better at a distance, and triangulation tends to perform better at close range. This leaves open the possibility of new sensor designs that make use of ideas from both ToF and structured light, with greatly improved robustness and accuracy. For example: chips could be designed with both “normal” and “ToF” pixels, the ToF light source could have a focusing lens and spatial pattern, the modulated light used with the ToF sensor could be similar to structured light patterns, the data from ToF could be used as a rough guess to disambiguate phase/depth in structured light when there are not enough patterns.

We are convinced that through a dialogue between hardware and algorithm designers, both sides can benefit. An example for a related research area in which such a dialogue has already resulted in great advancements is the area of computational photography. There, algorithm designers and hardware manufacturers have worked

together on new designs for optical systems and processing algorithms that open up new ways of digital imaging, e.g., through high dynamic range imaging, wave front coding etc. We believe that the advent of ToF depth imaging technology is a further boost to this development, as it was already shown by new ideas on space-time imaging [21]. We also believe that ToF designs can have a similar impact in the emerging field of computational videography where future video sensors and processing paradigms are developed. We believe that the seminar served as a platform to initiate such developments by bringing together key players in the field. In this context, the pros and cons of alternative depth measurement sensors, such as IR-based active stereo cameras, have also been discussed.

Goals and Achievements of the Seminar

The overall goal of this seminar was to bring together researchers from several TOF-related disciplines, review the state-of-the-art in ToF imaging on both the algorithmic and hardware side, and develop new concepts for algorithms and devices that advance the field as a whole. The seminar was not intended to be a classical workshop or conference where mostly finished research is presented. We wanted the seminar to be a platform for identifying and discussing the big open research questions and challenges. More specifically, the following is a list of challenges that have been discussed at the seminar, since they form the basis of low-level and applied research with Time-of-Flight cameras:

- Low-level processing
 - Basic mathematical modeling of ToF cameras: image formation model, noise modeling, calibration of the sensor and optics.
 - Low-level image processing problems: resolution enhancement through superresolution and sensor fusion, data filtering, feature extraction under random and systematic distortions.
- High-level processing
 - Static shape scanning: high-quality geometry scanning, 3d superresolution, alignment approaches, probabilistic methods for reconstruction and alignment under noise.
 - Dynamic shape scanning: Spatio-temporal filtering, multi-sensor fusion approaches, model-based dynamic scene reconstruction, unsupervised dynamic scene reconstruction (joint model-building and motion reconstruction), marker-less motion and performance capture, 3d video.
- Improvements of sensor design: pixel design, light source design and arrangement, Time-of-Flight measurement principles: amplitude modulation vs. shutter. In this context we will also discuss standardization questions.

The seminar was very successful with respect to the set goals and initiated great interaction between researchers from different domains which had never happened in this way at other conferences or workshops.

In order to best initiate this interaction, we decided to organize a multi-faceted scientific programme. It consisted of a variety of different presentation formats. In particular, we had a series of research talks on the different research problems which we wanted to address in the seminar. When selecting the research talks, we planned for having a mix of presentations by junior and senior researchers, as well as balance of different topics. Presenters dedicated at least half the presentation time to address open research problems in order to spawn new research projects and collaborations. In order to further initiate discussion between researchers with different backgrounds, and in order to very practically identify potential research projects, we also organized working groups in which small

teams discussed certain focus topics. Finally, the seminar participants organized very informal evening sessions in which special cross-disciplinary research topics were discussed in a very informal way. Finally, a demo session enables researchers and hardware specialists to showcase their latest results.

As an outcome of this, a very lively discussion and interaction was started between participants, and many concrete research projects were defined. Most fruitful discussions started on the topics of: 1) how to better exploit existing hardware and software systems; 2) the limitation of existing sensors and how to break them; 3) new

combinations of existing (heterogeneous) sensors; 4) technical and economical limitations of hardware.

To achieve sustainability beyond the seminar the organizers will edit a book summarizing the main methods, applications, and challenges in the context of ToF technology based on the presentations and discussions during the seminar. Such a book is currently missing in the community and the seminar itself shall also act as catalyst for such a project. For more rapid dissemination of ideas and results, the organizers also created Wiki⁵ which will be eventually relocated and maintained permanently.

References

- 1 M. Schmidt and B. Jähne. A physical model of time-of-flight 3d imaging systems, including suppression of ambient light. In Proc. of *3rd Workshop on Dynamic 3-D Imaging*, volume 5742 of *LNCS*, pp. 1–15. Springer, 2009.
- 2 M. Erz and B. Jähne. Radiometric and Spectrometric Calibrations, and Distance Noise Measurement of TOF Cameras. In Proc. of *3rd Workshop on Dynamic 3-D Imaging*, volume 5742 of *LNCS*, pp. 28–41. Springer, 2009.
- 3 M. Lindner, I. Schiller, A. Kolb, and R. Koch. Time-of-flight sensor calibration for accurate range sensing. *Computer Vision and Image Understanding*, 114(12):1318–1328, 2010. Special issue on Time-of-Flight Camera Based Computer Vision.
- 4 I. Schiller, C. Beder, and R. Koch. Calibration of a pmd camera using a planar calibration object together with a multi-camera setup. In *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume Vol. XXXVII. Part B3a, pp. 297–302, Beijing, China, 2008. XXI. ISPRS Congress.
- 5 M. Keller and A. Kolb. Real-time simulation of time-of-flight sensors. *J. Simulation Practice and Theory*, 17:967–978, 2009.
- 6 Q. Yang, R. Yang, Ja. Davis, and D. Nister. Spatial-depth super resolution for range images. *Proc. CVPR*, pp. 1–8, 2007.
- 7 D. Chan, H. Buisman, C. Theobalt, and S. Thrun. A noise-aware filter for real-time depth upsampling. In *Proc. of ECCV Workshop on Multi-camera and Multi-modal Sensor Fusion Algorithms and Applications*, 2008.
- 8 S. Schuon, C. Theobalt, J. Davis, and S. Thrun. High-quality scanning using time-of-flight depth superresolution. In *IEEE CVPR Workshop on Time-of-flight Computer Vision*, 2008.
- 9 S. Schuon, C. Theobalt, J. Davis, and S. Thrun. Lidarboost: Depth superresolution for tof 3d shape scanning. In Proc. of *IEEE Int'l Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- 10 E. Rodner, D. Hegazy, and J. Denzler. Multiple kernel gaussian process classification for generic 3d object recognition from time-of-flight images. In *Proc. of the Int'l Conf. on Image and Vision Computing*, 2010.
- 11 D. Hegazy and J. Denzler. Generic 3d object recognition from time-of-flight images using boosted combined shape features. In Proc. of *Int'l Conf. on Computer Vision, Theory and Applications (VISAPP 09)*, 2009.
- 12 M. Kemmler, E. Rodner, and J. Denzler. Global context extraction for object recognition using a combination of range and visual features. In Proc. of the *Dynamic 3D Imaging Workshop*, volume 5742 of *LNCS*, pp. 96–109. Springer, 2009.
- 13 Ol. Kähler, E. Rodner, and J. Denzler. On fusion of range and intensity information using graph-cut for planar patch segmentation. *Int'l Journal of Intelligent Systems Technologies and Applications*, 5(3/4):365–373, 2008.
- 14 M. Rapus, S. Munder, G. Barattoff, and J. Denzler. Pedestrian recognition using combined low-resolution depth and intensity images. In Proc. of *IEEE Intelligent Vehicles Symposium*, pp. 632–636, Eindhoven University of Technology Eindhoven, The Netherlands, June 2008.
- 15 C. Munkelt, M. Trummer, P. Kuehnmstedt, G. Notni, and J. Denzler. View planning for 3d reconstruction using time-of-flight camera data. In Proc. of *DAGM 2009*, volume 5748 of *LNCS*, pp. 352–361. Springer, 2009.
- 16 A Frick, B Bartczack, and R Koch. 3d-tv ldv content generation with a hybrid tof-multicamera rig. *3DTV-Conference: The True Vision – Capture, Transmission and Display of 3D Video (3DTV-CON)*, 2010, pp. 1–4, 2010.
- 17 A. Kolb, E. Barth, R. Koch, and R. Larsen. Time-of-flight cameras in computer graphics. *COMPUTER GRAPHICS forum*, vol. 29, no. 1, pp. 141–159, 2010.
- 18 C. Schaller, A. Adelt, J. Penne, and J. Hornegger. Time-of-flight sensor for patient positioning. In Proc. of *SPIE*, volume 7258, 2009.
- 19 C. Schaller, J. Penne, and J. Hornegger. Time-of-Flight Sensor for Respiratory Motion Gating. *Medical Physics*, 35(7):3090–3093, 2008.
- 20 J. Penne, C. Schaller, J. Hornegger, and T. Kuwert. Robust Real-Time 3D Respiratory Motion Detection Using Time-of-Flight Cameras. *Computer Assisted Radiology and Surgery 2008*, 3(5):427–431, 2008.
- 21 A. Kirmani, T. Hutchison, J. Davis, and R. Raskar. Looking around the corner using transient imaging. In Proc. of *ICCV*, pp. 159–166, 2009.

⁵ <http://www.dagstuhl.de/wiki/index.php?title=12431>

4.49 Foundations and Challenges of Change and Evolution in Ontologies

Organizers: James Delgrande, Thomas Meyer, and Ulrike Sattler
Seminar No. 12441

Date: 28. October–02. November, 2012 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.2.10.105

© © © Creative Commons BY-NC-ND 3.0 Unported license
 © James Delgrande, Thomas Meyer, and Ulrike Sattler



Participants: Marcelo Arenas, Franz Baader, Leopoldo Bertossi, Meghyn Bienvenu, Alexander Bochman, Piero Andrea Bonatti, Diego Calvanese, David Carral Martínez, Giovanni Casini, James P. Delgrande, Thomas Eiter, Eduardo Fermé, Giorgos Flouris, Laura Giordano, Birte Glimm, Andreas Herzig, Ian Horrocks, Aaron M. Hunter, Ryutaro Ichise, Gabriele Kern-Isberner, Marcel Lippmann, Carsten Lutz, Robert A. Meersman, Thomas Meyer, Jeff Z. Pan, Peter F. Patel-Schneider, Rafael Penaloza, Marcio Moretto Ribeiro, Riccardo Rosati, Ulrike Sattler, Stefan Schlobach, Matthias Thimm, David Toman, Leon van der Torre, Ivan José Varzinczak, Kewen Wang, Zhe Wang, Renata Wassermann, Grant Weddell, Emil Weydert, Frank Wolter

An ontology in computer science is an explicit, formal specification of the terms of a domain of application, along with the relations among these terms. An ontology provides a (structured) vocabulary which forms the basis for the representation of general knowledge. Ontologies have found extensive application in Artificial Intelligence and the Semantic Web, as well as in areas such as software engineering, bioinformatics, and database systems.

Research in ontologies in Artificial Intelligence has focussed on description logics (DL), where a description logic can be regarded as a (decidable) fragment of first order logic. Historically a DL is divided into two components, a so-called TBox, for expressing concepts and their interrelationships, and an ABox that contains assertions about specific individuals and instances. Thus, the TBox characterises a domain of application while the ABox contains information on a specific instance of a domain. A key point in description logics is that, via their limited expressiveness, one obtains “good”, ideally tractable, inference algorithms. The number of description logics is large, with several prominent families of logics, and the complexity of description logics has been well studied. Research in ontology languages and related reasoning services, most notably in description logics, has also spurred work into logics that are weaker than classical systems, as well provided a substantial impetus for research into modal logic. Moreover, there has been substantial interaction with the database community.

The success of this work has led to an increasing demand for a variety of reasoning services, both classical and non-classical. Crucially, an ontology will be expected to evolve, either as domain information is corrected and refined, or in response to a change in the underlying domain. In a description logic, such change may come in two different forms: the background knowledge, traditionally stored in the TBox, may require modification, or the ground facts or data, traditionally stored in the ABox, may be modified. In the former case, the process is akin to theory revision, in that the underlying background theory is subject to change. In the latter case,

one cannot simply update instances, as is done in a relational database, since any set of instances must accord with the potentially rich structure imposed by the TBox. The result is that one must be able to deal with changing ontologies, as well as related notions from commonsense reasoning, including nonmonotonic reasoning and paraconsistent reasoning.

The issues mentioned are of common interest to the ontology, belief change, and database communities. While there has been some interaction between researchers in these communities, there has not been a comprehensive meeting to address notions of change in ontologies in a broad or comprehensive fashion.

The aim of the workshop was to bring together researchers working in the areas of logic-based ontologies, belief change, and database systems, along with researchers working in relevant areas in nonmonotonic reasoning, commonsense reasoning, and paraconsistent reasoning. Hence the workshop’s goal was to facilitate discussions on the application of existing work in belief change, nonmonotonic reasoning, commonsense reasoning, and related areas on the one hand, to logic-based ontologies on the other. There has been extensive input and interest from the database community, which also has in interest in these problems. Overall the intent was to provide an interdisciplinary (with respect to computer science and mathematics) workshop for addressing both theoretical and computational issues in managing change in ontologies. In particular, the workshop has given participants a deeper understanding of the concepts, terminologies, and paradigms used in the three areas involved, and in their latest achievements and challenges. Examples of these were the distinction between data and schema level, the relation between different revision operators and justifications, the role of less expressive description logics, to name a few.

The workshop consisted of a five-day event with the following program: On the first day there were three introductory talks by a representative in each of the areas of belief change and nonmonotonic reasoning, description logics, and databases. The pur-

pose of these introductory talks was to come to a shared understanding (and terminology) of these areas, and provide a glimpse of the state-of-the-art and current research challenges in all three areas. On day 2, three breakout groups were created and participants were assigned to them based on their expertise but also in such a way as to have representatives of the three main areas in each group. The groups were 'Foundations and Techniques', 'Applications', and 'Perspectives and Future Directions', and their purpose was that of fostering discussions on the three fundamental components at the intersection of the above mentioned areas. Day 3 consisted of a report back from each of the groups followed by further discussion. On the fourth day there were presentations on overlapping areas and discussions of problems and issues of mutual interest for the different communities. Day 5 had a wrap-up session with a discussion on the overlap among the different areas, future challenges and next steps in this workshop series.

4.50 Requirements Management – Novel Perspectives and Challenges

Organizers: Jane Cleland-Huang, Matthias Jarke, Lin Liu, and Kalle Lyytinen
Seminar No. 12442

Date: 28.–31. October, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.10.117

© Creative Commons BY-NC-ND 3.0 Unported license

© Matthias Jarke



Participants: Nicholas Berente, Richard Berntsson Svensson, Jörg Dörr, Xavier Franch, Gilbert Fridgen, Anna Hannemann, Sean Hansen, Julia Heidemann, Jane Cleland-Huang, Jane Huffmann Hayes, Matthias Jarke, John Leslie King, Kim Lauenroth, Julio Cesar Leite, Lin Liu, Kalle Lyytinen, Patrick Mäder, John Mylopoulos, Andreas Oberweis, Barbara Paech, Xin Peng, Barbara Pernici, Balasubramaniam Ramesh, Isabelle Reymen, William N. Robinson, Christoph Rosenkranz, Matti Rossi, Alistair G. Sutcliffe, Fan Yang-Turner, Eric S. Yu, Haiyan Zhao, Andrea Zisman, Didar Zowghi

Since its inception in the 1970s, much of the research in requirements engineering (RE) has focused on the development of formal notations and protocols to represent requirements and to analyze their properties, such as consistency, correctness, completeness, and validity. Some work has analyzed the impacts of these requirements on downstream development tasks (e.g., traceability), or managing and reconciling conflicts in the requirements process. Much of requirements research has also assumed that the scope of RE is isolated to a specific project or even a specific stage of that project. The demand for a shift in focus is dictated by changes in computational paradigms and capabilities that draw upon platform strategies, web services, and virtualization of both application services and development platforms. These trends have significant implications for views of modularity and requirements evolution, complexity of RE tasks, and the economics and costs related to application and service use and development. The aim of the seminar was to bring together experts from multiple fields to discuss models and theories around these changes, focusing on a series of interrelated questions such as:

- How to theorize and study complexity within RE tasks?
- What theoretical perspectives can inform how and why requirements knowledge evolves as it is generated, validated, and distributed?
- How do requirements, system evolution, and environmental change interact?
- How do different types of knowledge interact to shape requirements and their evolution?
- What are the origins and flows of influence of requirements knowledge? How can non-linear influences be effectively managed in RE evolution?
- What is the effect of speed and scale in requirements processes?
- What is the role of goals and constraints and their complex interactions in RE?

In particular we sought better integration of theories of socio-tech-

nical system evolution, distributed cognition, models of RE and design knowledge and their economic effects, the impact of strategy and related knowledge endowments in RE processes (e.g., explorative vs. exploitative processes of requirements discovery), and the role of ambiguity, uncertainty and complexity in managing requirements knowledge. Attention was also placed on new research approaches and methods that can be brought to bear in addressing these problems. The seminar thus built and expanded on some of the critical themes that had been brought up five years earlier in two NSF-sponsored workshops in Cleveland [5] and Dagstuhl [4], [3]. The seminar brought together 33 researchers (exactly one third female) from 12 countries in four continents, with 22% industry participation. Participants felt that this unusually high diversity together with a good mix of junior and senior people of different disciplines, interests and expertise contributed strongly to lively and fruitful discussions. Several cooperative projects have emerged from these discussions. Selected results of the discussions and presentations will be published in a special issue of the ACM Transactions on Management Information Systems in 2014. The program of the seminar was organized into four panels with plenary talks and discussion, five parallel working groups with central reporting, and a final reflection session. With the parallel Dagstuhl seminar on “Foundations and Challenges of Change and Evolution of Ontology” we moreover organized a crossover plenary panel session in which we tried to converge to a better mutual understanding of the different perspectives on Evolution in AI and RE and explored possibilities for future cooperation. Several individual researchers later got together to agree on specific cooperative research. In the final reflection session, the main results, issues and challenges, also taking into account the ontology perspective, can be summarized as follows.

Jackson and Zave [2] have formulated an AI-inspired formalization of the traditional RE viewpoint as a kind of model-based diagnosis: Given a set of domain assumptions D and a set of requirements R , find a suitable specification S such that $S, D \Rightarrow R$.

In RE research, the *R* have often been interpreted as goals, to be refined and satisfied in some extended *AND-OR* graph structure.

From a social science and business informatics perspective, however, requirements engineering (RE) is in essence a boundary spanning task between the developers and the other stakeholders (users, management, regulators, ...) concerning the goals, functions and constraints of a system. The traditional viewpoint, where RE is just seen as an “early phase” (resulting in a contract) and the “last phase” (where acceptance testing takes place) is far too narrow. The following citation by Robert Glass wonderfully characterizes the situation RE has entered since the turn of the century: “*Walking on water, and programming according to specifications is easy – as long as both of them are frozen*” At least three key challenges to research and practice were identified in the seminar, together with counter-strategies where promising first steps for solutions were observed: Firstly, *large-scale projects* encounter changes in *D*, *R*, and the technology underlying *S* is shifting. As a consequence,

- 90% of these projects run over budget and time (this is similar to other big engineering projects, so not a drama in itself)
- One sixth so-called *Black Swan* projects show budget overruns of 70% and time overruns of 200% (this is true only for 1% of other engineering projects, so this is a real drama of software engineering).

A central cause is politically motivated over-ambitious goals with systematic under-estimation of the nature and scope of requirements, budget, and cost both on the side of customers and vendors, as well as poor change tracking. As a consequence, we strongly recommend to not just consider goals of stakeholders but also social structures and strategic dependencies in initial system analysis. Moreover, customer and other stakeholder requirements must be continuously monitored during the development process (and sometimes beyond). To ensure product and process compliance and effectively assess the impact of change, requirements traceability should be focused by using trace patterns to maintain transparency and keep the monitoring effort acceptable and feasible.

Secondly, we need architectural mechanisms that *constrain, but also leverage complexity*. In the seminar, John King pointed out the difference between “complicated” and “complex” problems. Complicated problems can be solved by experienced, highly competent engineers with foreseeable effort. In contrast, complex problems can only be explored with uncertain results; thus, taking on a project that tries to solve a complex problem in one shot is bound to lead to disaster – and apparently, it is exactly the tendency to take on such nice-sounding complex projects that leads to the unusually high share of Black Swans in software projects. Theories like Arthur’s theory

of Technology Evolution [1] or Thornton’s theory of institutional evolution [6] were cited in the seminar as showing a way forward, which we can also observe in practice. Platform strategies offer complicated but manageable base solutions that are now being offered both by open source communities and by big players in different sectors, such as IBM, Google, Facebook, and mobile phone vendors/operators. With a uniform infrastructure, they limit complexity. But by enabling innovation at the margin, e.g. end-user developed apps, they at the same time also leverage new complexity at the higher level. The easy entry, combined with ruthless selection of a very small percentage of truly successful apps, then offers a hotbed of complex evolutionary change. Which eventually will grow into, or be replaced again by yet another layer of platforms, as pointed out by both Brian Arthur and the earlier book by Thomas Friedman “The World is Flat”. Beyond such market selection mechanisms, software vendors employ various mechanisms to participate in this game in a more controlled way. We mention here the now broad area of software product families, but also Google’s 70 : 20 : 10 work rule where employees are free to spend a significant part of their work time on their own ideas, thus fostering continuous internal innovation. Methods for runtime requirements monitoring and requirements mining from usage patterns can be important contributions of the RE field in this context. Last not least, the future will not reduce the challenges of complexity and evolution. Our seminar understood information systems as socio-technical systems, but in fact many of today’s systems are neither truly social nor truly technical. From a social perspective, there is the new question for sustainability of systems, with the demand for re-optimization from the viewpoint of user rights (e.g. asymmetric information and market powers, privacy, data ownership, copyright vs. freedom of information), energy efficiency, and environmental footprint. From a technical perspective, the explosive expected growth of Cyberphysical Systems (Internet of Things) in business, engineering and science is not just an approach to monitor and actuate at a much more fine-grained level, but a significant source of more complexity and evolutionary challenges. Generating and implementing e.g. the visions of smart cities is just but one example. Rather than just talking grand new visions here, methods for how to get there step-by-step in an “only” complicated way – without exposing whole city to the chaos caused by over-ambitious “complex” systems – are urgently needed. In our world of more and more ubiquitous computing, where the impact and complexity of systems continuously seem to grow, *RE is the marketplace where responsibility is traded*, as communication, mutual understanding, and transparent well-structured information management are at the heart of this field.

References

- 1 B. Arthur (2009). *The Nature of Technology: What it is and how it Evolves*. Free Press.
- 2 M. Jackson, P. Zave (1995): *Deriving specifications from requirements: an example*. Proc. 17th ICSE.
- 3 M. Jarke, K. Lyytinen, eds. (2010): *High Impact Requirements Engineering*. Special Issue, *Wirtschaftsinformatik/BISE* 52, 3.
- 4 M. Jarke, P. Loucopoulos, K. Lyytinen, J. Mylopoulos, W.N. Robinson (2011): *The brave new world of design requirements*. *Information Systems* 36(7): 992–1008.
- 5 K. Lyytinen, P. Loucopoulos, J. Mylopoulos, W.N. Robinson, eds. (2009). *Design Requirements Engineering – A Ten-Year Perspective*. Springer LNBIP 14.
- 6 P. Thornton, W. Occasio, M. Lounsbury (2012). *The Institutional Logics Perspective: A New Approach to Culture, Structure, and Process*. Oxford University Press.

4.51 The Constraint Satisfaction Problem: Complexity and Approximability

Organizers: Johan Håstad, Andrei Krokhin, and Dániel Marx
Seminar No. 12451

Date: 4.–9. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Johan Håstad, Andrei Krokhin, and Dániel Marx



Participants: Albert Atserias, Per Austrin, Libor Barto, Manuel Bodirsky, Andrei A. Bulatov, Catarina Carvalho, Hubie Chen, David Cohen, Victor Dalmau, Martin Dyer, Leslie Ann Goldberg, Martin Grohe, Venkatesan Guruswami, Gregory Z. Gutin, Johan Hastad, Sangxia Huang, Anna Huber, Mark Jerrum, Peter Jonsson, Vladimir Kolmogorov, Marcin Kozik, Stefan Kratsch, Andrei Krokhin, Benoit Larose, Pin-Yan Lu, Konstantin Makarychev, Yury Makarychev, Petar Markovic, Barnaby Martin, Dániel Marx, Ralph McKenzie, Michael Pinsker, Prasad Raghavendra, Francesco Scarcello, Ola Svensson, Stefan Szeider, Suguru Tamaki, Johan Thapper, Matt Valeriote, Moshe Y. Vardi, Magnus Wahlström, Ross Willard, Yuichi Yoshida, Stanislav Zivny

The *constraint satisfaction problem*, or CSP in short, provides a unifying framework in which it is possible to express, in a natural way, a wide variety of computational problems dealing with mappings and assignments, including satisfiability, graph colorability, and systems of equations. The CSP framework originated 25-30 years ago independently in artificial intelligence, database theory, and graph theory, under three different guises, and it was realised only in the late 1990s that these are in fact different faces of the same fundamental problem. Nowadays, the CSP is extensively used in theoretical computer science, being a mathematical object with very rich structure that provides an excellent laboratory both for classification methods and for algorithmic techniques, while in AI and more applied areas of computer science this framework is widely regarded as a versatile and efficient way of modelling and solving a variety of real-world problems, such as planning and scheduling, software verification and natural language comprehension, to name just a few. An instance of CSP consists of a set of variables, a set of values for the variables, and a set of constraints that restrict the combinations of values that certain subsets of variables may take. Given such an instance, the possible questions include (a) deciding whether there is an assignment of values to the variables so that every constraint is satisfied, or optimising such assignments in various ways, or (b) finding an assignment satisfying as many constraints as possible. There are many important modifications and extensions of this basic framework, e.g. those that deal with soft or global constraints.

Constraint satisfaction has always played a central role in computational complexity theory; appropriate versions of CSPs are classical complete problems for most standard complexity classes. CSPs constitute a very rich and yet sufficiently manageable class of problems to give a good perspective on general computational phenomena. For instance, they help to understand which mathematical properties make a computational problem tractable (in a wide sense, e.g. polynomial-time solvable or non-trivially approximable,

fixed-parameter tractable or definable in a weak logic). It is only natural that CSPs play a role in many high-profile conjectures in complexity theory, exemplified by the Dichotomy Conjecture of Feder and Vardi and the Unique Games Conjecture of Khot.

The recent flurry of activity on the topic of the seminar is witnessed by two previous Dagstuhl seminars, titled “Complexity of constraints” (06401) and “The CSP: complexity and approximability” (09441), that were held in 2006 and 2009, respectively. This seminar was a follow-up to the 2009 seminar. Indeed, the exchange of ideas at the 2009 seminar has led to new ambitious research projects and to establishing regular communications channels, and there is a clear potential of a further systematic interaction that will keep on cross-fertilizing the areas and opening new research directions. The 2012 seminar brought together forty four researchers from different highly advanced areas of constraint satisfaction and involved many specialists who use universal-algebraic, combinatorial, geometric and probabilistic techniques to study CSP-related algorithmic problems.

The seminar included two substantial tutorials: one on the classification of the complexity of constraint languages via methods of logic and universal algebra (given by A. Krokhin from Durham U, UK and R. Willard from Waterloo U, CA), and the other on the approximability of CSP (given by P. Austrin from KTH Stockholm, SE). Other participants presented, in 28 further talks, their recent results on a number of important questions concerning the topic of the seminar.

Concluding Remarks and future plans. The seminar was well received as witnessed by the high rate of accepted invitations and the great degree of involvement by the participants. Because of the multitude of impressive results reported during the seminar and the active discussions between researchers with different expertise areas, the organisers regard this seminar as a great success. With steadily increasing interactions between such researchers, we

foresee a new seminar focussing on the interplay between different approaches to studying the complexity and approximability of the CSP. Finally, the organisers wish to express their gratitude to the Scientific Directors of the Dagstuhl Centre for their support of the seminar.

■ Description of the Topics of the Seminar

Classical computational complexity of CSPs. Despite the provable existence of intermediate (say, between P and NP-complete, assuming $P \neq NP$) problems, research in computational complexity has produced a widely known informal thesis that “natural problems are almost always complete for standard complexity classes”. CSPs have been actively used to support and refine this thesis. More precisely, several restricted forms of CSP have been investigated in depth. One of the main types of restrictions is the *constraint language* restriction, i.e., a restriction on the available types of constraints. By choosing an appropriate constraint language, one can obtain many well-known computational problems from graph theory, logic, and algebra. The study of the constraint language restriction is driven by the CSP *Dichotomy Conjecture* of Feder and Vardi which states that, for each fixed constraint language, the corresponding CSP is either in P or NP-complete. There are similar dichotomy conjectures concerning other complexity classes (e.g. L and NL). Recent breakthroughs in the complexity of CSP have been made possible by the introduction of the universal-algebraic approach, which extracts algebraic structure from the constraint language and uses it to analyse problem instances. McKenzie’s talk surveyed classes of algebras that arise in this context and Pinsker related this approach with infinite-valued CSPs. The algebraic approach has been applied to prove the Dichotomy Conjecture in many important special cases (e.g. Bulatov’s dichotomy theorems for 3-valued and conservative CSPs), but the general problem remains open. A powerful universal-algebraic theory of absorption has been developed in the last couple of years by Barto and Kozik, specifically motivated by CSP classification questions. This theory has already produced several spectacular classification results resolving long-standing problems (including a characterization of CSPs of bounded width, i.e. solvable by local propagation algorithms), and there is a clear sense that there is much more to come from it. Kozik presented new results on CSPs in NL that are based on the absorption theory.

Algebraic approaches to studying exact exponential and sublinear algorithms for CSPs were presented by Jonsson and Yoshida, respectively.

The complexity of Valued CSPs, which are a significant generalisation of Max CSP, was considered in the talks by Huber, Kolmogorov, Thapper, and Živný. Very strong results were reported, especially the full description of tractable cases by Thapper and Živný. Raghavendra presented results that might lead to closer interchange of ideas between algebraic and probabilistic approaches to CSPs.

The algebraic approach to the complexity of counting solutions for CSPs, with many results, was presented by Bulatov, Dyer, Goldberg, and Jerrum, while Lu reported recent progress on classifying the complexity of Holant problems.

Approximability of CSPs. The use of approximation algorithms is one of the most fruitful approaches to coping with NP-hardness. Hard optimization problems, however, exhibit different behavior with respect to approximability, making it an exciting, and by now, well-developed but far from fully understood, research area. The CSP has always played an important role in the study of approximability. For example, it is well known that the famous PCP theorem has an equivalent reformulation in terms of inapproximability of a certain CSP; moreover, the recent combinatorial proof of

this theorem by Dinur in 2006 deals entirely with CSPs. The first optimal inapproximability results by Håstad in 2001 were about certain CSPs, and they led to the study of a new hardness notion called *approximation resistance* (which, intuitively, means that a problem cannot be approximated beyond the approximation ratio given by picking an assignment uniformly at random, even on almost satisfiable instances). Many CSPs have been classified as to whether they are approximation resistant but there is not even a reasonable conjecture for a full classification. Håstad, Huang, and K. Makarychev presented new results on approximation resistance.

In a related development, Guruswami and Zhou have discussed, in 2010, a “hybrid” form of tractability for CSPs, where classical tractability is combined with good approximability on almost satisfiable instances, and they conjecture that CSPs of bounded width have this desirable property. This conjecture was proved by Barto and Kozik in 2012 (and presented by Barto at the seminar), with further results in this direction presented by Dalmau.

Arguably, the most exciting development in approximability in the past five to six years is the work around the *unique games conjecture* (UGC), which was introduced by Khot in 2002. It states that, for CSPs with a certain constraint language over a large enough domain, it is NP-hard to tell almost satisfiable instances from those where only a small fraction of constraints can be satisfied. This conjecture (if true) is known to imply optimal inapproximability results for many classical optimization problems. Moreover, if the UGC is true then, as shown by Raghavendra in 2008, a simple algorithm based on semidefinite programming provides the best possible approximation for all CSPs (though the exact quality of this approximation is unknown). In 2010, Arora *et al.* gave a sub-exponential time algorithm for unique games CSPs, which is based on a new graph decomposition method. This does not give strong evidence in favor or against the conjecture, but it shows that there are important new algorithmic ideas to be discovered. Y. Makarychev presented an asymptotically optimal (modulo UGC) approximation algorithm for the general Max CSP.

Parameterized complexity of CSPs. A different way to cope with NP-hardness is provided by parameterized complexity, which relaxes the notion of tractability as polynomial-time solvability to allow non-polynomial dependence on certain problem-specific parameters. A whole new set of interesting questions arises if we look at CSPs from this point of view. Most CSP dichotomy questions can be revisited by defining a parameterized version; so far, very little work was done in this direction compared to investigations in classical complexity. Interestingly, some of the most tantalizing open problems in parameterized algorithmics (e.g. the fixed-parameter tractability of the BCLIQUE problem) are directly related to complexity of CSPs, and Marx’s talk contained an overview of such problems. A new research direction (often called “parameterizing above the guaranteed tight bound”) led to unexpected positive results for Max r -SAT by Alon *et al.* in 2010. In this direction, the basic question is to decide the fixed-parameter tractability of the following type of problems: if we know that a random assignment satisfies at least E clauses/constraints in expectation (and hence such an assignment is easy to find), find an assignment that satisfies at least $E + k$ clauses/constraints. Gutin presented recent results in this direction.

Along with the constraint language restriction, another important restriction of CSPs that has been thoroughly investigated is the *structural* restriction, where the way in which the immediate interaction between variables in instances is restricted. In this direction, the notions of (hyper)graph decompositions and treewidth turned out to be particularly important. These notions are core concepts of parameterized algorithmics, and so, it is not surprising that parameterized complexity is an important tool in characterizing structural

restrictions that lead to tractable CSPs. In particular, many known classification results with respect to classical complexity in this direction (e.g. Grohe, 2007) use tools from parameterized complexity. Scarcello and Szeider described their new results in this direction.

Logic and the complexity of CSP. Starting from earlier work by Kolaitis and Vardi, concepts and techniques from logic have provided unifying explanations for many tractable CSPs. This has led to the pursuit of classifications of CSP with respect to *descriptive complexity*, i.e. definability in a given logic. Logics considered in this context include first order logic and its extensions, finite-variable logics, the logic programming language Datalog and its fragments. Kozik's talk described a contribution in this direction.

The CSP can be recast as the problem of deciding satisfiability of existential conjunctive formulas. Natural extensions of this framework that allow counting or universal quantifiers were considered in the talks by Martin and Chen, respectively. Atserias' talk related proof complexity, CSPs, and semidefinite programming.



Fig. 4.16

Simone Litschka – ohne Titel Part of the Dagstuhl art collection and donated by Rotary Club Saarbrücken, Architekturbüro Birtel, participants of AVACS-Klausurtagung (10362), and participants of the Dagstuhl Seminars 10071 and 10361.

4.52 Publication Culture in Computing Research

Organizers: Kurt Mehlhorn, Moshe Y. Vardi, and Marc Herbstritt
Seminar No. 12452

Date: 06.–09. November, 2012 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.2.11.20

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Kurt Mehlhorn, Moshe Y. Vardi, and Marc Herbstritt



Participants: Andrew P. Bernat, Ronald F. Boisvert, George Danezis, Lance Fortnow, Batya Friedman, Maarten Fröhlich, Jonathan Grudin, Marc Herbstritt, Manuel Hermenegildo, Alfred Hofmann, Nicolas Holzschuch, Srinivasan Keshav, Ursula Martin, Keith Marzullo, Friedemann Mattern, Kurt Mehlhorn, Bertrand Meyer, Jeff Mogul, M. Tamer Özsu, José Palazzo Moreira de Oliveira, Sweitze Roffel, Vladimiro Sassone, Fred B. Schneider, Douglas B. Terry, Jan van Leeuwen, Moshe Y. Vardi, Andrei Voronkov, Dan Wallach, Reinhard Wilhelm

The dissemination of research results is an integral part of research and hence a crucial component for any scientific discipline. While computing research has been phenomenally successful, there is a broad feeling that the publication models are quite often obstacles. Yet there is no agreement on whether the publication models need to be radically changed or fine-tuned, and there is no agreement on how such change may occur. Over the past few years, a vigorous discussion has been going on through editorials, Viewpoint articles, and blogs of the Communication of the ACM – see Jonathan Grudin’s overview available at <http://research.microsoft.com/en-us/UM/People/jgrudin/publications/publicationculture/CACMreferences.pdf>.

In spite of this ongoing debate, the community seems no closer to an agreement whether a change has to take place and how to effect such a change.

The workshop brought together key players in this debate for an intense three-day discussion and deliberation, with the aim of analyzing the issues and developing guidelines for the way forward. A specific focus of the workshop was to develop consensus around a set of guiding principles. An expected outcome of the workshop is a manifesto to be published afterwards.

■ Topics

The workshop addressed several topics that were part of the community’s collective conversation on publication culture during the last years:

1. The uniqueness of the publication model in computing research:
 - the emphasis on conference publishing and the decline of journal publishing;
 - the large and growing number of specialty conferences and workshops that are really conferences;
 - coping with established publication cultures in the (other)

sciences and with the different cultures of different computing sub-communities.

2. Cultural issues:

- the culture of hypercritical reviewing and the decline of thorough constructive reviewing;
- tenure and promotion practices that encourage short-term research;
- the influence of bibliometry on publication behavior and tenure practices and the quality of bibliometry.

3. New publication models:

- the tension between open access and reader-pays publishing, and the spectrum in between;
- the role of social media in scholarly publishing;
- the role of various actors: commercial publishers, scientific societies, academic publishers and archives;
- the place of self-publishing or publishing in public repositories;
- the need to develop new rules for data citation, sharing, and archiving.

■ Organization

The workshop was organized by Moshe Y. Vardi and Kurt Mehlhorn with coordinating support by Marc Herbstritt. Additionally, a program committee (PC) was set up, including Andrew P. Bernat, Jon Crowcroft, Jan van Leeuwen, Bertrand Meyer, Fred B. Schneider, and Douglas B. Terry. The PC helped in seeking suitable contributions and advising the organizers in shaping the program. Each invitee was asked to submit a position statement which was reviewed by the organizers and the PC. The collection of accepted position statements provided a broad and concise overview of the problems in the publication culture of computing research, disclosing a variety of different and competing viewpoints.

On Wednesday Nov. 7, 2012, the workshop started with a sessi-

on presenting standpoints from scholarly societies and commercial publishers, among them Ronald Boisvert (NIST/ACM), Dan Walach (Rice University/USENIX), Maarten Fröhlich (IOS Press), Alfred Hofmann (Springer Science Business+Media/LNCS), Sweitze Roffel (Elsevier), Andrew Bernat (Computing Research Association), and Moshe Y. Vardi (Rice University/ Editor-in-Chief of Comm. of the ACM). The afternoon session focussed on peer review and research dissemination, including the talks from Bertrand Meyer (ETH Zürich/Informatics Europe), Ursula Martin (Queen Mary University London), Lance Fortnow (Georgia Institute of Technology), Doug Terry (Microsoft – Mountain View), Nicolas Holzschuch (INRIA Rhône-Alpes), George Danezis (Microsoft Research – Cambridge), and José Palazzo Moreira de Oliveira (UFRGS).

On Thursday Nov. 8, 2012, the workshop continued with a morning session on “conferences versus journals” as well as on “open access”, with talks from Manuel Hermenegildo (IMDEA), Keith Marzullo (NSF), Kurt Mehlhorn (MPII), Jeff Mogul (HP), M. Tamer Özsu (University of Waterloo), and Vladimiro Sassone (University of Southampton). The afternoon session focussed also “conferences versus journals”, but also on indexing and general cultural issues; talks were given by Reinhard Wilhelm (Saarland University), Jan van Leeuwen (Utrecht University), Jonathan Grudin (Microsoft Research – Redmond), Andrei Voronkov (Manchester University), Srinivasan Keshav (University of Waterloo), Fred B. Schneider (Cornell University), and Batya Friedman (University of Washington).

Batya Friedman moderated the “Future Workshop”, which (1) interactively asked participants after the sessions to contribute brief descriptions of substantial shortcomings in our current publication culture, according to one’s own opinion, (2) asked participants to describe an idealized publication culture for computing research, and (3) finally, asked participants to provide brief accounts of potential solutions to the problems raised and ways to reach ideal outcomes.

The results of the “Future Workshop” were discussed on Friday Nov. 9, 2012, and served as basis for working groups. The working groups met in small teams and presented the results from their discussions to the audience. Finally, Moshe Y. Vardi gave a summary on the workshop and talked about future actions.

The organizers and the PC met on Friday afternoon to clarify core issues for the upcoming manifesto.

■ Outcomes

The main outcomes will be covered in the upcoming manifesto that will be published in the “Dagstuhl Manifestos” series⁶. However, as discussed during the organizers and PC meeting on Friday afternoon, a first sketch of a consensus list with regard to problems and desired solutions is as follows:

- Problems:
 - *Scaling*: The publishing ecosystem in computing research—conference and journals—has not scaled up.
 - *Policy*: We have no universally accepted norms and policies, and no single authority.
 - *Data*: We have many opinions but little data.
 - *Business model*: Huge gap between publishers and authors/readers.
 - *Incentives*: Large number of small papers.
 - *Measurements*: Highly imperfect metrics.
 - *Conferences*: Too many submissions and resubmissions, random-like decisions, too many conferences, too much travel, conferences as “journals that meet in hotels”.

- *Journals*: Not exciting, hard to find reviewers, poor journal management systems.
- *Reviewing*: Increasing burden, declining standards, declining competence.
- Wish list:
 - *Defragmented Community*: Learn to operate at scale.
 - *Rational reviewing*: Eliminate treadmill, eliminate hyper-criticality, reintroduce review rounds.
 - *Revitalized journals*: Perhaps through jourference/cournal hybrids.
 - *Reduce paper inflation*: Focus on quality, not quantity.
 - *Appropriate bibliometrics*: Recognize conferences, eliminate self-citation.
 - *Open Access*: Research results should be available to all to read and assess as soon as possible.
 - *Viable associations*: Strong associations that can enable, facilitate, and lead a better publication culture.

■ Resources

Position statements and slides from the presentations are available at <http://www.dagstuhl.de/mat/index.en.phtml?12452>.

⁶ <http://drops.dagstuhl.de/dagman>

4.53 Games and Decisions for Rigorous Systems Engineering

Organizers: Nikolaj Bjørner, Krishnendu Chatterjee, Laura Kovacs, and Rupak M. Majumdar
Seminar No. 12461

Date: 11.–16. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.45

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Nikolaj Bjørner, Krishnendu Chatterjee, Laura Kovacs, and Rupak M. Majumdar



Participants: Francesco Alberti, Dirk Beyer, Nikolaj Bjørner, Tomas Brazdil, Krishnendu Chatterjee, Swarat Chaudhuri, Laurent Doyen, Uwe Egly, Azadeh Farzan, Bernd Finkbeiner, Alberto Griggio, Aarti Gupta, Ashutosh Kumar Gupta, Arie Gurfinkel, Jochen Hoenicke, Radu Iosif, Moa Johansson, Igor Konnov, Konstantin Korovin, Laura Kovacs, Axel Legay, Rupak Majumdar, Anca Muscholl, Albert Oliveras, Joel Ouaknine, David Parker, Brigitte Pientka, Ruzica Piskac, Albert Rubio, Andrey Rybalchenko, Helmut Seidl, Martina Seidl, Natasha Sharygina, Ana Sokolova, Armando Solar-Lezama, Eugenia Ternovska, Helmut Veith, Tomas Vojnar, Andrei Voronkov, Georg Weissenbacher, Josef Widder, Thomas Wies, Florian Zuleger

Principled approaches to systems design offer several advantages, including developing safety-critical systems and scaling technological advances with multi-core processes and cloud computing. Rigorous mathematical techniques, such as model checking, decision procedures, and abstract interpretation, are dominantly used a posteriori in systems engineering: a program is formally analyzed after it has been developed. In the context of rigorous systems engineering, post-hoc verification is however very costly and error-prone. The explosion of concurrent computation in the new generation of embedded systems has therefore motivated the integration of established methods with novel techniques in the design process from day one.

Such an integration has been materialized in using game theoretic synthesis of reactive systems from higher level design requirements. In many synthesis algorithms, it is better to work with symbolic representations, where the state space is modeled using logical formulas. This enables techniques to scale to potentially infinite models, but requires decision procedures for checking the validity of sentences in the pertinent logical theories. The increasingly complex integration of model checking with complementary techniques such as software testing has imposed new requirements on decision procedures, such as proof generation, unsatisfiable core extraction, and interpolation.

The main goal of the Dagstuhl Seminar 12461 “Games and Decisions for Rigorous Systems Engineering” was to bring together researchers working in the field of rigorous systems engineering, the tool-supported application of mathematical reasoning principles to the design and verification of complex software and hardware systems. The seminar had a special focus on developing systems (reactive, concurrent, distributed) using recent advances in game theoretic synthesis and in decision procedures and automated deduction techniques.

The seminar covered the following three main areas:

- software verification (reactive, concurrent, distributed);

- game theory and reactive synthesis;
- decision procedures (SAT, SMT, QBF) and theorem proving (first and higher order).

Within the scope of these areas, the seminar addressed tooling around software testing, model checking, interpolation, decision procedures, and model finding methods in automated theorem proving.

In the spirit of advancing tools and theory in related areas of theorem proving and model checking, the seminar schedule included tutorials on games, synthesis, theorem proving; research talks on recent results; and discussion sessions on applications and exchange formats for benchmarking tools.

The seminar fell on 5 days in the week of November 12–16, 2012. All together, 43 researchers participated (11 women and 32 men).

4.54 Symbolic Methods for Chemical Reaction Networks

Organizers: Francois Boulier, Anne J. Shiu, Thomas Sturm, and Andreas Weber
Seminar No. 12462

Date: 11.–16. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.66

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Anne J. Shiu

Participants: Francois Boulier, Christopher W. Brown, Carsten Conradi, Gheorghe Craciun, Alicia Dickenstein, Andreas Eggers, Ralf Markus Eiswirth, Hassan Errami, François Fages, Elisenda Feliu, Holger Fröhlich, Vladimir Gerdt, Manoj Gopalkrishnan, Dima Grigoriev, Daniel Kaschek, Marek Kosta, Francois Lemaire, Stefan Müller, Sabine Peres, Adrien Poteaux, Nicole Radde, Georg Regensburger, Satya Swarup Samal, Karsten Scheibler, Igor Schreiber, Stefan Schuster, Werner M. Seiler, Anne J. Shiu, Thomas Sturm, Jan H. van Schuppen, Andreas Weber



Systems of differential equations and hybrid systems more generally are prevalent in chemical engineering and systems biology. The analysis of such systems focuses on resolving their dynamical properties, for instance, determining their equilibria and capacity for multistationarity or Hopf bifurcations. The additional tasks of parameter estimation, model reduction, and model inference are also relevant for these systems. These goals are difficult in general, especially due to the large size of these systems, especially those arising in systems biology. Non-numeric methods are essential in this context, because reaction parameters can vary over a wide range, parameter uncertainty is predominant in systems biology, and even the qualitative behavior of a system typically varies among different regions of the parameter space. Two major lines of research in this area are represented by chemical reaction network theory and stoichiometric network analysis. Our Dagstuhl seminar brought together researchers from both of these research areas, as well as researchers in symbolic computation and those on the application side (chemical engineering and systems biology). The aim of our seminar was twofold: to introduce practitioners to existing relevant theory and software from symbolic computation, and to allow participants to pose current computational challenges in this area, in order to spur development of symbolic computation methods to resolve these problems. To this end, collaborative working groups on various related topics were held throughout the week.

On Monday during the seminar, most of the participants gave short talks introducing their research interests. On Tuesday, long talks were given by Gheorghe Craciun and Francois Fages. A long talk on Wednesday was given by Stefan Schuster. Markus Eiswirth and Holger Fröhlich gave long talks on Thursday. Discussion groups met throughout the week.

4.55 SAT Interactions

Organizers: Nadia Creignou, Nicola Galesi, Oliver Kullmann, and Heribert Vollmer
Seminar No. 12471

Date: 18.–23. November, 2012 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.2.11.87

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Nadia Creignou and Heribert Vollmer



Participants: Olaf Beyersdorff, Uwe Bubeck, Catarina Carvalho, Nadia Creignou, Stefan Dantchev, Evgeny Dantsin, Anuj Dawar, Arnaud Durand, Johannes Ebbing, Uwe Egly, John Franco, Nicola Galesi, Heidi Gebauer, Andreas Goerdt, Miki Hermann, Timon Hertli, Edward A. Hirsch, Kazuo Iwama, Jan Johannsen, Lefteris M. Kirousis, Hans Kleine Büning, Donald Ervin Knuth, Juha Kontinen, Alexander S. Kulikov, Oliver Kullmann, Massimo Lauria, Victor W. Marek, Barnaby Martin, Arne Meier, Julian-Steffen Müller, Jakob Nordström, Ramamohan Paturi, Rahul Santhanam, Dominik Scheder, Henning Schnoor, Uwe Schöning, Martina Seidl, Robert H. Sloan, Ewald Speckenmeyer, Stefan Szeider, Jacobo Toran, Heribert Vollmer, Sean Weaver, Xishun Zhao

■ Brief Introduction to the Topic

Propositional satisfiability (or Boolean satisfiability) is the problem of determining whether the variables of a Boolean formula can be assigned truth values in such a way as to make the formula true. The satisfiability problem, SAT for short, stands at the crossroads of logic, graph theory, computer science, computer engineering and computational physics.

In particular SAT is of central importance in various areas of computer science including algorithmics, verification, planning and hardware design. It can express a wide range of combinatorial problems as well as many real-world ones. Due to its potential practical implications an intensive search has been done on how one could solve SAT problems in an automated fashion. The last decade has seen the development of practically-efficient algorithms for SAT, which can solve huge problems instances.

At the same time SAT is very significant from a theoretical point of view. Since the Cook-Levin's theorem, which has identified SAT as the first NP-complete problem, it has become a reference for an enormous variety of complexity statements. The most prominent one is the question "is P equal to NP?" Proving that SAT is not in P would answer this question negatively. Indeed, as stated by Richard Lipton on his blog *Gödel's Lost Letter and P = NP* (<http://rjlipton.wordpress.com>) such a proof matters since it would tell us why some computational problems are intrinsically more difficult than others, it would suggest new methods that would yield new insights on the fundamental nature of computation and it would help with goals of security for cryptographers.

During the past two decades, an impressive array of diverse techniques from mathematical fields, such as propositional logic, model theory, Boolean function theory, combinatorics, probability, and statistical physics has contributed to a better understanding of the SAT problem. Although significant progress has been made on several fronts, most of the central questions remain unsolved so far. One of the main aims of the Dagstuhl Seminar was to bring together researchers from different areas of activity in SAT (with an

emphasize on mathematical aspects), so that they can communicate state-of-the-art advances and embark on a systematic interaction that will enhance the synergy between the different areas.

■ Organization of the Seminar and Activities

The workshop brought together 44 researchers from different areas of computer science and mathematics such as logic, complexity theory, algorithms, and proof complexity with complementary expertise. 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. Twenty-five talks of various lengths took place over the five days of the workshop. Introductory and tutorial talks of 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:

- Olaf Beyersdorff, Proof complexity
- Arne Meier, Complexity classifications for different satisfiability problems
- Victor Marek, Erdős' dream; SAT and combinatorics
- Uwe Bubeck, Quantified Boolean formulas: complexity and expressiveness
- Oliver Kullmann, The combinatorics of minimal unsatisfiability
- Martina Seidl, QBF solvers

Most of the tutorials were given by young researchers, reflecting the fact that the SAT community is dynamic and fast evolving.

A highlight of the seminar was the talk by Donald E. Knuth, delivered Wednesday morning, on “Satisfiability and the Art of Computer Programming”. Knuth reported about his experiences while working on a chapter on satisfiability for the upcoming volume of his world-renowned series.

There were additionally 19 shorter talks. These talks covered a wide range of topics related to satisfiability. The different approaches discussed above in the seminar description were all very well represented by the different talks given during the five days of the seminar.

1. Combinatorics
 - Xishun Zhao, Finiteness conjecture on hitting minimal unsatisfiable formulas
 - Uwe Schöning, Probability distributions for local search and make versus break
 - Heidi Gebauer, Applications of (k, d) -trees
2. Complexity
 - Juha Kontinen, Dependence logic and complexity
 - Julian-Steffen Müller, A fragment of dependence logic characterizing PTIME
 - Alexander Kulikov, New lower and upper bounds for Boolean circuit complexity
 - Johannes Ebbing, Model checking for modal intuitionistic dependence logic
3. Proof complexity
 - Uwe Egly, Proof complexity for QBF
 - Jan Johannsen, Separating clause learning proof systems from (regular) resolution
 - Jakob Nordström, Relating proof complexity measures and practical hardness of SAT
 - Massimo Lauria, Open problems in proof complexity
4. Algorithms
 - Stefan Szeider, Fixed-parameter tractability and SAT
 - Mohan Paturi, Algorithmic expressivity and hardness of satisfiability
 - Dominik Scheder, Exponential lower bounds for the PPSZ k -SAT algorithm
5. Logic
 - Arnaud Durand, A criterion for tractability of counting solutions to uniform CSP
 - Hans Kleine Büning, On some configuration problems based on representations in propositional logic
6. Solvers
 - John Franco, Adding unsafe constraints to improve satisfiability performance
 - Sean Weaver, Satisfiability enhancements enabled by state machines

This classification is necessarily rough, as several talks crossed the boundaries between these areas, in keeping with the theme of the workshop. The broad scope of the talks extended even to areas not anticipated by the organizers, such as dependence logic. The workshop thus achieved its aim of bringing together researchers from various related communities to share state-of-the-art research.

■ Concluding Remarks and Future Plans

The organizers regard the workshop as a great success. Bringing together researchers from different areas of theoretical computer science 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 *SAT Interactions* in the future.

4.56 Is the Future of Preservation Cloudy?

Organizers: Erik Elmroth, Michael Factor, Ethan Miller, and Margo Seltzer
Seminar No. 12472

Date: 19.–21. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.102

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Erik Elmroth, Michael Factor, Ethan Miller, and Margo Seltzer



Participants: Ian F. Adams, Jean Bacon, Mary Baker, Christoph Becker, André Brinkmann, Nikos Chondros, Milena Dobrova, Erik Elmroth, Michael Factor, Sam Fineberg, David Giaretta, Matthias Grawinkel, Alexandru Iosup, Ross King, Hillel Kolodner, Ewnetu Bayuh Lakew, Natasa Milic-Frayling, Ethan Miller, Dirk Nitschke, Gillian Oliver, Peter R. Pietzuch, David S. H. Rosenthal, Raivo Ruusalepp, Gerhard Schneider, Margo Seltzer, Liuba Shrira, Joanne Syben, Lawrence You

Two significant trends in data management are emerging: data is moving to cloud infrastructures and an increasing fraction of data produced is born digital. We risk losing all record of born digital data if we do not take explicit steps to ensure its longevity. While each of these trends raises its own set of questions, our seminar began with two fundamental questions at the intersection of these trends: What role should the cloud play in preservation? What steps should we be taking now to preserve the future of today's digital artifacts?

We addressed these two questions by bringing together a diverse cohort of approximately thirty participants. Our participants consisted of researchers from both academia and industry, representatives from cloud providers, and archivists and librarians from memory institutions. Every participant was responsible for some aspect of the program, and the workshop was characterized by lively debate. There were four primary outcomes of the workshop:

1. We identified key functional requirements that are critical if cloud infrastructures are to be used for long-term digital preservation.
2. We identified topics where we were unable to reach agreement; since we are trying to look into the future, while not satisfactory, it seems likely we will need to wait until the future to resolve these debates.
3. We identified several specific problems requiring further work and brought together groups of people interested in pursuing those areas.
4. We identified several areas that we were not able to address, either because we lacked the expertise in the room or we ran out of time; these areas represent opportunities for subsequent workshops.

Perhaps the most pressing issue with respect to existing cloud infrastructures is the lack of standardized APIs. If data are to outlive any particular organization, then it is crucial that archives span organizational boundaries; standardized APIs make this dramatically easier and more robust. There was also agreement that some form

of automated appraisal was important, but there were no concrete ideas about how to do it.

We had lively debate around the long term cost of cloud storage, in particular public clouds; since this debate depended upon assumptions of future costs, the future will ultimately resolve the debate. We also had much discussion around the importance of logical preservation and whether the modern world, with readily available open source viewers has made the need for logical preservation obsolete.

Several small working groups coalesced around the areas of: archival exit (how do you get data out of an archive), the technical design of preservation-as-a-service (PaaS), technologies for ensuring that data is "forgotten", and searching distributed archives. We are hoping to see these small groups evolve into productive collaborations that continue the work begun at the seminar.

Finally, there were a number of areas related to using the cloud as a preservation service that we were unable to address. For example, what legal issues arise if companies undertake digital archival initiatives? Is there a legal definition of "deletion" of data, and is it practical? Where does "record management" end and "archival" begin? Who is the customer for long term preservation? Is it the data provider? Or perhaps it's the data consumer? What happens to archived data if payment cannot be made? What is the economic model behind long term archival? These and other questions provide ample opportunity for further workshops on this topic.

■ Organization

The workshop was organized around a series of 90-minute sessions, each of which began with one or more short presentations followed by a moderated discussion. We had one person scribe each session and the session moderators produced the session summaries that appear in this report documenting each session. We also devoted one session to smaller breakout groups, who reported back in our closing session.

4.57 Quantitative Security Analysis

Organizers: Boris Köpf, Pasquale Malacaria, and Catuscia Palamidessi
Seminar No. 12481

Date: 20.–25. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.135

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Boris Köpf, Pasquale Malacaria, and Catuscia Palamidessi

Participants: Alessandro Aldini, Mario Alvim, Anindya Banerjee, Béatrice Bérard, Arnar Birgisson, Michele Boreale, Kostas Chatzikokolakis, Tom Chothia, David Clark, Jorge Cuellar, Alessandra Di Piero, Ehab ElSalamouny, Sardaouna Hamadou, Holger Hermanns, Michael Hicks, Sebastian Hunt, Daniel Kifer, Boris Köpf, Matteo Maffei, Pasquale Malacaria, Fabio Martinelli, Michael W. Mislove, C. Carroll Morgan, John Mullins, Gethin Norman, Martin Ochoa, Catuscia Palamidessi, Quoc Sang Phan, Alexander Pretschner, Andrey Rybalchenko, Mathieu Sassolas, Vladimiro Sassone, Fabrizio Smeraldi, Geoffrey Smith, Marco Stronati, Tachio Terauchi, Carmela Troncoso, Herbert Wiklicky, Santiago Zanella Beguelin



The high amount of trust put into today's software systems calls for a rigorous analysis of their security. Unfortunately, security is often in conflict with requirements on the functionality or the performance of a system, making perfect security an impossible or overly expensive goal. Under such constraints, the relevant question is not whether a system is secure, but rather how much security it provides. Quantitative notions of security can express degrees of protection and thus enable reasoning about the trade-off between security and conflicting requirements. Corresponding quantitative security analyses bear the potential of becoming an important tool for the rigorous development of practical systems, and a formal foundation for the management of security risks.

While there has been significant progress in research on quantitative notions of security and tools for their analysis and enforcement, existing solutions are still partial. The focus of the seminar is to discuss the following key issues.

Quantitative Notions of Security: A single qualitative security property may give rise to a spectrum quantitative generalizations, each with different characteristics and application domains. For quantitative confidentiality, current research focuses on differential privacy and measures based on information-theoretic entropy. For other security properties such as integrity, availability, incoercibility, vote verifiability, etc., quantitative generalizations are only now emerging or have not even been proposed. One goal of this seminar is to advance the understanding of the relationship between existing quantitative security properties, and to join forces in the development of new ones.

Tools for Quantitative Security Analysis: Performing a quantitative security analysis of a realistic system is a challenging problem due to the complexity of modern software. It is mandatory to provide developers with tool support for this task. One goal of this seminar is to advance the understanding of the fundamental reasoning principles for quantitative notions of security, their connection to programming languages and verification

techniques, and the theoretical limits for automatically deriving quantitative security guarantees.

Novel Application Domains: Quantitative security analyses have been successfully applied, e.g., for quantifying the side-channel leakage in cryptographic algorithms, for capturing the loss of privacy in statistical data analysis, and for quantifying security in anonymity networks. In emerging application domains such as electronic voting or distributed usage control, the need for quantitative analyses has been recognized. It is a goal of this seminar to foster the collaboration between experts in emerging application domains and those in quantitative security analysis.

4.58 Analysis of Security APIs

Organizers: Mike Bond, Riccardo Focardi, Sibylle Fröschle, and Graham Steel
Seminar No. 12482

Date: 26.–28. November, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.11.155

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Mike Bond, Riccardo Focardi, Sibylle Fröschle, and Graham Steel



Participants: Pedro Adao, Ross Anderson, Daniel Angermeier, David R. Aspinall, Romain Bardou, Mike Bond, Veronique Cortier, Marion Daubignard, Stéphanie Delaune, Riccardo Focardi, Sibylle Fröschle, Steve Kremer, Robert Künnemann, Ralf Küsters, Flaminia L. Luccio, Matteo Maffei, Sebastian Mödersheim, Benjamin Morin, Andreas Philipp, Markus N. Rabe, Phillip Rogaway, Mark D. Ryan, Stefanie Schlegel, Jörg-Cornelius Schneider, Laurent Simon, Viorica Sofronie-Stokkermans, Marco Squarcina, Graham Steel, Petr Svenda, Susan Thompson, Frank Thunig, Ronald Toegl

The seminar brought together 32 participants from academia and industry in Europe and the USA. It featured a joint session with the concurrent seminar on quantitative security analysis (which included the keynote talk), a breakout session with demonstrations of software and practical classes, a discussion of the most important open problems in the field and a collection of talks spanning the breadth of the field from theoretical models to applications.

■ Research Context and Goals of the Seminar

A security API is an Application Program Interface that allows untrusted code to access sensitive resources in a secure way. It is the interface between processes running with different levels of trust. Examples of security APIs include the interface between the tamper-resistant chip on a smartcard (trusted) and the code running on the client application (untrusted), the interface between a cryptographic Hardware Security Module (or HSM, trusted) and the host machine (untrusted), and web service APIs (an interface between a server, trusted by the service providers, and the rest of the Internet).

The crucial aspect of a security API is that it is designed to enforce a policy, i.e. no matter what sequence of commands in the interface are called, and no matter what the parameters, certain security properties should continue to hold. This means that if the less trusted code turns out to be malicious (or just faulty), the carefully designed API should prevent compromise of critical data. Designing such an interface is extremely tricky and error prone, and over the last ten years, serious vulnerabilities in the security APIs deployed in HSMs in the ATM (cash machine) network and in commodity security devices like smartcards and USB keys have come to light.

A number of formal methods researchers have turned their attention to security APIs over the last five years. While significant advances have been made and notable results achieved, such as the discovery of several new attacks, the process of specifying and verifying the security policy for such APIs still lacks both satisfactory

foundations and efficient algorithms. At the same time, the security API paradigm is being proposed as a solution for more and more applications, from social networks to smartphones, with more complicated and less well understood security goals.

The objective of the seminar was to bring together researchers in academia and industry around the topic of security APIs and their analysis. There were three main aims:

1. To address the shortcomings of current API analysis techniques as applied to the relatively well explored domains of cryptographic key management and HSMs, in particular in their ability to deal with global mutable state and their models of cryptography.
2. To identify the security API requirements arising from the new generation of applications, in mobile device applications and web services, and map out the research problems that need to be solved in order that formal API analysis can be applied here.
3. To find ways to include the process and results of formal API analysis into the standards and certification procedures.

Some progress was made on all these points in the talks and the discussions late into the evening that followed in the conducive environment of Schloss Dagstuhl. On the first point, several talks presented models specifically aimed at modelling state in a more satisfactory way, and we had a tutorial on the verification methods used in program analysis. Several new application areas for API analysis were presented, including car to car communication and password protection. Some highly enlightening talks on the standards process helped to improve understanding of the problem, if not providing steps to an easy solution. The variety of open problems identified (see summary below) shows that this is a vibrant area of computer security research with much promise for the future.

4.59 Interpreting Observed Action

Organizers: Susanne Biundo-Stephan, Hans Werner Guesgen, Joachim Hertzberg, and Stephen R. Marsland

Seminar No. 12491

Date: 2.–6. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.1

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Susanne Biundo-Stephan, Hans Werner Guesgen, Joachim Hertzberg, Stephen R. Marsland

Participants: Sven Albrecht, Mehul Bhatt, Susanne Biundo-Stephan, Martin V. Butz, Amedeo Cesta, Krishna Sandeep Reddy Dubba, Tom Duckett, Frank Dylla, Simone Frintrop, Martin A. Giese, Hans Werner Guesgen, Verena V. Hafner, Joachim Hertzberg, Alexandra Kirsch, Geert-Jan M. Kruijff, Stephen R. Marsland, Bernd Neumann, Heiko Neumann, Lucas Paletta, Bernd Schattenberg, Aryana Tavanai, Sabine Timpf, Thomas Wiemann, Diedrich Wolter



For many applications of smart embedded software systems, the system should sense the footprint of a human or humans acting in the system's environment, interpret the sensor data in terms of some semantic model about what the human is doing, and respond appropriately in real time. Examples of such applications include smart homes, human-machine or human-robot interaction, assistance, surveillance, and tutoring systems; given the current trend towards ambient intelligence, ubiquitous computing, and sensor networks, the number of systems in these categories can certainly be expected to rise in the next ten years or so.

The problem shares many features with classical object recognition and scene reconstruction from sensor data in terms of a static scene model. Interpreting in semantic terms sensor data from the environment has a long tradition in AI – arguably, it has been one of the original core problems put forth by AI's founding fathers. However, the problem of interpreting observed action in the sense of this seminar differs in some aspects from what state-of-the-art AI or engineering approaches would allow to be tackled by routine:

- **Events in space-time** rather than static objects need to be characterized. This necessarily involves some representation and model of temporal and spatial data (e.g., the human put a sauceman *on* the cooker *and then* turned the cooker on).
- **Real-time processing** of the sensor data or percepts is required to keep track of what is happening. In fact, “real time” here is the pace of human action, i.e., relatively slow compared to CPU clock ticks. However, given a potentially rich stream of sensor data and a potentially large body of background knowledge, even this pace is demanding for the respective knowledge processing methods.
- **Willed human action**, be it planned, intended, or customary, is the domain of interpretation. In knowledge representation, this appears to be a relatively unexplored area, compared to, say, upper ontologies of household items, red wine, or pizza varieties.

Contemplating the three words that make up the title of this seminar (“interpret”, “observe”, and “act”), it becomes clear that there are a number of issues that need to be addressed in this context. Firstly, any interpretation is to some degree subjective and uses a particular repertoire of basic actions in its language. Secondly, an observation uses a particular type of sensor data and often is not possible without interpretation at the same time. Thirdly, there are issues around what actions are to be considered:

- Are only willed and physical action to be considered?
- Is avoidance an action?
- What constitutes an action in the first place?
- When does a particular action end?
- Is an unsuccessful action an action?

In summary, what precisely is observed action interpretation and what would be benchmark data for it?

To find an answer to this question, the participants of the seminar emerged themselves in a variety of activities: technical talks, working groups, plenary discussions, and a number of informal discussions. In the rest of this report, some of these activities and their results are discussed in more detail.

4.60 Human Activity Recognition in Smart Environments

Organizers: James L. Crowley, Kai Kunze, Paul Lukowicz, and Albrecht Schmidt
Seminar No. 12492

Date: 02.–07. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.17

© © © Creative Commons BY-NC-ND 3.0 Unported license

© James Lawrence Crowley, Paul Lukowicz, and Albrecht Schmidt



Participants: Gernot Bahle, Michael Beigl, Mehul Bhatt, Ulf Blanke, Oliver Brdiczka, Nora Broy, Andreas Bulling, Jingyuan Cheng, Joelle Coutaz, James L. Crowley, Alois Ferscha, James Fogarty, Dawud Gordon, Georg Groh, Niels Henze, Jochen Kerdels, Koichi Kise, Daniel Kohlsdorf, Kai Kunze, Alexander Lehmann, Paul Lukowicz, Diana Nowacka, Thomas Phan, Gerald Pirkl, Thomas Plötz, Alireza Sahami, Bernt Schiele, Albrecht Schmidt, Hedda R. Schmidtke, Bernhard Sick, Stephan Sigg, Hideyuki Tokuda, Kristof Van Laerhoven, Katharina A. Zweig

Today, commercial systems have become popular that utilize a broad range of sensors to facilitate gesture and motion-based interaction. Examples range from multi-touch surfaces, through tilt control common in mobile phone applications, and complex motion-based games controllers, e.g. Nintendo Wii and Microsoft Kinect. While these systems are mainstream, the next basic research challenge is activity-driven, implicit interaction. Two key differences to existing systems are:

1. the interpretation of complex human activities, and
2. the extension of interaction from periods where a user consciously performs control gestures to permanent monitoring of user activity.

Conceptually, activity-driven interaction builds on the vision of context awareness developed since the 1990 [1–3]. Applications range from sports, through mobile gaming, information retrieval, personal healthcare to industrial work support [4–6]. For example, monitoring certain activities can support therapy in areas ranging from cardiovascular diseases to psychiatric disorders and cognitive disabilities. Activity based support (automatically showing correct manual pages, pointing out mistakes) can speed up industrial maintenance tasks by up to 30%.

Despite demonstrated potential, currently only very simple activity based applications such as physical activity monitoring have managed to go beyond early stage lab demonstrations. From the scientific point of view the question is how to map information available from unreliable, often simple sources onto complex human activities. The main challenges stem from the combination of three factors:

- In every day situations sensor choice, placement and configuration is often dictated by practicability, usability, and user acceptance constraints rather by the requirements of the recognition system. In addition, the system configuration may dynamically change [7, 8].

- The diversity of human behavior. Even the simplest activities can be performed in a multitude of ways differing not only between people, but also between individual execution instances of a single person. (e.g. using different arms, different hand positions, or even the hip to close a drawer)
- The complexity of human behavior. Relevant human actions are seldom atomic and independent. Instead, a complex hierarchy of actions that may be executed in parallel, overlap and interleave is to be considered by the recognition system.

Beyond the technological challenges involved in the recognition system, there are additional unsolved problems including application design, usability, user acceptance, and business models for commercialization.

The field lacks also definitions for many commonly used terms including “action,” “sensor,” “evidence,” and even “activity” itself, leading to ambiguity in scientific discourse. The conceptual grounding provided by Nardi and Kaptelinin’s definition of Activity Theory are perfectly understandable to a human [9]. Yet, they are not easily codified into machine programmable constructs. The theory recognizes that elements of “activity” such as goal and motive are socially constructed, depending on the perspectives of the actors in the system. Despite the complexities of “activity” at the human cognitive level, researchers demonstrated that some notions of activity can be utilized in computer systems, but meanings of terms differ among the various research groups. Currently, many different communities are involved in research related to activity recognition, including the core ubicomp community, human computer interaction, computer vision, cognitive science and artificial intelligence.

Privacy concerns are a critical barrier to adoption of activity-based technologies [10]. These concerns range from risk of criminal activities (e.g., stalking and identity theft), to social issues of managing personal relationships. Technological approaches to addressing the concerns must also be based on deep understanding

of the psychological, sociological and political constraints under which people will operate activity-based systems.

The top level objective of the workshop was to define and establish the scientific community and associated research questions/methodologies related to the broad area of activity recogni-

on. The major tangible outcomes are the start of the creation of an activity recognition repository accessible under <http://activity-recognition.github.com> and the plan of writing a standard book about activity recognition.

References

- 1 A Schmidt, K Aidoo, and A Takaluoma. Advanced interaction in context. *Handheld and Mobile Computing*, 1999.
- 2 T. Starner, B. Schiele, and A. Pentland. Visual contextual awareness in wearable computing. In *Proc. of the Second Int'l Symp. on Wearable Computing, Pittsburgh, October, 1998*.
- 3 A. Schmidt. Implicit human computer interaction through context. *Personal and Ubiquitous Computing*, 4(2):191–199, 2000.
- 4 J. Hoey, T. Plötz, D. Jackson, A. Monk, C. Pham, and P. Olivier. Rapid specification and automated generation of prompting systems to assist people with dementia. *Pervasive and Mobile Computing*, 7(3):299–318, 2011.
- 5 J Lester, T Choudhury, and G Borriello. A practical approach to recognizing physical activities. *Proc. of Pervasive*, Jan 2006.
- 6 K. Kunze, F. Wagner, E. Kartal, E. Morales Kluge, and P. Lukowicz. Does context matter?-a quantitative evaluation in a real world maintenance scenario. *Pervasive Computing*, pp. 372–389, 2009.
- 7 K. Kunze and P. Lukowicz. Dealing with sensor displacement in motion-based onbody activity recognition systems. In *Proc. of the 10th Int'l Conf. on Ubiquitous Computing*, pages 20–29. ACM, 2008.
- 8 K. Kunze and P. Lukowicz. Symbolic Object Localization Through Active Sampling of Acceleration and Sound Signatures. *Lecture Notes in Computer Science*, 4717:163, 2007.
- 9 V. Kaptelinin and B.A. Nardi. Activity theory in a nutshell. *Acting with Technology: Activity Theory and Interaction Design*, pp. 29–72, 2006.
- 10 S. Sigg. Context-based security: State of the art, open research topics and a case study. In *Proc. of the 5th ACM Int'l Workshop on Context-Awareness for Self-Managing Systems*, pp. 17–23. ACM, 2011.

4.61 Organizational Processes for Supporting Sustainable Security

Organizers: Lizzie Coles-Kemp, Carrie Gates, Dieter Gollmann, Sean Peisert, and Christian Probst

Seminar No. 12501

Date: 9.–12. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.37

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Lizzie Coles-Kemp, Carrie Gates, Dieter Gollmann, Sean Peisert, and Christian Probst



Participants: Julie Boxwell Ard, Florian Arnold, Debi Ashenden, Arshid Bashir, Sören Bleikertz, Rainer Böhme, Lizzie Coles-Kemp, Sophie Engle, Vaibhav Garg, Carrie Gates, Dieter Gollmann, Marit Hansen, Cormac Herley, Michael Huth, Jean-Lous Huynen, Dan Ionita, Florian Kammüller, Ana Margarida Leite de Almeida Ferreira, Makayla Miranda Lewis, Kai-Uwe Loser, Anne-Marie Oostveen, Wolter Pieters, Joachim Posegga, Marco Prandini, Christian W. Probst, Ingrid Schirmer, Sven Übelacker, Sam Weber, Sean Whalen, Trish Williams, Alf Zugenmaier

The Dagstuhl seminar “Designing for process resilience to insider threats” was held on December 10–12th December, 2012 (Seminar #12501) to advance our understanding of ways of reducing insider threats through the design of resilient organizational processes.

The 2012 seminar built on the results of its predecessor from 2010 (Insider Threats: Strategies for Prevention, Mitigation, and Response, #10341, Seminar Homepage, Seminar Report). In this seminar we developed a shared, inter-disciplinary definition of the insider and a good formulation for a taxonomy or framework that characterizes insider threats. The seminar also began to explore how organizational considerations might better be incorporated into addressing insider threats.

The purpose of the 2012 seminar was to build on the understanding of the classification of the insider threat as a type of informed threat and the design requirements for tools and policies to respond to this category of threat that we had gained from the 2008 and 2010 Dagstuhl seminars on insider threats (Countering Insider Threats, #08302, and Insider Threats: Strategies for Prevention, Mitigation, and Response, #10341). Our goal was to explore what makes organizational processes resilient to insider threats. The exploration of organizational processes required us to consider the fluid set of informed actors against organizations whose processes and boundaries can be dynamic. It also required us to conceptualise threats and vulnerabilities as “emergent”. The conclusions from the previous seminars had resulted in the insight that resilient organizational processes are more resilient with respect to insider threats and more capable of limiting the damage from insider attacks. We also had the insight that resiliency appears to stem from usable, effective, and efficient security having been built into the organizational processes.

The seminar participants contained a carefully balanced mix of social and computer scientists and practitioners in order to explore the technological, organizational and social dimensions of the orga-

nizational process and its implementation. In order to productively combine the skills of the different disciplines and perspectives presented, the seminar started with a series of provocations. Debi Ashenden presented a provocation about the competing and sometimes conflicting uses of gamification in the UK military setting. Kai-Uwe Loser presented a grounded example of personal data management practices and the conflicting perceptions of policy compliance that emerged within the example. Trish Williams presented a provocation about the value of big data in the case of electronic health data.

These design principles reflect a start point for future work on the design of organizational processes that are sustainably secure. Seminar organizers intend to produce a book that extends and explores these principles.



Fig. 4.17
Chess board of Ingo Wegener. Donated to Schloss Dagstuhl by Ms Wegener-Mürbe.

4.62 Securing Critical Infrastructures from Targeted Attacks

Organizers: Marc Dacier, Frank Kargl, and Alfonso Valdes

Seminar No. 12502

Date: 09.–12. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.49

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Frank Kargl, Mark Dacier, and Alfonso Valdes



Participants: Magnus Almgren, Nils Aschenbruck, Davide Balzarotti, Rafael Barbosa, Gunnar Bjoerkman, Damiano Bolzoni, Levente Buttyan, Alvaro Cárdenas Mora, Marco Caselli, Marc Dacier, Sandro Etalle, Felix C. Freiling, Jakob Fritz, Elmar Gerhards-Padilla, Dina Hadziosmanovic, Frank Kargl, Stefan Katzenbeisser, Erwin Kooi, Maryna Krotofil, Klaus Kursawe, Corrado Leita, Tobias Limmer, Michael Munzert, Heiko Patzlaff, Andreas Paul, Franka Schuster, Valentin Tudor, Alfonso Valdes, Stephen Wolthusen, Emmanuele Zambon

The last years have highlighted the fact that our ICT security precautions in many critical infrastructure (CI) systems are clearly insufficient, especially if considering targeted attacks carried out by resourceful and motivated individuals or organizations. Critical infrastructures, like energy or water provisioning, transportation, telecommunication, or health support are relying to an ever-larger extent on ICT, often being monitored or controlled in a semi or fully automated way. Disruption of these control processes could turn out to be disastrous, especially as many of these systems are cyber-physical systems that interact with the real world through sensors and actuators and can thus have a direct influence on the physical world not mediated by the common sense of a human being.

Rendering ICT systems in such critical infrastructure unusable or malfunctioning can cause huge economical damages or even endanger human lives. Some examples: it is reported by the Institute for Science and International Security (ISIS) in December 2010⁷ that the Stuxnet malware actually damaged around 1000 Uranium enrichment centrifuges in the Iranian enrichment facility in Natanz (which was possibly its goal). If the same would happen in a European Uranium enrichment facility, the economical damage would be significant and danger to population due to failure of systems could not be ruled out completely. In 2000, an insider attack on a sewage treatment facility in Queensland, Australia caused millions of liters of raw sewage to spill out into local parks and rivers⁸. The CIP Vigilance Blog collects a long list of such issues⁹.

There are many similar examples where Industrial Control Systems (ICS) have been affected due to insufficient security precautions. Moreover, the apparent success in infiltrating Critical Infra-

structure environments is calling attention on the ineffectiveness of standard security mechanisms in detecting similar attacks. Stuxnet is believed to have been operating undetected for almost one year leveraging multiple vulnerabilities that were previously unknown, and has been discovered only as a consequence to an operational anomaly that triggered the attention of the field operators. This fact clearly shows that not only our security mechanisms in ICS are insufficient, but that even our methods to find vulnerabilities and detect ongoing or successful attacks in Critical Infrastructure environments are not up to their task. It is very likely that Stuxnet could be the “first of a kind”, as demonstrated by the recent apparition of the so-called Duqu threat, apparently based on the same code (see the Symantec thorough analysis for more on this topic¹⁰).

Similar argumentation can be applied to other forms of control systems like Intelligent Transport Systems, modern health systems, Smart electric grids, and many more. The advanced metering infrastructures (AMI) now being deployed in some form on the electric grids of many countries offers potential benefits in terms of reduction of peak load, which in turn enables green house gas reduction and various economic benefits. However, it introduces potentially hundreds of millions of computationally limited networked endpoints outside of a defensible physical or electronic perimeter. Moreover, smart grids may be subject to attacks that do not require an adversary to compromise a device, whether a smart meter on a residence or a phasor measurement unit (PMU) that contributes to wide area measurement or state estimation. Real-time price signals communicated to smart meters may induce volatility, and if spoofed may lead to destabilizing load fluctuation (see [1]). Spoofing of GPS

⁷ <http://www.isis-online.org/isis-reports/detail/did-stuxnet-take-out-1000-centrifuges-at-the-natanz-enrichment-plant/>

⁸ http://www.theregister.co.uk/2001/10/31/hacker_jailed_for_revenge_sewage/

⁹ <http://ciip.wordpress.com/2009/06/21/a-list-of-reported-scada-incidents/>

¹⁰ http://www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_duqu_the_precursor_to_the_next_stuxnet.pdf

signals can cause PMUs to lose synchronization, resulting in threats to real-time control and corrupt grid state estimation.

There are many challenges involved in this, especially the heterogeneity of the systems that often involve legacy and proprietary system where not even all specification might be available to security engineers. High dependability and availability requirements of such systems often do not allow fast update cycles in case of security vulnerabilities are disclosed. The trend to use more COTS hardware and software in such systems creates problems and opportunities at the same time. A problem is that all malware that is available in such systems suddenly also becomes available to attackers on Critical Infrastructure ICT and that a lot of known vulnerabilities become exploitable. On the pro side, many established security mechanisms like firewalls, Intrusion Detection Systems, or OS security mechanisms like malware scanners can be applied. However, you often need to specifically adjust them for the new domain (e.g., by having SCADA specific signatures for an IDS). At the same time, the different (dependability) requirements and different applications in Critical Infrastructure Systems often require new or updated approaches, e.g., regarding security updating or security testing methodologies.

The research community has taken up this challenge, as can be seen by the emergence of specific research projects (e.g., EU projects like ReSIST, IRIIS, VIKING, SERSCIS, INSPIRE, CRUTIAL, CRISALIS), and regular contributions on the topic at conferences and workshop (RAID, DIMVA, CCS, LEET, IEEE SSP, NDSS, Usenix Security, etc.). The US Department of Homeland Security and Department of Energy fund numerous projects under programs such as the National SCADA Test Bed (NSTB) and Cyber Security for Energy Delivery Systems (CSEDS). However, we identified that the research community would benefit from being better connected, having identified a clear list of major research challenges, and knowing to what extent they have been addressed so far. Stemming from this motivation, we proposed this Dagstuhl research seminar with the goal to bring together leading researchers both from academia and industry to discuss and evaluate the state of the art 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.

One of the most important aspects was to identify whether security challenges and solutions apply to all different areas of CI, be it water, electricity, gas, transport, health-support, public safety infrastructures, or tele-communication. Our initial expectation was that there would be clusters of domains with very similar profiles on the one hand, but also large differences between clusters. This, however, was not clear previously, as many security researchers focused on specific areas or specific aspects of security.

Beyond, during the seminar we also focused on the question how targeted attacks on CI differ from ubiquitous unspecific attacks by malware or occasional hackers. As the later do not focus specifically on CI, they will typically not create large-scale damages — if damages occur, this is typically the consequence of computer systems being down. In contrast, the Stuxnet example illustrates how targeted malware can be injected into target systems in a very stealthy way and can cause subtle damage that can go unnoticed for a long time. Consequently, security countermeasures, reactions, and forensic methods have to differ as well. However, the research community has just started to address the area of targeted attacks.

The seminar started from a set of questions related to this:

- What are the specific security challenges and requirements that are ubiquitous throughout different Critical Infrastructure domains? Where do those domains differ in terms of security?
- What is the status with respect to protection from, detection, and analysis of targeted attacks on Critical Infrastructures? What solutions can be transferred from general ICT? Where have

new solutions already been found? Where is further research needed?

- Do these solutions apply to Critical Infrastructures in general, or do we need to work on domain-specific solutions?
- How can the negative effects of successful attacks be contained?
- How can CI be made resilient to attack, and able to maintain critical (possibly degraded) function in the presence of attack?
- How can we bridge the gap between low-level research on the granularity of individual ICT devices or single networks, e.g., to conduct forensic analysis or deploy IDS, on the one hand and on the other the research that assesses the system-wide effects of targeted attacks, e.g., on effect propagation?
- How can technical solutions and organizational policies be aligned and enhanced in a consistent way?
- How can we bridge the gap of knowledge between security experts rarely aware of the specific characteristics of CI systems and CI experts not necessarily up to date with the latest security research outcomes.
- How can we shed some light on CI insecurity without running the risks of opening a Pandora box? What are the consequences of such risks? Are there legal implications to consider?
- How do the approaches of academia and industry in addressing targeted attacks on CI differ?

Many of these questions were addressed during our two and a half day Dagstuhl seminar from December 9 to 12, 2012. We had the envisioned nice mix of participants with an industry participation of over 30 % and experts from various domains of critical infrastructures.

The agenda featured two main plenary talks, nine short presentations, and regular working group breakout sessions. The plenary talks were given by Alvaro Cárdenas Mora from Fujitsu Laboratories of America / UT Dallas who spoke on “Short-term and Long-term research Challenges for Securing Cyber-Physical Systems” and Levente Buttyán from the BME CrySyS Lab who gave us a first hand insight into analysing targeted attack malware in his presentation on “The cousins of Stuxnet: Duqu, Flame, and Gauss”.

The short presentations focussed on a broad variety of topics, some giving broader updates on research agendas and activities like the European CRISALIS project, some others addressing specific areas like train control systems or smart grids. One short talk by Felix Freiling asked the challenging question whether detecting targeted attacks has to be considered impossible by their very nature, a discussion that working group 2 later continued in depth. Other topics addressed in the short presentations included intrusion detection mechanisms for industrial control systems, a report on the CERT run by Siemens, and on societal consequences of cyber attacks on electrical supply systems.

These talks provided perfect input for our working groups. We initially envisioned four working groups with the titles (1) Business Aspects of Security for CI in Different Domains, (2) Attacker Models and Risk Analysis for Targeted Attacks on CI, (3) CI Security in different CI Domains, and (4) Technical Security Approaches for Intrusion Detection and Network Monitoring. However, during initial discussions and working group assignments, groups (1) and (3) found their topics to be closely related and decided to merge.

The merged working groups (1) and (3) first identified the challenge that definitions of what a *Critical Infrastructure* actually is are quite diverse and fuzzy which led to a narrowed down definition of *Critical Cyber Infrastructure* which provided a working definition to then reason about the nature of security incidents and solutions in such systems. One conclusion from their work was that there is a gap between generic IT security and the large number of different CI domains that may be bridged by providing clearer reference sce-

narios that security researchers can focus on. That way, one could then identify whether generic security solutions can be applied to such scenarios or even cross-scenario or whether specific solutions need to be found.

Working group (2) mostly investigated attacker models for targeted attacks. Again, the term targeted attack was not clear and the first result of the working group was a attack classification scheme to be able to narrow down on this term and distinguish various types of attacks. The group even went beyond the targeted attack term and suggested PEST (persistent, sophisticated and targeted) as a categorization of the most critical types of attacks. In a second meeting, the working group discussed attacker motivations and identified a clear lack of intelligence regarding such motivations. Therefore, a lot today is more guesswork than based on clear facts and more investigations into the nature of PEST attacks seems to be required.

Finally, working group (4) focussed on the technical topic of intrusion detection and network monitoring in ICS, coming up with

a list of attack scenarios, technical challenges and ideas for enhancement of countermeasures.

In a final plenary wrap-up discussion, all participants agreed that the seminar's topic was definitely a very challenging one. As both the definition of Critical Infrastructure and Targeted Attack are not even clearly agreed upon and as CIs are so diverse, we were not even able to cover all possible instantiations of CIs by dedicated experts. Especially the work in the working groups provided important first steps towards clearer definition and a common understanding of these issues and as such the seminar has to be considered a success that should be followed up by future activities.

The question whether joint security approaches and solutions for targeted attacks on critical infrastructures can be found can therefore not finally be answered. However, the research community and industry would definitely benefit from a closer cooperation of researchers and practitioners that work on *PEST attacks on Critical Cyber Infrastructures*.

■ References

1 Roozbehani, M.; Dahleh, M.A.; Mitter, S.K., "Volatility of Power Grids Under Real-Time Pricing", IEEE Transacti-

ons on Power Systems, vol. 27, no. 4, pp. 1926–1940, Nov. 2012



Fig. 4.18

Quilt of Jill Knuth. Part of the Dagstuhl art collection and donated by Jill and Don Knuth. The picture was taken when Don Knuth attended Dagstuhl Seminar 12471 in November 2012.

4.63 Divide and Conquer: the Quest for Compositional Design and Analysis

Organizers: Marieke Huisman, Barbara Jobstmann, Ina Schaefer, and Marielle Stoelinga
Seminar No. 12511

Date: 16.–21. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.64

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Marieke Huisman, Barbara Jobstmann, Ina Schaefer, and Marielle Stoelinga



Participants: Wolfgang Ahrendt, Simon Bliudze, Einar Broch Johnsen, Ferruccio Damiani, Rayna Dimitrova, Christian Eisentraut, Bernd Finkbeiner, Kathi Fisler, Susanne Graf, Dilian Gurov, Reiner Hähnle, Ludovic Henrio, Marieke Huisman, Bart Jacobs, Barbara Jobstmann, Christian Kästner, Shmuel Katz, Shriram Krishnamurthi, Malte Lochau, Oscar M. Nierstrasz, Doron A. Peled, Arnd Poetzsch-Heffter, Erik Poll, Sophie Quinton, Johannes Reich, Arend Rensink, Ina Schaefer, Malte Schwerhoff, Vasiliki Sfyrla, Marjan Sirjani, Lei Song, Martin Steffen, Marielle Stoelinga, Ufuk Topcu, Stavros Tripakis, Andrzej Wasowski

Compositionality is a key concept in computer science: only by breaking down a large system into smaller pieces, we can build today's complex software and hardware systems. The same holds true for verification and analysis: realistic systems can only be analyzed by chopping them up into smaller parts. Thus, compositionality has been widely studied in various different settings, and by different communities: people in programming languages, software verification, and model checking have all come up with their own techniques and solutions.

Thus, the goal of this workshop has been to bring together these fields and communities, so that they can learn from and cross-fertilize each other. We have succeeded in doing so: through three extensive tutorials, longer and shorter presentations, and working sessions, researchers from different areas have learned about each others problems, techniques, and approaches.

The scientific programme was built around four corners stones

1. Personal introductions.
2. Three well-received tutorials:
 - Compositional programming by Oscar Nierstrasz
 - Compositional verification by Arnd Poetzsch-Heffter
 - Compositional modelling by Arend Rensink
3. Regular presentations, presenting in-depth technical knowledge on:
 - Verification of programming languages
 - Automatic synthesis
 - Interface theories
 - Model checking
 - Contract-based design
 - Software product lines
4. Working group sessions:
 - Working group on software product lines
 - Working group on Benchmark for Industrial Verification/Synthesis Problems

- Working group on Modular Full Functional Specification and Verification of C and Java programs that Perform I/O
- Model checking vs deductive verification
- Compositional Synthesis of Reactive Systems

4.64 Representation, Analysis and Visualization of Moving Objects

Organizers: Joachim Gudmundsson, Patrick Laube, and Emiel Van Loon

Seminar No. 12512

Date: 16.–21. December, 2012 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.2.12.89

© © © Creative Commons BY-NC-ND 3.0 Unported license

© Joachim Gudmundsson, Patrick Laube, and Emiel Van Loon

Participants: Susanne Bleisch, Kevin Buchin, Maike Buchin, Urska Demsar, Matt Duckham, Győző Gidófalvi, Joachim Gudmundsson, Bin Jiang, Andrea Kölzsch, Patrick Laube, Ben Loke, Mirco Nanni, Ran Nathan, Ross Purves, Chiara Renso, Jörg-Rüdiger Sack, Kamran Safi, Roeland Scheepens, Rodrigo I. Silveira, Aidan Slingsby, Bettina Speckmann, Egemen Tanin, Marc van Kreveld, Emiel Van Loon, Robert Weibel, Daniel Weiskopf, Georgina Wilcox, Jo Wood



The seminar brought together researchers and domain experts involved in developing and utilizing methods for knowledge extraction from massive amounts of data from moving objects. This knowledge is essential to substantiate decision making in public and private sectors, in application domains such as transportation modeling, urban planning, tourism, wildlife ecology, spatial epidemiology, location-based services, flight safety, and marine safety. Moving object data typically include trajectories of discrete spatial objects (e.g. humans, vehicles, animals, and goods), continuous phenomena (e.g. storms, ocean currents) as well as trajectories of abstract concepts (e.g. information flow, moving data points in attribute space) or even vectors of spreading diseases. Technologies for object tracking are low cost and increasingly reliable in terms of coverage and accuracy, hence movement records are nowadays generated in huge volumes on a routine basis, using diverse technologies such as radio frequency mapping, Global Navigation Satellite Systems, video sequences and Doppler radar.

The computational analysis of movement data has seen a successful first decade with progress made in capturing, preprocessing, storing, indexing and querying movement data, combined with promising results in visualizing movement and detecting movement patterns. However, whereas such basic progress in handling movement data was needed for establishing a new field and attracting funding, attention must now move on towards the extraction of useful information and process knowledge from tracking data.

In many application fields the need for analysing large sets of trajectories is evident and crucial; however, only very rudimentary automated analysis tools are available and anything more advanced is currently analyzed manually. As an example there are several companies tracking the movement of football players through video with a frequency of at least 25Hz and accuracy of approximately 10cm. But most of the analysis and the annotation (passes, throw-ins, goals etc) of the data are still made manually. Thus the analysis part has been neglected and in comparison with the image processing part it is technologically far behind in the development. A reason for this is the obvious lack of theoretical and practical so-

lutions for many crucial fundamental problems.

For that reason, this seminar focussed on formalizing methods for algorithmic analysis, visual analytics, data mining and knowledge discovery, defined by a multidisciplinary team of researchers and practitioners.

Participants and Format. The seminar brought together researchers from several disciplines involved in developing and utilizing computational techniques for spatiotemporal object representation, data mining, and visualization. This community encompasses an interdisciplinary mix of methodologically oriented as well as application oriented researchers. The methods-oriented researchers are from fields such as theoretical computer science, spatial databases, knowledge discovery and data mining, visual analytics, and geographic information science. They were complemented by application scientists from various fields, especially behavioral ecology and urban planning.

Drawing upon positive experiences in previous seminars of this series, oral presentation sessions were complemented by special sessions dedicated to open research questions and project ideas, as well as to discussions in small, concurrent break-out groups focussing on a specific domain or case studies. Since the participants of the seminar came from quite different backgrounds, concise survey talks on the first two days were delivered on movement ecology (Ran Nathan), industry movement research (Ben Loke) and spatial data mining (Győző Gidófalvi).

A data challenge was organized for the participants prior to the seminar. This ensured that the participants were well aware of the application domain which was the focus of the seminar and it gave the domain experts a possibility to see the potential use of various different approaches. The data challenge included a bird migration data set provided by Emiel van Loon.

Many interesting research results were presented, demonstrating the progress in this field. The participants were highly satisfied with the quality of the seminar; especially the involvement of domain specialists from relevant application domains was highly appreciated.

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 Dagstuhl abrufbar unter <http://www.dagstuhl.de/de/ueber-dagstuhl/presse/>.

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 2012 fand er statt vom 3. bis 6. Juni parallel zu dem Dagstuhl-Seminar “Future Internet for eHealth”. 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. Weitere Informationen sind auf der Webseite unter <http://www.dagstuhl.de/12232> abrufbar.

In 2012 wurden mehrere Seminar-bezogene Pressemeldungen herausgegeben, siehe Fig. 5.1. Beiträge in den Medien über Seminare und sonstige Aktivitäten von Schloss Dagstuhl sind auf der Webseite <http://www.dagstuhl.de/de/ueber-dagstuhl/presse/presse-ueber-dagstuhl/> nach Jahren geordnet recherchierbar.

Die Pressemeldungen dienen oftmals als Initiator für Berichterstattungen durch die Medien. So wurden auch in 2012 zahlreiche journalistische Beiträge produziert (Fachpresse, Radio, TV), die einige komplexe Seminar-Themen allgemein verständlich aufbereitet haben.

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 website at <http://www.dagstuhl.de/en/about-dagstuhl/press/>.

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 2012, the workshop took place from June 3–6 in parallel with Dagstuhl Seminar 12231, “Future Internet for eHealth”. 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” at Saarland University). Participants as well as trainers and referees were very satisfied with the workshop. See the event webpage at <http://www.dagstuhl.de/12232> for further details.

In 2012, 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 at the website at <http://www.dagstuhl.de/de/ueber-dagstuhl/presse/presse-ueber-dagstuhl/>.

The press releases are often picked up by the media. In 2012 journalists prepared a number of media reports (press, radio, TV) on complex topics from the seminars, making them intelligible to the general public.

Wie Netzwerke unsere Welt verändern

Social, Supply-Chain, Administrative, Business, Commerce, Political Networks: a multi-discipline perspective | Dagstuhl Perspectives Workshop 12182
<http://www.dagstuhl.de/12182>

Mit Smartphone und “Internet der Zukunft” Wohlstandskrankheiten vorbeugen

Future Internet for eHealth | Dagstuhl Seminar 12231
<http://www.dagstuhl.de/12231>

Wie Computer lernen, sich selbstständig gegen Hackerangriffe zu verteidigen

Machine Learning Methods for Computer Security | Dagstuhl Perspectives Workshop 12371
<http://www.dagstuhl.de/12371>

Wie Sensoren und Computer das Verhalten von Menschen erkennen und Umgebungen steuern

Human Activity Recognition in Smart Environment | Dagstuhl Seminar 12492
<http://www.dagstuhl.de/12492>

Fig. 5.1

Seminar-related press releases in 2012

Schloss Dagstuhl verbreitet Neuigkeiten rund um sein Programm über soziale Netzwerkdienste wie Twitter (@Dagstuhl) und LinkedIn. Über Twitter werden primär Programmankündigungen an aktuell knapp 400 Follower verbreitet, zunehmend nutzen es aber auch Seminar-Teilnehmer, um ihre Eindrücke vom Seminar mitzuteilen. Bei LinkedIn wird eine eigene Gruppe "Friends of Schloss Dagstuhl" unterhalten (mit derzeit etwa 550 Mitgliedern), mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Weiterhin werden dort interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben.

News on the program of Schloss Dagstuhl are disseminated via social networks such as Twitter (@Dagstuhl) and LinkedIn. Twitter is used primarily to disseminate program announcements to nearly 400 followers, but is increasingly used also by Dagstuhl Seminar participants to share their impressions. A "Friends of Schloss Dagstuhl" group is maintained at LinkedIn (with about 550 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.

Fortbildung

5.2

Educational Training

Schloss Dagstuhl engagiert sich im schulischen Bereich durch Organisation einer jährlichen Lehrerfortbildung, die sich an Informatik-Lehrer 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 22. "Lehrerfortbildung in Informatik", die vom 12.–14. Dezember 2012 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 aber Anfragen zur Teilnahme von Lehrern aus anderen Bundesländern. In 2012 konnte Schloss Dagstuhl zusammen mit den beiden pädagogischen Instituten fünf Lehrern von außerhalb des Saarlands und Rheinlands-Pfalz die Teilnahme an der Fortbildung ermöglichen. Mehr Informationen zur Veranstaltung 2012 gibt es unter <http://www.dagstuhl.de/12503>.

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 22th annual Dagstuhl Teacher Training Workshop, held at Schloss Dagstuhl on December 12–14, 2012, 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. In 2012, the center was able to invite five of these extra-regional teachers to attend the workshop thanks to support from both educational institutes. Details about the workshop in 2012 are available at <http://www.dagstuhl.de/12503>.

6 **Dagstuhl Publishing** *Dagstuhl Publishing*

Portfolio

6.1

Portfolio

Die Open-Access-Verlagsdienstleistungen von Schloss Dagstuhl werden mittlerweile systematisch betrieben, auch wenn die Prozesse sich weiterhin in der Weiterentwicklung befinden. Im Portfolio des Angebots gibt es zum einen Publikationsserien, die sich auf Veranstaltungen beziehen, die auf Schloss Dagstuhl abgehalten wurden (*Dagstuhl Reports*, *Dagstuhl Manifestos*, und *Dagstuhl Follow-Ups*), zum anderen Serien, die Konferenzen und Workshops “außerhalb” von Schloss Dagstuhl bedienen. Zudem wurde kürzlich die wissenschaftliche Zeitschrift *LITES* gegründet.

■ Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports* dokumentiert (<http://drops.dagstuhl.de/dagrep>). Die Zeitschrift wurde Anfang 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Seminaren und Workshops, die im jeweiligen Monat stattgefunden haben. Die veröffentlichten Artikel dienen dazu, eine zitierbare Dokumentation zur Verfügung zu stellen, wobei der Inhalt der Berichte nicht begutachtet wird. Um jedoch umfassende Zusammenstellungen von begutachteten Artikeln auf Basis eines Dagstuhl-Seminars oder -Perspektiven-Workshops zu ermöglichen, wurde die Buchreihe *Dagstuhl Follow-Ups* gegründet.

Von den 64 Seminaren und Workshops in 2012 haben 62 einen Bericht in den Dagstuhl Reports veröffentlicht. Für zwei Seminare steht leider kein Bericht zur Verfügung.

Das wissenschaftliche Direktorium (siehe Fig. 12.5) bildet das Editorial Board.

■ Dagstuhl Manifestos

In 2011 wurde ebenfalls die Zeitschrift *Dagstuhl Manifestos* (<http://drops.dagstuhl.de/dagman>) als Open-Access-Zeitschrift eingerichtet, um die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Workshops gehört – zu veröffentlichen. Das wissenschaftliche Direktorium (siehe Fig. 12.5) fungiert dabei als Editorial Board. Die Ausgabe für 2012 enthält zwei Manifestos, siehe Fig. 6.1.

The Open Access publishing services offered by Dagstuhl, while fully functional, are still in development. The portfolio covers series related to events at Schloss Dagstuhl (*Dagstuhl Reports*, *Dagstuhl Manifestos*, and *Dagstuhl Follow-Ups*), series for conferences and workshops held outside of Schloss Dagstuhl (*OASICs* and *LIPICs*), and the recently established journal *LITES*.

■ Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports* (<http://drops.dagstuhl.de/dagrep>). 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. Reports in this periodical serve mainly as citable documentation; the content is not peer-reviewed. 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*.

Of the 64 seminars and workshops held in 2012, 62 have published a report in the Dagstuhl Reports series. Two seminars have not prepared a report.

The Scientific Directorate (see Fig. 12.5) acts as editorial board.

■ Dagstuhl Manifestos

In 2011 we set up the *Dagstuhl Manifestos* series (<http://drops.dagstuhl.de/dagman>) as an OpenAccess journal for publishing the manifestos that are an expected result of Dagstuhl Perspectives Workshops. The Scientific Directorate (see Fig. 12.5) acts as the editorial board of the journal. The 2012 volume includes two Dagstuhl Manifestos, see Fig. 6.1.

Social, Supply-Chain, Administrative, Business, Commerce, Political networks: a multi-discipline perspective

<http://dx.doi.org/10.4230/DagMan.2.1.1>

based on Dagstuhl Perspectives Workshop 12182

Computation and Palaeography: Potentials and Limits

<http://dx.doi.org/10.4230/DagMan.2.1.14>

based on Dagstuhl Perspectives Workshop 12382

Fig. 6.1

Manifestos published in 2012 in the journal *Dagstuhl Manifestos*

■ Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups* (<http://drops.dagstuhl.de/dfu>) ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert.

In 2012 wurde ein Buch mit dem Titel *Multimodal Music Processing* veröffentlicht (online verfügbar unter <http://www.dagstuhl.de/dagpub/978-3-939897-37-8>); das Buch wurde von Teilnehmern des Dagstuhl Seminars 11041 geschrieben. Mit Hilfe einer finanziellen Unterstützung des Exzellenzclusters “Multimodal Computing and Interaction” (MMCI), der an der Universität des Saarlandes beheimatet ist (siehe <http://www.mmci.uni-saarland.de>), wurde zudem eine Druckversion erstellt (siehe Fig. 6.2).

■ Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups* book series (<http://drops.dagstuhl.de/dfu>) is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Perspectives Workshop.

In 2012 one volume entitled *Multimodal Music Processing* was published (online available at <http://www.dagstuhl.de/dagpub/978-3-939897-37-8>); the book was written by participants of the Dagstuhl Seminar 11041. With financial support from the Cluster of Excellence “Multimodal Computing and Interaction” (MMCI) located at Saarland University (see <http://www.mmci.uni-saarland.de>), a print version of the book was prepared (see Fig. 6.2).

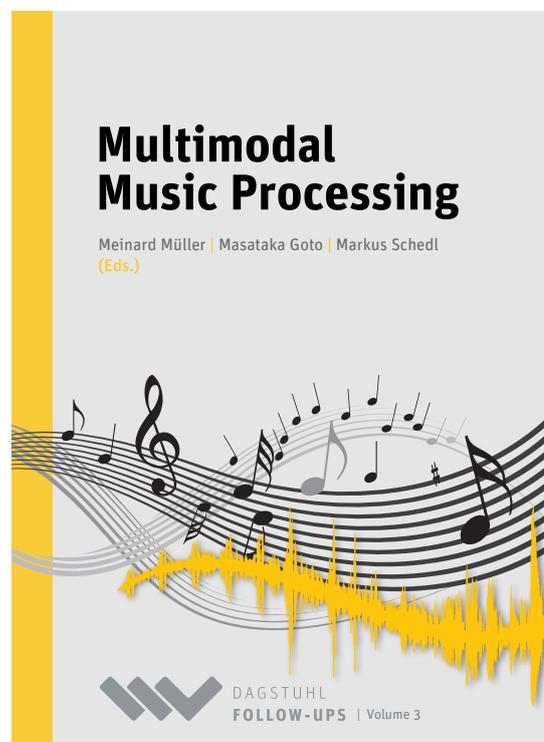


Fig. 6.2
Cover of the printed book of “Multimodal Music Processing” (Dagstuhl Follow-Ups, Vol. 3), edited by Meinard Müller, Masataka Goto, and Markus Schedl

■ OASlcs: OpenAccess Series in Informatics

Die Reihe *OASlcs* (<http://drops.dagstuhl.de/oasics>) veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Prof. Dr. Dorothea Wagner ist seit 2010 Hauptherausgeberin der Reihe. In 2011 haben Prof. Dr. Daniel Cremers, Prof. Dr. Barbara Hammer, und Prof. Dr. Marc Langheinrich die Einladung zur Mitarbeit im Herausgebergremium angenommen; siehe Fig. 6.3.

Es wurden 8 Bände in 2012 von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 6.4.

Für 2013 wurden zudem bereits attraktive Anträge eingereicht und durch das Editorial Board begutachtet.

■ OASlcs: OpenAccess Series in Informatics

The *OASlcs* series (<http://drops.dagstuhl.de/oasics>) aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. In 2010, Prof. Dr. Dorothea Wagner was named editor-in-chief. In 2011, Prof. Dr. Daniel Cremers, Prof. Dr. Barbara Hammer, and Prof. Dr. Marc Langheinrich accepted the invitation to work as editorial board members; see Fig. 6.3.

In 2012, Dagstuhl published eight OASlcs volumes covering the proceedings of topically widespread workshops and conferences, see Fig. 6.4.

OASlcs also received attractive proposals for 2013 which were reviewed by the Editorial Board.

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 (as of December 2012)

Vol. 21 SLATE'12 1st Symposium on Languages, Applications and Technologies http://www.dagstuhl.de/dagpub/978-3-939897-40-8
Vol. 22 SCOR'12 3rd Student Conference on Operational Research http://www.dagstuhl.de/dagpub/978-3-939897-39-2
Vol. 23 WCET'12 12th International Workshop on Worst-Case Execution Time Analysis http://www.dagstuhl.de/dagpub/978-3-939897-41-5
Vol. 24 SSV'12 6th International Workshop on Systems Software Verification http://www.dagstuhl.de/dagpub/978-3-939897-36-1
Vol. 25 ATMOS'12 12th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems http://www.dagstuhl.de/dagpub/978-3-939897-45-3
Vol. 26 GCB'12 German Conference on Bioinformatics 2012 http://www.dagstuhl.de/dagpub/978-3-939897-44-6
Vol. 27 IRTG1131'11 Visualization of Large and Unstructured Data Sets: Applications in Geospatial Planning, Modeling and Engineering – IRTG 1131 Workshop 2011 http://www.dagstuhl.de/dagpub/978-3-939897-46-0
Vol. 28 ICCSW'12 2012 Imperial College Computing Student Workshop http://www.dagstuhl.de/dagpub/978-3-939897-48-4

Fig. 6.4

OASlcs volumes published in 2012

■ LIPIcs: Leibniz International Proceedings in Informatics

Die LIPIcs-Reihe (<http://drops.dagstuhl.de/lipics>) veröffentlicht Tagungsbände von international renommierten 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. Wie schon in 2011, wurden auch in 2012 die Tagungsbände von 5 großen Konferenzen veröffentlicht: CSL, FSTTCS, ICLP, RTA und STACS. Desweiteren wurde der Antrag der TYPES-Konferenz vom Herausbergremium begutachtet und – beginnend mit dem Post-Proceedings-Band von 2011 – akzeptiert. Die STACS-Konferenz wurde in 2012 nach 5-jähriger Laufzeit erneut evaluiert und für weitere 5 Jahre (2014–2018) akzeptiert. Die ICLP-Konferenz hat nach experimentellen Testphase von 2 Jahren die LIPIcs-Reihe wieder verlassen, um in Zukunft die gesamten Konferenzbeiträge in der Zeitschrift TPLP zu veröffentlichen.

■ LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series (<http://drops.dagstuhl.de/lipics>) 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. As in 2011, in 2012 the series published the proceedings of five major conferences: CSL, FSTTCS, ICLP, RTA, and STACS. Additionally, the editorial board accepted the publication of the proceedings of the TYPES conference, beginning with post-proceedings for 2011. STACS was re-evaluated by the LIPIcs editorial board and accepted for another 5-year period (2014–2018). ICLP quit the LIPIcs series after having published the technical communications of ICLP in 2011 and 2012, mainly because it now wants to focus on publishing all papers in the TPLP journal.

Prof. Dr. Susanne Albers Humboldt University Berlin, 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
Prof. Dr. Wolfgang Thomas RWTH Aachen, Germany
Dr. V. Vinay Chennai Mathematical Institute, India LimberLink Technologies Pvt. Ltd.
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 (as of December 2012)

Vol. 14 STACS'12 29th International Symposium on Theoretical Aspects of Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-35-4
Vol. 15 RTA'12 23rd International Conference on Rewriting Techniques and Applications http://www.dagstuhl.de/dagpub/978-3-939897-38-5
Vol. 16 CSL'12 Computer Science Logic 2012 – 26th International Workshop/21st Annual Conference of the EACSL http://www.dagstuhl.de/dagpub/978-3-939897-42-2
Vol. 17 ICLP'12 Technical Communications of the 28th International Conference on Logic Programming http://www.dagstuhl.de/dagpub/978-3-939897-43-9
Vol. 18 FSTTCS'12 IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science 2012 http://www.dagstuhl.de/dagpub/978-3-939897-47-7

Fig. 6.6

LIPIcs volumes published in 2012

■ LITES: Leibniz Transactions on Embedded Systems

Die OpenAccess-Fachzeitschrift *LITES* (<http://drops.dagstuhl.de/lites>) veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift initiiert und aufgesetzt, wobei der eigentliche Start für Anfang 2013 geplant ist. Alan Burns konnte als Hauptherausgeber gewonnen werden sowie ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die sich für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 6.7). Im Gegensatz zu anderen Zeitschriften im Bereich eingebetteter Systeme, steht bei *LITES* eine moderate Veröffentlichungsgebühr (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. Zusätzlich wurde in 2012 ein Antrag bei der Deutschen Forschungsgemeinschaft (DFG) eingereicht, um weitere finanzielle Unterstützung für 3 Jahre zu erhalten.

■ LITES: Leibniz Transactions on Embedded Systems

The *LITES* journal (<http://drops.dagstuhl.de/lites>) publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was set up in 2012 with the plan to start operating in early 2013. In addition to winning Alan Burns as editor-in-chief, the journal was able to recruit broad team of experienced researchers to act as section-area editors (see Fig. 6.7). 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. In addition, a proposal requesting financial support for the journal during its first three years was prepared and submitted to the Deutsche Forschungsgemeinschaft (DFG) in 2012.

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 (as of December 2012)

Infrastruktur

6.2

Infrastructure

6

■ Indizierung

Alle Reihen 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. Der CPCI soll in Zukunft ähnliche bibliometrische Analysen wie der Impact Factor beim Zeitschriften-orientierten Journal Citation Index bieten. Die Einreichungsprozedur beim CPCI ist allerdings sehr intransparent und es gibt keinerlei Rückmeldung hinsichtlich Aufnahme in den Index. Die Reihen LIPIcs und OASICs sind zudem im Directory of Open Access Journals (DOAJ) gelistet, siehe Fig. 6.8.

■ LeibnizOpen

Die Leibniz-Gemeinschaft hat kürzlich *LeibnizOpen* (<http://www.leibnizopen.de/>) als ein Online-Repository gestartet, um OpenAccess-Veröffentlichungen von Leibniz-Instituten und deren Wissenschaftlern zu unterstützen. 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.

■ AG Open Access

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop "Effizientes Journal-Management: Potential durch Open Access!" initiiert und koordiniert. Der Workshop findet am 31. Januar und 1. Februar 2013 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt. Es werden circa 45 Teilnehmer aus den Verlagsabteilungen von ungefähr 20 Leibniz-Instituten erwartet.

■ Indexing

All series 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. In the future the CPCI is to offer similar bibliometric evaluations such as the "impact factor" of the Journal Citation Index. Yet the CPCI submission procedure is very non-transparent and provides no feedback on inclusion in the index. The LIPIcs and OASICs series are also listed in the Directory of Open Access Journals (DOAJ), see Fig. 6.8

■ LeibnizOpen

The Leibniz Association recently launched the *LeibnizOpen* repository (<http://www.leibnizopen.de/>) to promote the open-access publications 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

A workshop entitled "Effizientes Journal-Management: Potential durch Open Access!" was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop takes place at the Leibniz Association headquarters in Berlin from January 31 to February 1, 2013 and will bring together approximately 45 professionals in charge of publishing activities at about 20 Leibniz institutes.

DBLP	
Dagstuhl Reports	http://dblp.uni-trier.de/db/journals/dagstuhl-reports/index.html
Dagstuhl Manifestos	http://dblp.uni-trier.de/db/journals/dagstuhl-manifestos/index.html
Dagstuhl Follow-Ups	http://dblp.uni-trier.de/db/series/dfu/index.html
OASICs	http://dblp.uni-trier.de/db/series/oasics/index.html
LIPIcs	http://dblp.uni-trier.de/db/series/lipics/index.html
DOAJ	
OASICs	http://www.doaj.org/doi?func=openurl&genre=journal&issn=21906807
LIPIcs	http://www.doaj.org/doi?func=openurl&genre=journal&issn=18688969

Fig. 6.8
Indexing of Dagstuhl Publishing series in dblp and DOAJ

■ Technisches Backend: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS) werden alle Publikationen, die durch Schloss Dagstuhl veröffentlicht werden, verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der *Dublin Core-Initiative* (<http://dublincore.org/>) 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 (<http://elib.uni-stuttgart.de/opus/doku/about.php>). Weitere Informationen zu DROPS finden sich unter <http://www.dagstuhl.de/drops>.

■ Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB) zur (digitalen) Langzeitarchivierung eingereicht.

■ Mirroring

Um dem Verlust von Daten vorzubeugen, wurden in 2010 zwei Kooperationen zur Spiegelung (Mirroring) von Inhalten des Publikationsservers DROPS initiiert:

- 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, see <http://www.io-port.net> (Digital Library/LIPIcs). 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.

■ Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS). The general guidelines of the *Dublin Core initiative* (<http://dublincore.org/>) 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 and to be accessible by a wide readership. The technical basis for this is an adapted version of the OPUS system (<http://elib.uni-stuttgart.de/opus/doku/about.php>). For more information on DROPS, please refer to <http://www.dagstuhl.de/drops>.

■ Long-term Archiving

All publications are submitted to the German National Library (D-NB) 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 series. For more information, see <http://www.io-port.net> (Digital Library/LIPIcs). 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 offers a home for numerous software archives and publications. All the DROPS assets are now mirrored at regular intervals on the Aachen SunSite.

7 LZI+DBLP

LZI+DBLP

Projektbeschreibung

7.1

Die Informatik benötigt eine belastbare Datenbasis zum Nachweis und zur Evaluierung wissenschaftlicher Literatur. In dem von der Leibniz-Gemeinschaft geförderten Projekt *LZI+DBLP* entwickeln Schloss Dagstuhl – Leibniz-Zentrum für Informatik und die Literaturdatenbank *dblp* seit 2011 gemeinsam die dafür notwendigen technischen, inhaltlichen und organisatorischen Strukturen für einen Ausbau von *dblp* als die umfassende und hochqualitative bibliographische Datenbasis der Informatik. Das Projekt wird zudem durch die Klaus Tschira Stiftung gefördert. Dieses Kapitel dokumentiert die im Jahr 2012 erzielten Fortschritten bei der Etablierung neuer Richtlinien und Strukturen zur Sicherung der Datenqualität von *dblp* und der kontinuierlichen Weiterentwicklung der Datenbasis der Literaturdatenbank *dblp*.

Project Description

The documentation of research work is instrumental for identifying successful research. In a collaborative venture, the Leibniz Center for Informatics and the *dblp* computer science bibliography have come to the aid of the computing community by developing the technical and organisational structures for a consolidation of *dblp* as the comprehensive and reliable open data bibliographic resource in computer science. This project is financed by the Leibniz Association, as well as by a donation of the Klaus Tschira Foundation. This chapter gives an overview of the progress made by the *dblp* computer science bibliography in 2012 with respect to the establishment of quality assurance mechanisms and the ongoing curation and expansion of the data stock.

Ergebnisse

7.2

Results

■ Datenakquise und Statistik

Die Literaturdatenbank *dblp* indexiert Publikation an Hand vollständiger Inhaltsverzeichnisse von Konferenzbänden oder Journalausgaben. Mit Hilfe einer eigens entwickelten Software zur Datenextraktion werden Metadaten von den Webseiten der Herausgeber ausgelesen und zur weiteren Bearbeitung vorbereitet. Die Metadaten werden anschließend vom *dblp*-Team redaktionell bearbeitet und eventuelle Fehler und Ungenauigkeiten korrigiert. Diese Datenpflege wird zwar von Hilfssoftware unterstützt, erfolgt aber vornehmlich von Hand durch den jeweiligen Redakteur.

2012 wurde die Datenbank auf diese Weise um 324 089 Publikationseinträge erweitert und indexierte gegen Ende des Jahres inzwischen 2 165 813 Datensätze. Die Aufnahmequote konnte somit um weitere fünf Prozent gegenüber der ohnehin schon recht hohen Quote des Vorjahres gesteigert werden.

Ein Ziel des Projektes war die Verbesserung der Indexierung von Monographien. Der Fokus des Jahres 2012 lag hier bei der gezielten Aufnahme von Dissertationen. Dank Daten der Deutschen Nationalbibliothek konnten dabei 6 084 (vornehmlich deutschsprachige) Dissertationen in den *dblp*-Katalog übernommen werden, so dass nun 6 916 Dissertation in *dblp* indexiert werden.

Einen Überblick über die Entwicklung der Datenakquise kann Fig. 7.1 und 7.2 entnommen werden.

■ Herausgeber als Datenlieferanten

Obwohl *dblp* Daten primär selbst von Herausgeberwebseiten akquiriert hat dieser Ansatz einige Nachteile. Das größte Problem liegt dabei in dem Fehlen einer verlässlichen Benachrichtigung bei Erscheinen neuer Publikationen. Idealerweise werden diese Informationen sowie die Metadaten von den Herausgebern selbst direkt an *dblp* übersandt. Nachdem bereits 2011 derartige Absprachen mit Springer, IEEE

■ Data Acquisition and Statistics

The *dblp* computer science bibliography indexes conferences and journals on a per-volume basis. Using web harvesting software, bibliographic meta data of journal or proceedings volumes are extracted from the publisher's website. This meta data is diligently checked and corrected by the *dblp* team. The data-cleaning process is assisted by algorithms, but is executed almost exclusively by hand by a member of the *dblp* team.

In 2012, the *dblp* database grew by 324 089 publication records to reach a total of 2 165 813 records, thereby consolidating and increasing by a further five percent the high rate of new inclusions reached in 2011.

One goal of the project was to improve the coverage of monographs in *dblp*. The specific focus of 2012 was to include computer science dissertations. A total of 6 084 dissertations from predominantly German universities were added to *dblp* in 2012 thanks to data provided by the *Deutsche Nationalbibliothek* (German National Library). The database now indexes a total of 6 916 dissertations.

Figures 7.1 and 7.2 indicate the total new records by year and new records 2012 by type.

■ Publishers as Suppliers of Bibliographic Data

Usually *dblp* obtains new publication meta data by extracting data from the publisher's website. This approach has some drawbacks; in particular, there is no reliable notification mechanism for new publications. Ideally, publishers deliver the bibliographic meta data of their newest publications directly to *dblp*. After Springer, IEEE, and Schloss Dagstuhl started delivering data in the early stages of the project, in

und Schloss Dagstuhl getroffen wurden, konnten in diesem Jahr IOS Press als weiterer Partner gewonnen werden (beginnend im Frühjahr 2013). Ferner konnte auch mit der ACM Digital Library und der USENIX Association bereits Einigkeit erzielt werden; Details der Zusammenarbeit werden derzeit noch ausgearbeitet.

2012, dblp made arrangements with IOS Press to start delivering data in early 2013. The negotiations for a cooperation with the ACM Digital Library and the USENIX Association are currently ongoing.

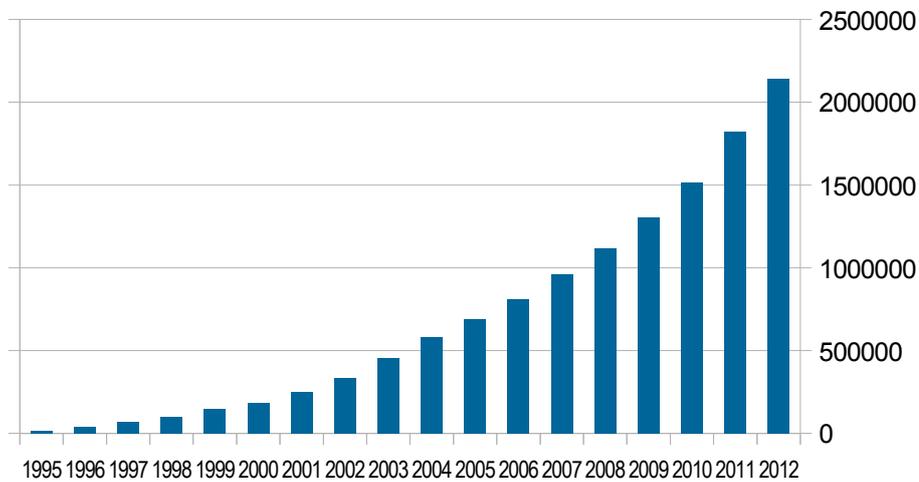


Fig. 7.1
Total number of records by year

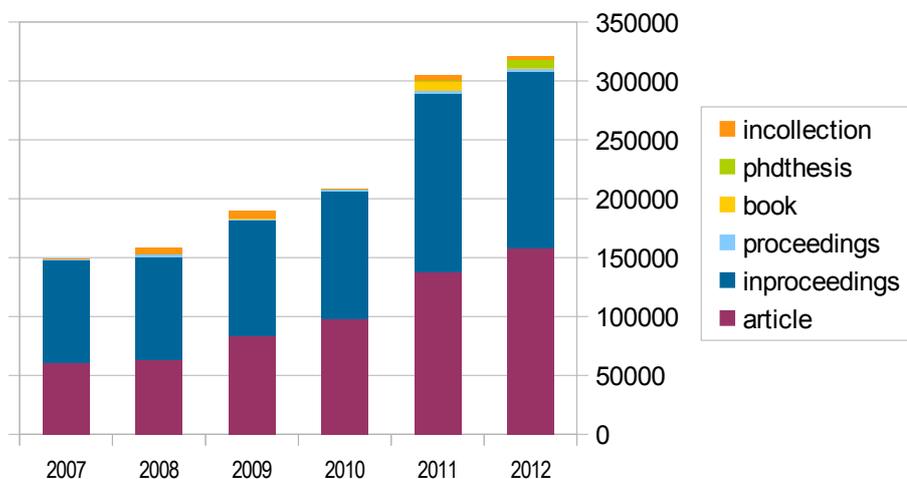


Fig. 7.2
New records by year and type

■ Verstetigung der Arbeit des dblp-Beirates

Seit seiner Konstitution im November 2011 hat der dblp-Beirat (siehe Fig. 7.3) inzwischen seine Arbeit aufgenommen. Dabei standen zunächst zwei Themen im Mittelpunkt: Die Definition von Mindeststandards für die Aufnahme von Publikationsreihen in dblp, sowie Rahmenrichtlinien für die Ausgestaltung einer bibliometrischen Infrastruktur für dblp. Beide Themen werden in Abschnitten weiter unten behandelt.

Darüber hinaus hat der dblp-Beirat das dblp-Team mit Rat und Expertise unterstützt. Dabei konnte insbesondere die Abdeckung in interdisziplinären Feldern der Informatik (wie etwa Wirtschaftsinformatik, Bioinformatik oder medizinische Informatik) erheblich verbessert werden.

Prof. Dr. Andreas Butz LMU Munich, Germany
Prof. Dr. Dietmar Saupe University of Konstanz, Germany
Prof. Dr. Hannah Bast University of Freiburg, Germany Speaker
Prof. Dr. Hans-Peter Lenhof Saarland University, Germany
Prof. Dr.-Ing. Jürgen Teich University of Erlangen-Nuremberg, Germany
Prof. Dr. Mila Majster-Cederbaum University of Mannheim, Germany
Prof. Oliver Günther Ph.D. University of Potsdam, Germany
Prof. Dr. Dr. h.c. Otto Spaniol RWTH Aachen University, Germany
Prof. Dr.-Ing. Rüdiger Dillmann Karlsruhe Institute of Technology, Germany
Prof. Dr. Rüdiger Reischuk University of Lübeck

Fig. 7.3
dblp Advisory Board (as of December 2012)

■ Mindeststandards für die Aufnahme von Publikationsreihen in dblp

Durch den dblp-Beirat wurde ein Katalog von Mindeststandards für die Aufnahme von Publikationsreihen in den dblp Datenbestand erarbeitet. Dieser Katalog enthält die folgenden Kriterien:

- Aspekte der Publikationsreihe, wie der Ausrichtung auf ein klar definiertes Themenfeld innerhalb der Informatik, eine etablierte und regelmäßige Publikationshistorie, sowie die Unterstützung durch eine etablierte wissenschaftliche Fachgesellschaft oder einen etablierten Herausgeber.
- Aspekte des Editorial Boards und der Autoren, wie etwa der wissenschaftlichen Stellung und der Internationalität des Boards und des Autorenkreises.

■ Consolidation of the dblp Advisory Board

After its constituent meeting in November 2011, the dblp Advisory Board (see Fig. 7.3) commenced work in 2012. The board focused on two topics during this first year: setting minimum standards for the inclusion of new publication venues to dblp, and preparing guidelines and an agenda for the development of a bibliometric infrastructure at dblp. Both topics are addressed in sections below.

The dblp Advisory Board also supported the dblp team by providing expertise in specific research areas, allowing dblp to significantly improve its coverage of interdisciplinary areas of computer science such as business information systems, bioinformatics, and health informatics.

■ Minimum Standards for the Inclusion of New Publication Venues

A catalog of requirements for the inclusion of new publication venues in dblp was compiled by the dblp Advisory Board. These minimum standards include the following criteria:

- venue criteria, such as the possession of a discernible thematic focus in computer science, a long-lasting regular publishing history, and the support of an established publisher or professional society,
- editor- and author-related criteria, such as the reputation and internationality of editors and authors,

- Publikationsstandards, wie der Anwendung eines ernsthaften Peer-Review-Verfahrens und von Replikations-Richtlinien, sowie die Einhaltung gängiger Konventionen bezüglich Layout, Typographie und Struktur der Beiträge.
- Aspekte der Daten-Zugänglichkeit, wie die freie Zugänglichkeit von Metadaten im Web, die Registrierung persistenter URN und das Vorhandensein eines Konzepts zur Langzeitarchivierung elektronischer Versionen.

Es ist geplant, ein Review-Verfahren für neue Publikationsreihen zu installieren, welches sich an den definierten Kriterien orientiert und die erforderlichen Mindeststandards kommunizieren soll.

■ Überarbeitung des dblp-Webdienstes

Als reines Online-Angebot ist die Funktionalität und die Bedienbarkeit seiner Webseite für dblp von zentraler Bedeutung. Daher stellte die Modernisierung und Verbesserung des dblp-Webauftrittes einen Schwerpunkt der Arbeiten in 2012 dar.

Das Layout der dblp-Webseite hatte sich im Prinzip seit seiner Erschaffung 1993 kaum gewandelt. Dieses Layout wurde nun modernisiert und an die Möglichkeiten eines zeitgemäßen CSS/HTML5-Layouts angepasst, siehe Fig. 7.4. Eine Reihe neuer Funktionen wurden hinzugefügt, so etwa eine verbesserte Suche mit Filteroptionen, ein ergonomisches Farb- und Icon-Schema, sowie neue Datenexportformate (inklusive einer Anbindung an das semantische Web). Zudem wurde beim Entwurf auf eine modulare Struktur geachtet, die zukünftig eine einfachere Integration neuer Inhalts- und Funktionskomponenten gestattet. Die neue Webseite befand sich 2012 in einer öffentlichen Testphase und soll 2013 schrittweise das alte Layout ablösen.

- publication standards, such as the existence of a serious peer reviewing process, republication policies, as well as adequate layout and typography,
- accessibility criteria, such as the long-term availability of articles, the availability of high-quality bibliographic meta data, and reachability by persistent URN identifiers.

To ensure that every new publication venue is assessed in accordance with these criteria, and to communicate our quality-control standards to the interested public, it is intended to install a venue application and review process for dblp.

■ 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 2012, the redesign and improvement of the dblp web infrastructure was a main focus of the dblp team.

Prior to 2012 the layout of the dblp web site had remained essentially unchanged since its inception in 1993. This layout was modernized and improved during the year under review to meet the capabilities of a state-of-the-art CSS based HTML5 layout, see Fig. 7.4. New functionalities were added to the new site, such as improved search and filter capabilities, an ergonomic color and icon scheme, new data export options (including linking dblp to the semantic web), and a modular structuring of the site's design components that allows for an easier integration of future content options and functionalities. The new website was available in 2012 in a public testing state and will gradually replace the old pages in 2013.

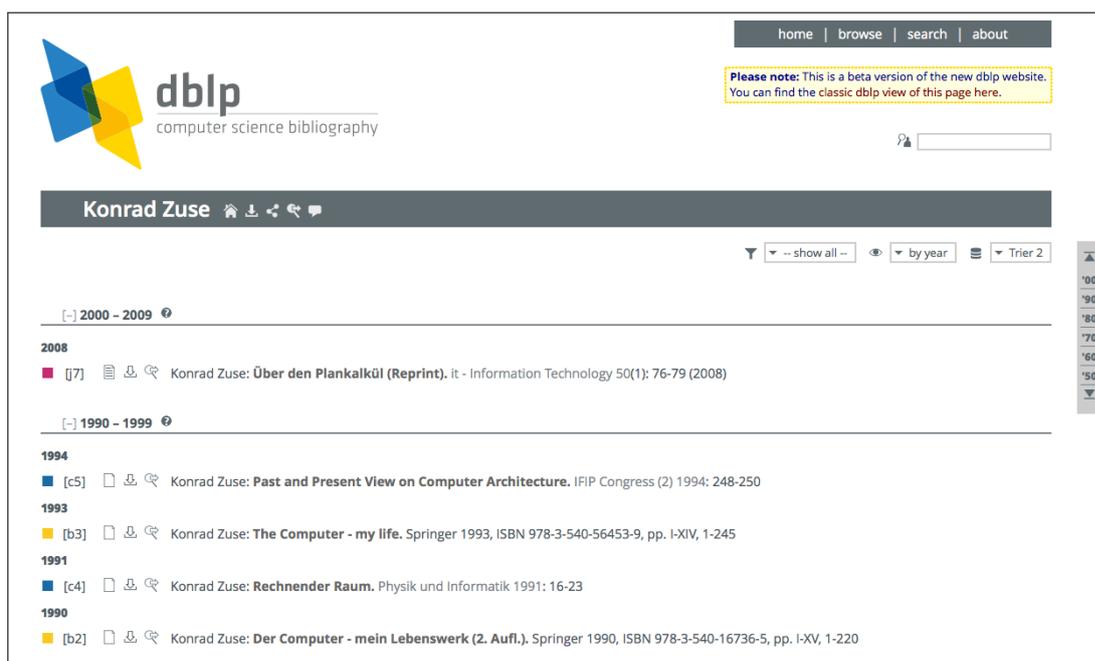


Fig. 7.4
Screenshot of the new dblp webdesign

Des Weiteren wurde die interne Datenspeicherung und -aggregation von dblp von Grund auf überarbeitet. Der komplette Datenbestand wurde auf ein valides XML-Format umgestellt und eine Reihe von historisch bedingten Problemfällen konnten aufgelöst werden. Die Routinen zur internen Datenaggregation wurden in Java neu geschrieben, um so eine bessere Wartbarkeit und Erweiterbarkeit zu gewährleisten. Die vormals statischen dblp-Autorenprofile werden nun mittels CGI-Skripte instanziiert.

Zudem wurde zur besseren Lastverteilung ein weiterer dblp-Mirror-Dienst in Schloss Dagstuhl aufgesetzt (<http://dblp.dagstuhl.de>). Es ist beabsichtigt, dass dieser Server zusammen mit dem etablierten Server in Trier (<http://dblp.uni-trier.de>) transparent unter der zukünftig primären Webadresse dblp.org fungieren soll. Die neue URL soll in 2013 stärker in den Vordergrund gerückt werden; die alten URLs werden jedoch weiterhin erreichbar bleiben.

■ Entwicklung einer bibliometrischen Infrastruktur

In seiner Sitzung im November 2009 hat sich der Fakultätentag Informatik dafür ausgesprochen, den dblp Datenbestand als Grundlage für scientometrische Analysen innerhalb der Informatik heranzuziehen. Als notwendigen ersten Schritt bedarf es dazu jedoch zunächst einer Bestandsaufnahme der tatsächlichen Abdeckung der Kern-Informatik-Publikationsreihen innerhalb von dblp. Zu diesem Zweck ist eine breit angelegte Befragung unter den deutschen Informatik-Forschern geplant. Die Ergebnisse der Studie sollen zudem dabei helfen, eine Bewertung der verschiedenen Konferenz- und Journalreihen aus Sicht der deutschen Informatik zu erarbeiten sowie die Prioritäten bei der Aufnahme neuer Reihen in dblp auszurichten.

Die Befragung wird dabei in Zusammenarbeit mit der Gesellschaft für Informatik e.V. (GI) und dem Fakultätentag Informatik organisiert. Die wissenschaftliche Aufsicht über die Befragung wird von einem Editorial Board, bestehend aus Prof. Dr. Hannah Bast (dblp-Beirat), Prof. Dr. Peter Liggesmeyer (GI), Prof. Dr. Rüdiger Reischuk (Fakultätentag Informatik), Prof. Dr. Reinhard Wilhelm (Schloss Dagstuhl) und Dr. Michael Wagner (dblp-Team), wahrgenommen. Das Board entwickelt 2012 einen Katalog an Kriterien, an Hand derer die Ergebnisse bewertet werden sollen. Ausschlaggebend hierfür soll dabei die wissenschaftliche Qualität sowie die Relevanz innerhalb der internationalen Informatik-Forschung sein. Reihen mit eher nationaler Bedeutung oder kleinere Spezialworkshops stehen nicht im Mittelpunkt der Studie.

Die Ergebnisse und Rohdaten der Studie sollen veröffentlicht werden, um ein Höchstmaß an Transparenz zu ermöglichen. Die Daten werden zudem intensiv vom dblp-Team ausgewertet werden, um Erkenntnisse in die Steigerung der Qualität des Webdienstes einfließen zu lassen. Es ist geplant, die Befragung in einigem zeitlichen Abstand periodisch zu wiederholen und somit zu verstetigen.

Furthermore, the data storage and aggregation infrastructure of the dblp back end was overhauled. The whole data stock is now stored as valid XML and several inherited quirks of the storage format have been abandoned. The routines and scripts for the websites data aggregation were entirely rewritten in Java to benefit maintainability and extensibility. The HTML of the author profile pages is now instanced using CGI scripts.

For the sake of load balancing and failure safety, a new dblp server was installed at Schloss Dagstuhl (<http://dblp.dagstuhl.de>) in 2012. It is intended that this server, together with the established server at the University of Trier (<http://dblp.uni-trier.de>), will act as the primary dblp web servers using the common domain dblp.org. This new URL will be promoted more actively starting in 2013, although all current URLs will remain active.

■ Development of a Bibliometric Infrastructure

In November 2009, the *Fakultätentag Informatik* declared that it would consider using dblp as a basis for scientometric analyses in computer science. As a mandatory first step towards this goal, the adequacy of dblp's coverage of the relevant core computer science publications will be evaluated by means of a forthcoming large-scale survey of German computer scientists. The results of this survey will also be used to develop a rating of international computer science conferences and journals from a German perspective, and to assign priorities to possible further additions to dblp.

The survey will be organized in cooperation with the Gesellschaft für Informatik e.V. (GI) and the Fakultätentag Informatik and supervised by an editorial board consisting of Prof. Dr. Hannah Bast (dblp Advisory Board), Prof. Dr. Peter Liggesmeyer (GI), Prof. Dr. Rüdiger Reischuk (Fakultätentag Informatik), Prof. Dr. Reinhard Wilhelm (Schloss Dagstuhl), and Dr. Michael Wagner (dblp team). The board developed in 2012 a catalogue of criteria for conducting the survey and interpreting its results. Primary criteria for the interpretation will be the scientific quality of a venue and the importance of its publications to the international computer science community. Venues with a national focus or small community workshops will be excluded from this survey.

It is planned to publish survey results, ratings, and catalogue of criteria in order to guarantee a maximum of transparency. The results will also be exhaustively studied by the dblp team to help improve the overall quality of its service. It is intended that the survey will be repeated periodically in the future.

■ Kolloquiumsreihe Bibliometrie

Als Ausgangspunkt für die Erörterung einer bibliometrischen Infrastruktur wurde von Schloss Dagstuhl in Zusammenarbeit mit der Universität des Saarlandes und der Universität Trier eine Kolloquiumsreihe zum Thema “Bibliometrie in der Informatik” organisiert. Die Reihe fand an der Universität des Saarlandes statt und umfasste die Vorträge wie in Fig. 7.5 aufgeführt.

Die Vorträge waren sehr gut besucht (im Schnitt etwa 30–40 Personen) und boten die Möglichkeit für interessante Diskussionen. Eine Fortführung der Reihe ist zunächst jedoch nicht in Planung.

■ Colloquia Series on Bibliometrics

As a starting point for the investigations on bibliometrics, Schloss Dagstuhl organized the colloquia series “Bibliometrics in Computer Science” during the summer of 2012 in cooperation with the University of Trier/dblp and Saarland University. The series took place at Saarland University and included the presentations listed in Fig. 7.5.

All presentations were well received by the audience (30–40 persons on average) and ended in fruitful discussions. However, the continuation of the colloquia series is currently not intended.

Stefanie Hausteil (FZ Jülich) | May 7, 2012

Möglichkeiten und Grenzen bibliometrischer Analysen

Peter van den Besselaar (VU Amsterdam) | June 4, 2012

If traditional bibliometrics does not work in computer science, what is the alternative?

Debora Weber-Wulff (HTW Berlin) | July 16, 2012

Pseudokonferenzen, Pseudozeitschriften, Pseudoverlage

Fig. 7.5

Talks in the colloquia series on bibliometrics

Ausblick

7.3

Outlook

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 dblp computer science bibliography, 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-Language-Processing (NLP) Forschungsgruppe des HITS um Prof. Dr. Michael Strube bringt dabei ihre Erfahrung mit graph- und netzwerk-basierten NLP-Methoden bei der Co-Referenz-Resolution und der Konzept- bzw. Entitäts-Disambiguierung ein.

2012 wurde ein Projektantrag vorbereitet, der 2013 im Leibniz-Wettbewerbsverfahren in der Förderlinie “Nationale und internationale Vernetzung” eingereicht wurde. Das Projekt wird von Schloss Dagstuhl (dblp-Team) koordiniert und soll von 2014 bis 2017 laufen.

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, dblp 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 Prof. Dr. Michael Strube, joins the project by providing its extensive experience with graph-based and network methods for NLP tasks such as coreference resolution, cross-document coreference resolution, concept and entity disambiguation.

In 2012, a project proposal was prepared for submission to the “National and international networking” SAW funding line of the Leibniz Association in early 2013. The project is coordinated by Schloss Dagstuhl (i.e., the dblp team) and is intended to run from 2014 to 2017.

8

Einrichtung und Service *Facilities and Services*

Tagungsräume

8.1

Conference Facilities

Schloss Dagstuhl bietet 3 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-Vorfürhungen auch verteilter Systeme optimal unterstützt. Mittels einem Presenter können Vortragende ihre vorbereiteten Materialien präsentieren, ohne zum Laptop oder Arbeitsplatz zurückzukehren.

Neben den Hörsälen bietet Dagstuhl 6 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.

The center 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 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 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.

Computer und Vernetzung

8.2

Computers and Networks

Schloss Dagstuhl bietet seinen Gästen eine adäquate Anbindung an das Internet. In 2012 erfolgte die Anbindung an das Internet über das DFN mit zwei Leitungen, einer 100Mbit/s Leitung sowie als Backup eine 20Mbit/s Leitung. 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, deren Zugangsdaten jeder Gast auf seinem Namensschild findet. 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. Nachdem wegen stark zurückgegangener Nachfrage 2012 ein Rechnerraum zu einem weiteren Seminarraum umgebaut wurde, bietet Schloss Dagstuhl momentan noch einen Rechnerraum an. In diesem stehen 9 feste Arbeitsplätze zur Verfügung. Davon sind 3 Arbeitsplätze mit Apple Macs ausgerüstet, 2 sind dedizierte MS-Windows-Arbeitsplätze. 4 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 Mac-

Schloss Dagstuhl offers its guests an adequate connection to the Internet. In 2012 the center was connected to the Internet by the DFN (German Research Network) using a 100 Mbit/s line and a 20 Mbit/s backup line, which is accessible via Wi-Fi (IEEE 802.11 b,g,n) throughout the grounds. Access is either via EDUROAM or a Dagstuhl-hosted private account whose access data is printed on the personal name tag of each guest. 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. Due to the decreasing demand for workstations among our guests, one of the two computer rooms was converted to a meeting room in 2012. The remaining room includes 9 workstations with 3 Apple Macs, 2 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

book 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 5 Recherche-Arbeitsplätze in der Bibliothek sowie 3 fest installierte Rechner in den Seminarräumen.

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

8.3

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.

Dagstuhl's Web-based Services

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.

Dagstuhl-Website

8.4

Schloss Dagstuhls Internetauftritt bietet nicht nur seinen Gästen sondern allen Nutzern 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, 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.

The Dagstuhl Website

In keeping with the center's philosophy, its Internet 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 motivation text, list of participants, publications, etc.

Freizeit und Ambiente

8.5

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 und einen Freizeitkeller mit einem Tischfußball 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, and a recreation room with gym equipment and table football facilities. 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.

Kinderbetreuung

8.6

Childcare

Viele unserer internationalen Gäste möchten 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.

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.



Fig. 8.1

Dagstuhl's baroque music room was the scene of several concerts given by our guests in 2012

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 31.12.2012 umfasste der Bibliotheksbestand 59 991 bibliographische Einheiten, der vollständig im Online-Katalog verzeichnet ist. Eine Besonderheit ist der umfangreiche Zeitschriftenbestand, der fast komplett elektronisch bezogen wird. Neben den vom Zentrum selbst 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 sowie in der Elektronischen Zeitschriftenbibliothek 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 Bibliothek sind u.a. der Online-Bibliothekskatalog, die Zeitschriftenbestandsliste mit Zugang zu den in Dagstuhl online verfügbaren Zeitschriften sowie weitere Informationsangebote der Bibliothek zu erreichen. Siehe: <http://www.dagstuhl.de/de/library/>.

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 31 December 2012, the library's assets totaled 59 991 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 the journals subscribed by the center, 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 loans. The library's entire holdings of journals and periodicals are additionally listed in the ZDB, the world's largest specialized database for serial titles, and in the EZB, the Electronic Journals Library. 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. Every week all the books authored by the participants of 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 journals as well as other information offerings can be accessed by the library's webpage. See: <http://www.dagstuhl.de/de/library/>.

Spenden an die Bibliothek

9.2

Library Donations

9

Die Bibliothek von Schloss Dagstuhl profitiert durch zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2012 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 1 176 Bände als Spenden von Verlagen und Seminarteilnehmern.

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. During 2012 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 1 176 volumes during the year 2012 as donations from publishing houses and seminar participants.

Eurographics – European Association for Computer Graphics

<https://www.eg.org>

IOS Press

<http://www.iospress.nl>

O'Reilly Media, Inc.

<http://oreilly.com>

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

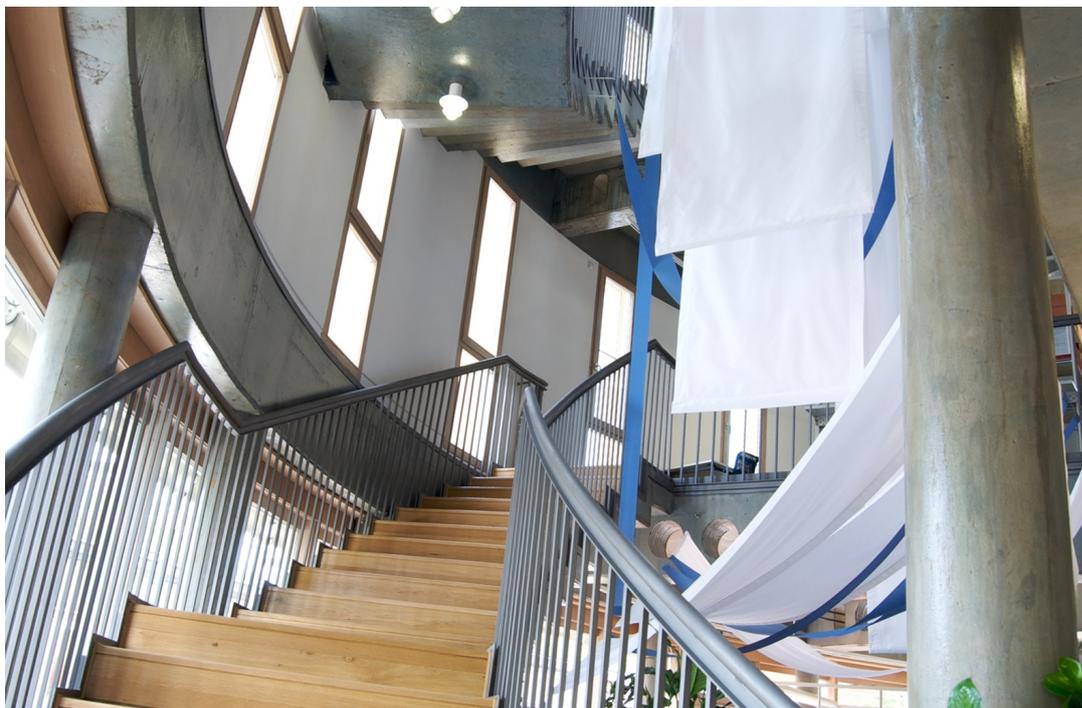


Fig. 9.2

The stairs leading up to the second-floor library stacks at Schloss Dagstuhl

10 Kunst *Art*

Dagstuhl als Galerie

10.1

Dagstuhl as Art Gallery

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Ausstellungen von Künstlern organisiert. Das großzügige Raumangebot sowie die hervorragende Ausleuchtung mit großen Kontrasten zwischen Tag und Nacht bieten den Künstlern sehr gute Möglichkeiten, ihre Werke darzustellen. Allgemein gelobt wird die Ausstrahlung, die von dem Kunstangebot ausgeht. Sie durchbricht die Nüchternheit des Neubaus in anregender und angenehmer Weise.

■ Ausstellungen 2012

Es wurden 5 Ausstellungen in 2012 organisiert, siehe Fig. 10.1. Künstlerprofile, Portfolio, Konzepte der einzelnen Ausstellungen sind erreichbar unter der Seite <http://www.dagstuhl.de/Kunst/>.

Exhibitions of artists are regularly organized in the so-called cloister of the new building. The spacious surroundings and excellent lighting provide for striking contrasts between day and night, thus offering the artists an excellent venue to exhibit their work. Compliments are frequently heard with regard to the fascinating atmosphere created by the art offerings. They provide an intriguing juxtaposition to the otherwise ascetic nature of the new building.

■ 2012 Exhibits

In 2012, Schloss Dagstuhl organized five exhibitions, see Fig. 10.1. For artist profiles, portfolios, and exhibit concepts, please refer to: <http://www.dagstuhl.de/Art/>.

Kunstankauf durch Spenden

10.2

Purchasing of Art through Donations

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 zum Spenden angeboten, dabei wird der Preis in kostengünstige Anteile aufgeteilt. Sobald ein Bild voll gezeichnet ist, werden die Teilnehmer aufgefordert, den Gegenwert der bestellten Anteile als Spende einzuzahlen, wodurch dann das Objekt angekauft werden kann. Dieses Verfahren ist gleichzeitig reizvoll für die Stifter, die sowohl in der virtuellen Internet-Galerie von Dagstuhl als auch an dem realen Objekt genannt werden, und fruchtbar für das Zentrum, da es ihm ermöglicht, von Künstlern, die Ausstellungen ausgerichtet haben, Werke anzukaufen.

Weitere Information gibt es unter <http://www.dagstuhl.de/Kunst/>. Allen Förderern sei an dieser Stelle herzlich gedankt.

Dank einer privaten Spende konnte das Selbstportrait von Octavie de Lasalle von Louisenthal, welches am "Tag der offenen Tür" im Jahr 2011 ausgestellt war, angekauft werden. Das Bild wurde sorgfältig restauriert und im Dezember 2012 an Schloss Dagstuhl übergeben; siehe Fig. 10.2.

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, with donations being made by acquiring shares at affordable prices. As soon as a picture is fully subscribed for, the donors are asked to pay in the value of the shares subscribed by them, thus enabling the art item to be purchased. This procedure provides an incentive for the donors as they are not only mentioned in Dagstuhl's online art gallery but also mentioned on the art item itself. The art donation program also benefits the center, enabling Schloss Dagstuhl to purchase works of art from the artists who arrange exhibitions there.

Further information about art can be found at <http://www.dagstuhl.de/Art/>. We would like to take this opportunity to thank all those who have made art donations in 2012.

Thanks to a private donation, the center was able to purchase in 2011 the self-portrait of Octavie de Lasalle von Louisenthal. The painting was carefully restored and presented to Schloss Dagstuhl in December of 2012, see Fig. 10.2.

Marlene Reucher Works by artist Marlene Reucher January 9 to March 9, 2012
Transit Works by artist Mane Hellenthal March 19 to April 27, 2012
Und alles voll Gefunkel ist Works by artist Birgit Ginkel May 7 to June 16, 2012
Beyond the Landscape Organized by Maxwell Roberts and Alexander Wolff as part of Dagstuhl Seminar 12261 "Putting Data on the Map" June 26 to November 9, 2012
moving source, paintings & objects Works by artists Volker Sieben and Ila Wingen ¹ November 21, 2012 to February 7, 2013

Fig. 10.1

Art exhibitions in 2012



Fig. 10.2
Self-portrait of Octavie de Lasalle von Louisenthal on her historical easel

11

**Stiftung “Informatikzentrum
Schloss Dagstuhl”**
The Dagstuhl Foundation

Zielsetzung

11.1

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 2012 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.

Zur Verwaltung von Spenden besitzt die Gesellschaft seit 1995 eine nicht rechtsfähige Stiftung "Informatikzentrum Schloss Dagstuhl", deren Vermögen sie als Sondervermögen verwaltet. Die Stiftung fördert aus Kapitalerträgen die gemeinnützigen Ziele, die im Gesellschaftsvertrag der GmbH festgeschrieben sind. Sie fördert

- die Informatikforschung auf international anerkanntem Niveau
- die interdisziplinäre Forschungsdiskussion und Forschungsk Kooperation
- 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.

Aims

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 2012 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 founded in 1995 the "Informatikzentrum Schloss Dagstuhl" foundation which is wholly part of the Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH and whose endowment Schloss Dagstuhl manages. The foundation supports identical aims as Schloss Dagstuhl:

- 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

Support of junior scientists is a special priority of the Dagstuhl Foundation.

Förderung

11.2

Ergänzend zu Einzelspenden bitten wir Personen und Firmen oder Institutionen um eine jährliche Spende.

■ Fördernde Einzelpersonen

Fördernde Einzelpersonen spenden jährlich mindestens 60 €. Sie erhalten zum Dank jährlich exklusiv eine Zusammenstellung der wissenschaftlichen Ergebnisse zurückliegender Seminare, sowie den Tätigkeitsbericht der Gesellschaft.

■ Fördernde Firmen oder Institutionen

Fördernde Firmen oder Institutionen unterstützen die Stiftung jährlich mit 600 € als Universität oder 1 200 € als Wirtschaftsunternehmen.

Contributions

In addition to individual donations, we welcome yearly donations from people and organizations who are willing to support the foundation on a regular basis.

■ Individual Donors

Individual donors contribute annually 60 € or more. As a sign of gratitude, they receive the *Dagstuhl News*, a topical summary recent seminars at Dagstuhl, as well as the Schloss Dagstuhl's annual report.

■ Institutional Donors

Institutional or company donations are defined as 600 € in the case of a university and 1 200 € in the case of a commercial company.

■ Spenden im Rahmen von Seminaren

Schloss Dagstuhl unterstützt die Verbindung zwischen Forschung und Anwendung und so nehmen von Unternehmen mit starkem Bezug zur Forschung wie Google, IBM oder Microsoft regelmäßig Mitarbeiter an Seminaren und Perspektiven-Workshops teil. Da wir erwarten, dass diese Unternehmen kommerziell von Forschungsergebnissen profitieren, bitten wir sie um eine Spende. Solche Spenden werden einem von der Firma ausgesuchten Seminar oder Workshop zugerechnet und auf der entsprechenden Webseite publiziert. Im Jahr 2012 wurde die Stiftung von zwei Firmen auf diese Weise unterstützt, siehe Fig. 11.1.

■ Donations in Conjunction with Seminars

Schloss Dagstuhl fosters exchange between research and industry. Employees from companies with a strong research departments such as Google, IBM or Microsoft often participate in Dagstuhl Seminars and Perspectives Workshops. Since we assume that these companies gain economically from the results, we ask them for a donation. Such donations are recognized on the web page of a seminar of the company's choice. In 2012 we received donations from two such companies; see Fig. 11.1.

Ausblick

11.3

Schloss Dagstuhl strebt eine stärkere Eigenständigkeit der Stiftung an und erwägt die Gründung eines Fördervereins, der die Stiftung in Zukunft führen könnte.

Outlook

Schloss Dagstuhl would like the foundation to become independent and is considering founding a legally independent structure for its management.

ESTECO, Italy

Dagstuhl Seminar 12041 | Learning in Multiobjective Optimization | January 22–27, 2012

Organizers: Salvatore Greco (Universita di Catania, IT), Joshua D. Knowles (University of Manchester, GB), Kaisa Miettinen (University of Jyväskylä and KTH Royal Institute of Technology – Stockholm), Eckart Zitzler (PH Bern, CH)

Microsoft Research, Cambridge

Dagstuhl Seminar 12161 | Abstractions for scalable multi-core computing | April 15–20, 2012

Organizers: Faith Ellen (University of Toronto, CA), Christof Fetzer (TU Dresden, DE), Tim Harris (Microsoft Research UK – Cambridge, GB), Nir Shavit (Tel Aviv University, IL)

Fig. 11.1

Donations to the Dagstuhl Foundation in 2012, in conjunction with seminars

12 **Organe und Gremien** *Dagstuhl Bodies*

Struktur der Gesellschaft

12.1

Das Zentrum 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 voor Wiskunde en 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.

Structure of the Center

The center is operated as a non-profit organization whose associates include the Gesellschaft für Informatik e.V.¹ (GI), the Universität des Saarlandes, the Technische 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 voor Wiskunde en Informatica (CWI, The Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the "Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung" (BLK)² 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"³ to "Schloss Dagstuhl – Leibniz-Zentrum für Informatik"⁴.

Dagstuhl was evaluated for the first time in July of 2009. 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.

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 Vertreter

Dagstuhl Bodies

The following six 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 (Universität

¹ engl.: German Informatics Society

² engl.: Federal Government–State Commission for Educational Planning and Research Promotion

³ engl.: International Conference and Research Center for Computer Science

⁴ engl.: Schloss Dagstuhl – Leibniz Center for Informatics

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

Nachdem Wolfgang Lorenz im Mai 2012 ausgeschieden ist, hat die Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH noch 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 einen Statusbericht nebst Stellungnahme und Empfehlungen (Audit) für den Senatsausschuss Evaluierung der Leibniz-Gemeinschaft erstellen. Der Beirat tagt einmal im Jahr zusammen mit dem industriellen Kuratorium. Siehe Fig. 12.3.

■ 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

After Wolfgang Lorenz retired from his position as Technical Administrative Director of Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH in May 2012, Scientific Director Professor Dr. Dr. h.c. Reinhard Wilhelm and Technical Administrative Director Dr. Christian Lindig continued to serve as the remaining joint representatives of the company.

■ 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 and to draft a status report including position paper and recommendations (audit) for the Senate Evaluation Committee of the Leibniz Association. The board convenes once a year together with the Industrial Curatory Board. See Fig. 12.3.

■ 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 Perspectives Work-

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. Die Mitglieder wählen aus ihrer Mitte den Wissenschaftlichen Direktor. Das Amt wird seit Bestehen des Zentrums von Professor Dr. Reinhard Wilhelm wahrgenommen.

shops 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. 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 voor Wiskunde en Informatica (CWI), Netherlands
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V., Germany

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. <i>tenure started on June 5, 2012</i>
Prof. Dr.-Ing. Dr. h.c. Theo Härder Technische Universität Kaiserslautern, Germany Representative of Technische Universität Kaiserslautern
Prof. Dr.-Ing. 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 <i>tenure started on April 30, 2012</i>
Dr. Ralph Dieter Bundesministerium für Bildung und Forschung, Bonn, Germany Representative of the German federal government <i>tenure ended on March 31, 2012</i>
Wolfgang Habelitz Ministeriums für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany Representative of the Rhineland-Palatinate
Peter Hauptmann Ministerium für Wirtschaft und Wissenschaft, Saarbrücken, Germany Representative of the Saarland state <i>tenure ended on June 5, 2012</i>
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany Representative of the Saarland state <i>tenure started on September 25, 2012</i>

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 <i>tenure started on February, 2012</i>
Prof. Dr. Claude Kirchner Institut National de Recherche en Informatique et en Automatique (INRIA), France
Prof. Dr. Jan Karel Lenstra Centrum voor Wiskunde en Informatica (CWI), Netherlands <i>tenure ended on May, 2012</i>
Prof. Dr. Mila E. Majster-Cederbaum Universität Mannheim, Germany <i>tenure started on February, 2012</i>
Prof. Dr. Dr. h.c. mult. Kurt Mehlhorn Max-Planck-Institut für Informatik, Saarbrücken, Germany <i>tenure ended on May, 2012</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
Dr. Anne Norekian Staatskanzlei des Saarlandes, Referat T/3, Saarbrücken, Germany <i>Guest</i>

Fig. 12.3

Scientific Advisory Board

Industrielles Kuratorium Industrial Curatory Board
Dr. Udo Bub Deutsche Telekom AG, Berlin, Germany
Dr. Jorge R. Cuéllar Siemens AG, CT IC 3, München, Germany
Dr.-Ing. Elmar Dörner SAP Research, Karlsruhe, Germany
Dr. Uwe Dumschlaff sd&m, Troisdorf, Germany <i>tenure ended on May 2012</i>
Dr. Jo Ebergen Oracle Labs, United States
Dr. Goetz Graefe HP Labs, United States <i>tenure started on February 2012</i>
Prof. Dr. Ralf Guido Herrtwich Daimler AG, Böblingen, Germany
Dr. Thomas Hofmann Google Research, Zürich, Switzerland
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
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>)
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 TU 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 voor Wiskunde en Informatica (CWI), Amsterdam, The Netherlands Delegate of Centrum voor Wiskunde en Informatica (CWI)
Prof. Dr. Frank Leymann Universität Stuttgart, Germany Delegate of Universität Stuttgart <i>tenure ended on 05/2012</i>
Dr. Stephan Merz INRIA – Nancy, France Delegate of INRIA
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany Delegate of Universität Stuttgart <i>tenure started on 05/2012</i>
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 <i>tenure started on 05/2012</i>
Prof. Dr. Gerhard Weikum Max-Planck-Institut für Informatik, Saarbrücken, Germany Delegate of Max-Planck-Gesellschaft <i>tenure ended on 05/2012</i>
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 University of Bologna, Italy
Prof. Dr. Jan-Olof Eklundh Royal Institute of Technology, Stockholm, Sweden
Prof. Dr. David Notkin University of Washington, United States

Fig. 12.5
Scientific Directorate

13 Statistik 2012

Statistics 2012

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.

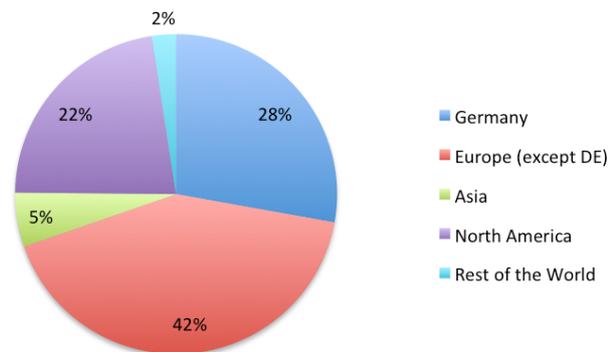
Country	A	B	Total
Argentina	5	1	6
Australia	33	6	39
Austria	60	4	64
Belgium	39	10	49
Bosnia and Herzegovina	1	0	1
Brazil	9	1	10
Canada	69	2	71
Chile	3	0	3
China	18	0	18
Croatia	3	0	3
Czech Republic	19	3	22
Denmark	19	2	21
Estonia	1	4	5
Finland	21	0	21
France	171	5	176
Georgia	1	0	1
Germany	653	893	1546
Greece	12	4	16
Hong Kong	3	0	3
Hungary	5	1	6
Iceland	1	0	1
India	21	0	21
Ireland	12	1	13
Israel	43	0	43
Italy	62	15	77
Japan	32	16	48
Lebanon	1	0	1
Luxembourg	16	39	55
Malta	2	0	2
Mexico	2	0	2
Netherlands	107	32	139
New Zealand	19	0	19
Norway	11	2	13
Oman	1	0	1
Pakistan	1	9	10
Poland	15	2	17
Portugal	19	2	21
Republic of Korea	5	0	5
Russian Federation	9	3	12
Serbia	1	0	1
Singapore	2	0	2
Slovenia	5	0	5
South Africa	5	0	5
Spain	36	2	38
Sweden	55	3	58
Switzerland	59	11	70
Thailand	0	2	2
Turkey	2	7	9
Ukraine	1	0	1
United Arab Emirates	1	0	1
United Kingdom	217	9	226
United States	436	17	453
Venezuela	2	0	2
Total	2346	1108	3454

(a)Details by country

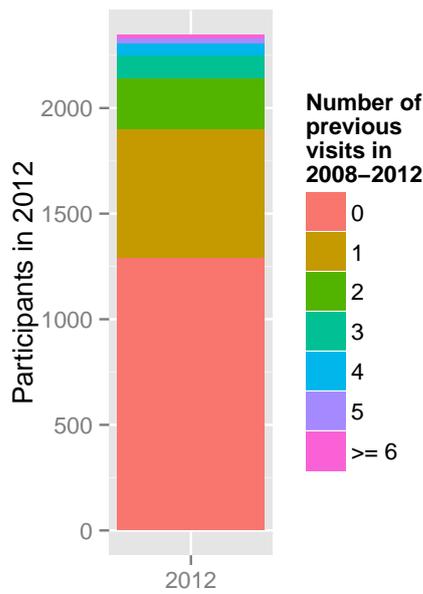
Fig. 13.1
Number of Dagstuhl guests by country of origin in 2012. A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and other events).

Region	A		B		Total	
	#	%	#	%	#	%
Germany	653	27.8	893	80.6	1546	44.8
Europe excl. Germany	982	41.9	161	14.5	1143	33.1
North and South America	526	22.4	21	1.9	547	15.8
Asia	128	5.5	27	2.4	155	4.5
Rest of world	57	2.4	6	0.6	63	1.8
Total	2346	100.0	1108	100.0	3454	100.0

(b)Details by region



(c)Graphical distribution for Dagstuhl Seminars and Dagstuhl Perspectives Workshops by region



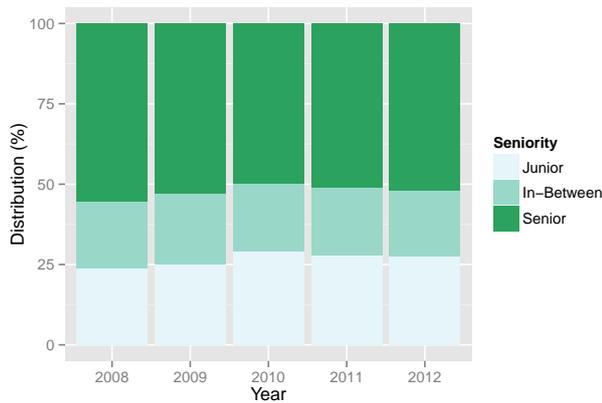
Previous visits 2008-2012	Participants	
	#	%
0	1293	55.1
1	611	26.0
2	237	10.1
3	111	4.7
4	57	2.4
5	22	0.9
≥ 6	15	0.8

(a)Graphical distribution

(b)Distribution of previous visits

Fig. 13.2

Dagstuhl participants in 2012 and their previous attendances in a Dagstuhl Seminar or Dagstuhl Perspectives Workshop from 2008 to 2012. In 2012, 55.1% of the participants took part in a seminar for the first time.



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

(a)Graphical distribution

(b)Detailed numbers

Fig. 13.3

Self-assigned seniority of Dagstuhl Seminar participants

Year	Small	Large	Total
2008	14	38	52
2009	22	38	60
2010	19	40	59
2011	18	34	55
2012	22	42	64

Fig. 13.4

Small vs. large Dagstuhl Seminars that took place in 2012. 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

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

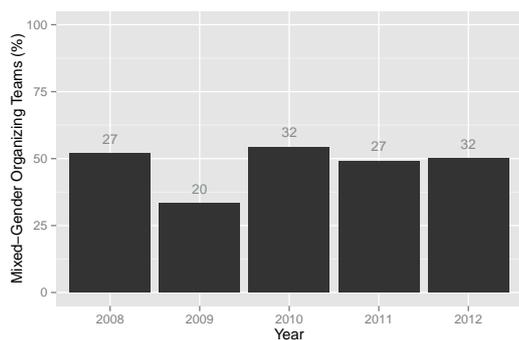
Fig. 13.6
Number of participants by event type and year. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl-Seminar, EDU = educational event, OE = other event.

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

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	69.2	29	30.9	
2011	80	54	67.5	26	32.5	
2012	90	69	76.7	21	23.3	

Fig. 13.8
Dagstuhl Seminar proposals and acceptance rates



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

(a) Graphical distribution

(b) Detailed numbers

Fig. 13.9

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	295	13.6
2010	2053	293	14.3
2011	1958	294	15.0
2012	2346	377	16.1

Fig. 13.10

Female participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops by year

Year	Invitees	Female Invitees	Decliners	Female Decliners
2008	4268	593	2467	349
2009	4671	643	2503	348
2010	4499	630	2446	337
2011	4223	599	2265	305
2012	5033	821	2687	444

Fig. 13.11

Gender of Dagstuhl Seminar invitees and decliners

14 **Veranstaltungen 2012** *Schedule of Events 2012*

Dagstuhl-Seminare**14.1****Dagstuhl Seminars****12011 – Foundations for Scripting Languages**

Robert Hirschfeld (Hasso-Plattner-Institut – Potsdam, DE), Shriram Krishnamurthi (Brown University – Providence, US), Jan Vitek (Purdue University, US)

January 2 – 6, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12011>

12021 – Computability, Complexity and Randomness

Veronica Becher (University of Buenos Aires, AR), Laurent Bienvenu (University Paris-Diderot, FR), Rodney Downey (Victoria University – Wellington, NZ), Elvira Mayordomo (University of Zaragoza, ES)

January 8 – 13, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12021>

12031 – Symmetric Cryptography

Frederik Armknecht (Universität Mannheim, DE), Stefan Lucks (Bauhaus-Universität Weimar, DE), Bart Preneel (KU Leuven, BE), Phillip Rogaway (University of California – Davis, US)

January 15 – 20, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12031>

12401 – Learning in Multiobjective Optimization

Salvatore Greco (Università di Catania, IT), Joshua D. Knowles (University of Manchester, GB), Kaisa Miettinen (University of Jyväskylä and KTH Royal Institute of Technology – Stockholm), Eckart Zitzler (PH Bern, CH)

January 22 – 27, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12401>

12051 – Analysis of Executables: Benefits and Challenges

Andy M. King (University of Kent, GB), Alan Mycroft (University of Cambridge, GB), Thomas W. Reps (University of Wisconsin – Madison, US), Axel Simon (TU München, DE), Andrea Flexeder (TWT GmbH, DE)

January 29 to February 3, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12051>

12061 – Network Attack Detection and Defense Early Warning Systems – Challenges and Perspectives

Georg Carle (TU München, DE), Hervé Debar (Télécom SudParis – Evry, FR), Hartmut König (BTU Cottbus, DE)

February 5 – 10, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12061>

12071 – Software Clone Management Towards Industrial Application

Ira D. Baxter (Semantic Designs – Austin, US), Michael Conradt (Google – München, DE), James R. Cordy (Queen's University – Kingston, CA), Rainer Koschke (Universität Bremen, DE)

February 12 – 17, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12071>

12081 – Information Visualization, Visual Data Mining and Machine Learning

Daniel A. Keim (Universität Konstanz, DE), Fabrice Rossi (Université Paris I, FR), Michel Verleysen (Université Catholique de Louvain, BE)

February 19 – 24, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12081>

12091 – Principles of Provenance

James Cheney (University of Edinburgh, GB), Bertram Ludaecher (University of California – Davis, US), Stijn Vansummeren (Université Libre de Bruxelles, BE)

February 26 to March 2, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12091>

12101 – Computation and Incentives in Social Choice

Edith Elkind (Nanyang TU – Singapore, SG), Christian Klamler (Universität Graz, AT)

March 4 – 9, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12101>

12111 – Normative Multi-Agent Systems

Giulia Andrighetto (European University Institute, IT), Guido Governatori (NICTA – St. Lucia, AU), Pablo Noriega (IIIA – CSIC – Barcelona, ES), Leon van der Torre (University of Luxembourg, LU)

March 11 – 16, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12111>

12121 – Applications of Combinatorial Topology to Computer Science

Lisbeth Fajstrup (Aalborg University, DK), Dmitry Feichtner-Kozlov (Universität Bremen, DE), Maurice Herlihy (Brown University – Providence, US)

March 18 – 23, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12121>

12131 – Open Models as a Foundation of Future Enterprise Systems

Robert B. France (Colorado State University, US), Ulrich Frank (Universität Duisburg-Essen, DE), Andreas Oberweis (KIT – Karlsruhe Institute of Technology, DE), Matti Rossi (Aalto University, FI), Stefan Strecker (FernUniversität in Hagen, DE)

March 25 – 30, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12131>

12151 – Touching the 3rd Dimension

Jean-Baptiste de la Rivière (Immersion SAS – Bordeaux, FR), Daniel Keefe (University of Minnesota, US), Antonio Krüger (DFKI – Saarbrücken, DE), Frank Steinicke (Universität Würzburg, DE)

April 9 – 12, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12151>

12152 – Software Synthesis

Rastislav Bodik (University of California – Berkeley, US), Sumit Gulwani (Microsoft – Redmond, US), Eran Yahav (Technion – Haifa, IL)

April 9 – 13, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12152>

12161 – Abstractions for scalable multi-core computing

Faith Ellen (University of Toronto, CA), Christof Fetzer (TU Dresden, DE), Tim Harris (Microsoft Research UK – Cambridge, GB), Nir Shavit (Tel Aviv University, IL)

April 15 – 20, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12161>

12171 – Semantic Data Management

Karl Aberer (EPFL – Lausanne, CH), Grigoris Antoniou (University of Huddersfield, GB), Oscar Corcho (Univ. Politec. de Madrid, ES), Rudi Studer (KIT – Karlsruhe Institute of Technology, DE), Elena Simperl (KIT – Karlsruhe Institute of Technology, DE)

April 22 – 27, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12171>

12181 – Quality of Experience: From User Perception to Instrumental Metrics

Markus Fiedler (Blekinge Institute of Technology – Karlskrona, SE), Sebastian Möller (TU Berlin, DE), Peter Reichl (FZ Telekommunikation Wien, AT)

May 1 – 4, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12181>

12191 – Artificial and Computational Intelligence in Games

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

May 6 – 11, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12191>

12221 – Cognitive Approaches for the Semantic Web

Dedre Gentner (Northwestern University – Evanston, US), Pascal Hitzler (Wright State University – Dayton, US), Kai-Uwe Kühnberger (Universität Osnabrück, DE), Frank van Harmelen (VU – Amsterdam, NL), Krzysztof Janowicz (University of California – Santa Barbara, US)

May 28 to June 1, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12221>

12231 – Future Internet for eHealth

Paolo Bonato (Harvard Medical School – Boston, US), Markus Fiedler (Blekinge Institute of Technology – Karlskrona, SE), Katarzyna Wac (University of Geneva, CH)

June 3 – 6, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12231>

12241 – Data Reduction and Problem Kernels

Michael R. Fellows (Charles Darwin University – Darwin, AU), Jiong Guo (Universität des Saarlandes, DE), Dániel Marx (MTA – Budapest, HU), Saket Saurabh (University of Bergen, NO)

June 10 – 15, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12241>

12261 – Putting Data on the Map

Stephen G. Kobourov (University of Arizona – Tucson, US), Frank van Ham (IBM Software Group – Netherlands, NL), Alexander Wolff (Universität Würzburg, DE)

June 24 – 29, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12261>

12271 – AI meets Formal Software Development

Alan Bundy (University of Edinburgh, GB), Dieter Hutter (DFKI Bremen, DE), Cliff B. Jones (Newcastle University, GB), J Strother Moore (University of Texas at Austin, US)

July 1 – 6, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12271>

12272 – Architecture-Driven Semantic Analysis of Embedded Systems

Jérôme Hugues (ISAE – Toulouse, FR), Oleg Sokolsky (University of Pennsylvania, US)

July 1 – 6, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12272>

12281 – Security and Dependability for Federated Cloud Platforms

Matthias Schunter (INTEL ICRI – Darmstadt, DE), Marc Shapiro (UPMC, Lab. LIP6 – Paris, FR), Paulo Jorge Verissimo (University of Lisboa, PT), Michael Waidner (TU Darmstadt, DE), Rüdiger Kapitza (TU Braunschweig, DE)

July 8 – 13, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12281>

12282 – Database Workload Management

Shivnath Babu (Duke University – Durham, US), Goetz Graefe (HP Labs – Madison, US), Harumi Anne Kuno (HP Labs – Palo Alto, US)

July 8 – 13, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12282>

12291 – Structure Discovery in Biology: Motifs, Networks and Phylogenies

Alberto Apostolico (Georgia Institute of Technology, US), Andreas Dress (Shanghai Institutes for Biological Sciences, CN), Laxmi Parida (IBM TJ Watson Research Center – Yorktown Heights, US)

July 15 – 20, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12291>

12321 – Robust Query Processing

Goetz Graefe (HP Labs – Madison, US), Wey Guy (Independent, US), Glenn Pauley (Conestoga College – Kitchener, CA), Harumi Anne Kuno (HP Labs – Palo Alto, US)

August 5 – 10, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12321>

12331 – Mobility Data Mining and Privacy

Christopher W. Clifton (Purdue University, US), Bart Kuijpers (Hasselt University – Diepenbeek, BE)

August 12 – 17, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12331>

12341 – Verifying Reliability

Görschwin Fey (Universität Bremen, DE), Masahiro Fujita (University of Tokyo, JP), Kaushik Roy (Purdue University, US), Matteo Sonza Reorda (Politecnico di Torino, IT)

August 19 – 24, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12341>

12342 – Engineering Multi-Agent Systems

Jürgen Dix (TU Clausthal, DE), Koen V. Hindriks (TU Delft, NL), Brian Logan (University of Nottingham, GB), Wayne Wobcke (UNSW – Sydney, AU)

August 19 – 24, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12342>

12351 – Interaction Beyond the Desktop

Alan Dix (University of Birmingham, GB), James D. Hollan (UC – San Diego, US), Albrecht Schmidt (Universität Stuttgart, DE), Jürgen Steimle (MIT – Cambridge, US)

August 26 – 31, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12351>

12352 – Information Flow and Its Applications

Samson Abramsky (University of Oxford, GB), Jean Krivine (University Paris-Diderot, FR), Michael W. Mislove (Tulane University, US)

August 26 – 31, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12352>

12361 – Information-centric networking – Ready for the real world?

Börje Ohlman (Ericsson Research – Stockholm, SE), Ignacio Solis (PARC – Palo Alto, US)

September 2 – 5, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12361>

12362 – The Multilingual Semantic Web

Paul Buitelaar (National University of Ireland – Galway, IE), Key-Sun Choi (KAIST – Daejeon, KR), Philipp Cimiano (Universität Bielefeld, DE)

September 2 – 7, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12362>

12363 – Software Defined Networking

Pan Hui (T-labs/TU Berlin, DE), Teemu Koponen (Nicira Networks Inc. – Palo Alto, US)

September 5 – 8, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12363>

12372 – Biological Data Visualization

Carsten Goerg (University of Colorado, US), Lawrence Hunter (University of Colorado, US), Jessie Kennedy (Edinburgh Napier University, GB), Sean O'Donoghue (CSIRO – North Ryde, AU), Jarke J. Van Wijk (TU Eindhoven, NL)

September 9 – 14, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12372>

12381 – Privacy-Oriented Cryptography

Jan Camenisch (IBM Research – Zürich, CH), Mark Manulis (University of Surrey, GB), Gene Tsudik (University of California – Irvine, US), Rebecca Wright (Rutgers University – Piscataway, US)

September 16 – 21, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12381>

12391 – Algorithms and Complexity for Continuous Problems

Frances Kuo (UNSW – Sydney, AU), Joseph F. Traub (Columbia University, US)

September 23 – 28, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12391>

12401 – Web Application Security

Lieven Desmet (KU Leuven, BE), Martin Johns (SAP Research – Karlsruhe, DE), Benjamin Livshits (Microsoft – Redmond, US), Andrei Sabelfeld (Chalmers UT – Göteborg, SE)

September 30 to October 5, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12401>

12411 – Coalgebraic Logics

Ernst-Erich Doberkat (TU Dortmund, DE), Alexander Kurz (University of Leicester, GB)

October 7 – 12, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12411>

12421 – Algebraic and Combinatorial Methods in Computational Complexity

Manindra Agrawal (Indian Inst. of Technology – Kanpur, IN), Thomas Thierauf (Hochschule Aalen, DE), Christopher Umans (CalTech – Pasadena, US)

October 14 – 19, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12421>

12431 – Time-of-Flight Imaging: Algorithms, Sensors and Applications

James E. Davis (University of California – Santa Cruz, US), Bernd Jähne (Universität Heidelberg, DE), Andreas Kolb (Universität Siegen, DE), Ramesh Raskar (MIT – Cambridge, US), Christian Theobalt (MPI für Informatik – Saarbrücken, DE)

October 21 – 26, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12431>

12441 – Foundations and Challenges of Change and Evolution in Ontologies

James P. Delgrande (Simon Fraser University – Burnaby, CA), Thomas Meyer (CSIR Meraka & University of KwaZulu-Natal, ZA), Ulrike Sattler (University of Manchester, GB)

October 28 to November 2, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12441>

12442 – Requirements Management – Novel Perspectives and Challenges

Jane Huang (DePaul University – Chicago, US), Matthias Jarke (RWTH Aachen, DE), Lin Liu (Tsinghua University Beijing, CN), Kalle Lyytinen (Case Western Reserve University – Cleveland, US)

October 28 – 31, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12442>

12451 – The Constraint Satisfaction Problem: Complexity and Approximability

Johan Hastad (KTH Stockholm, SE), Andrei Krokhin (University of Durham, GB), Dániel Marx (MTA – Budapest, HU)

November 4 – 9, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12451>

12461 – Games and Decisions for Rigorous Systems Engineering

Nikolaj Björner (Microsoft – Redmond, US), Krishnendu Chatterjee (IST Austria – Klosterneuburg, AT), Laura Kovacs (TU Wien, AT), Rupak Majumdar (MPI for Software Systems – Kaiserslautern, DE)

November 11 – 16, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12461>

12462 – Symbolic Methods for Chemical Reaction Networks

Francois Boulrier (Université de Lille I, FR), Anne J. Shiu (University of Chicago, US), Thomas Sturm (MPI für Informatik – Saarbrücken, DE), Andreas Weber (Universität Bonn, DE)

November 11 – 16, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12462>

12471 – SAT Interactions

Nadia Creignou (Université de Marseille, FR), Nicola Galesi (University of Rome “La Sapienza”, IT), Oliver Kullmann (Swansea University, GB), Heribert Vollmer (Leibniz Universität Hannover, DE)

November 18 – 23, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12471>

12472 – Is the Future of Preservation Cloudy?

Erik Elmroth (University of Umeå, SE), Michael Factor (IBM – Haifa, IL), Ethan Miller (University of California – Santa Cruz, US), Margo Seltzer (Harvard University, US)

November 18 – 21, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12472>

12481 – Quantitative Security Analysis

Boris Köpf (IMDEA Software Institute, ES), Pasquale Malacaria (Queen Mary University of London, GB), Catuscia Palamidessi (Ecole Polytechnique – Palaiseau, FR)

November 25 – 30, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12481>

12482 – Analysis of Security APIs

Mike Bond (University of Cambridge, GB), Riccardo Focardi (Università Ca’ Foscari di Venezia, IT), Sibylle Fröschle (OFFIS – Oldenburg, DE), Graham Steel (ENS – Cachan, FR)

November 25 – 28, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12482>

12491 – Interpreting Observed Action

Susanne Biundo-Stephan (Universität Ulm, DE), Hans Werner Guesgen (Massey University, NZ), Joachim Hertzberg (Universität Osnabrück, DE), Stephen R. Marsland (Massey University, NZ)

December 2 – 7, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12491>

12492 – Human Activity Recognition in Smart Environments

James L. Crowley (INRIA Rhône-Alpes, FR), Paul Lukowicz (DFKI – Kaiserslautern, DE), Albrecht Schmidt (Universität Stuttgart, DE), Kai Kunze (Osaka Prefecture University, JP)

December 2 – 7, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12492>

12501 – Organizational Processes for Supporting Sustainable Security

Lizzie Coles-Kemp (RHUL – London, GB), Carrie Gates (CA Labs -New York, US), Dieter Gollmann (TU Hamburg-Harburg, DE)

December 9 – 12, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12501>

12502 – Securing Critical Infrastructures from Targeted Attacks

Marc Dacier (Symantec Research Labs – Sophia Antipolis, FR), Frank Kargl (Uni Twente, NL & Uni Ulm, DE), Alfonso Valdes (University of Illinois – Urbana, US)

December 9 – 12, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12502>

12511 – Divide and Conquer: the Quest for Compositional Design and Analysis

Marieke Huisman (University of Twente, NL), Barbara Jobstmann (VERIMAG – Gières, FR), Ina Schaefer (TU Braunschweig, DE), Marielle Stoelinga (University of Twente, NL)

December 16 – 21, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12511>

12512 – Representation, Analysis and Visualization of Moving Objects

Joachim Gudmundsson (The University of Sydney, AU), Patrick Laube (Universität Zürich, CH), Emiel Van Loon (University of Amsterdam, NL)

December 16 – 21, 2012 | Dagstuhl Seminar | <http://www.dagstuhl.de/12512>

Dagstuhl-Perspektiven- Workshops

14.2

Dagstuhl Perspectives Workshops

12182 – Social, Supply-Chain, Administrative, Business, Commerce, Political networks: a multi-discipline perspective

Matthias Häsel (Otto Group – Hamburg, DE), Thorsten Quandt (Universität Münster, DE), Gottfried Vossen (Universität Münster, DE)

May 1 – 4, 2012 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/12182>

12212 – Co-Design of Systems and Applications for Exascale

Arndt Bode (TU München & LRZ Garching, DE), Adolfy Hoisie (Pacific Northwest National Lab., US), Dieter Kranzlmüller (LMU München & LRZ Garching, DE), Wolfgang E. Nagel (TU Dresden, DE)

May 20 – 25, 2012 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/12212>

12371 – Machine Learning Methods for Computer Security

Anthony D. Joseph (University of California – Berkeley, US), Pavel Laskov (Universität Tübingen, DE), Blaine Nelson (Universität Tübingen, DE), Fabio Roli (Università di Cagliari, IT), Doug Tygar (University of California – Berkeley, US)

September 9 – 14, 2012 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/12371>

12382 – Computation and Palaeography: Potentials and Limits

Tal Hassner (Open University – Israel, IL), Peter A. Stokes (King's College, DDH – London, GB), Lior Wolf (Tel Aviv University, IL)

September 18 – 21, 2012 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/12382>

12452 – Publication Culture in Computing Research

Kurt Mehlhorn (MPI für Informatik – Saarbrücken, DE), Moshe Y. Vardi (Rice University, US), Marc Herbstritt (Schloss Dagstuhl, DE)

November 6 – 9, 2012 | Dagstuhl Perspectives Workshop | <http://www.dagstuhl.de/12452>

GI-Dagstuhl-Seminare

14.3

GI-Dagstuhl Seminars

12211 – Quality-of-Service Attributes in Service- and Cloud-based Systems: Specification, Modelling, Monitoring, Prediction, and Optimisation

Lars Grunske (Universität Stuttgart, DE), Samuel Kounev (KIT – Karlsruhe Institute of Technology, DE), Ina Schaefer (TU Braunschweig, DE)

May 20 – 25, 2012 | GI-Dagstuhl Seminar | <http://www.dagstuhl.de/12211>

12222 – Multi-Core Software Performance Engineering

Steffen Becker (Universität Paderborn, DE), Jens Happe (SAP Research – Karlsruhe, DE), Heiko Koziolk (ABB Forschungszentrum Deutschland – Ladenburg, DE), Petr Tuma (Charles University – Prague, CZ)
 May 28 – 31, 2012 | GI-Dagstuhl Seminar | <http://www.dagstuhl.de/12222>

Lehrveranstaltungen

14.4

Educational Events**12232 – Workshop Wissenschaftsjournalismus “Schreiben über Informatik”**

Roswitha Bardohl (Schloss Dagstuhl, DE), Gordon Bolduan (Universität des Saarlandes, DE), Tim Schröder (Oldenburg, DE)

June 3 – 6, 2012 | Meeting | <http://www.dagstuhl.de/12232>

12292 – Summer School SecVote 2012

Hugo Jonker (University of Luxembourg, LU), Olivier Pereira (UC Louvain, BE), Mark D. Ryan (University of Birmingham, GB), Peter Y. A. Ryan (University of Luxembourg, LU)

July 15 – 20, 2012 | Meeting | <http://www.dagstuhl.de/12292>

12402 – Herbstschule “Information Retrieval”

Andreas Henrich (Universität Bamberg, DE)

September 30 to October 5, 2012 | Meeting | <http://www.dagstuhl.de/12402>

12503 – Lehrerfortbildung in Informatik

Roswitha Bardohl (Schloss Dagstuhl, DE), Heinz Dabrock (LPM Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)

December 12 – 14, 2012 | Meeting | <http://www.dagstuhl.de/12503>

Sonstige Veranstaltungen

14.5

Other Events**12032 – Klausurtagung GESIS**

York Sure-Vetter (GESIS – Leibniz Institute for the Social Sciences, DE)

January 15 – 18, 2012 | Meeting | <http://www.dagstuhl.de/12032>

12042 – Erneuerbare Mobilität – Konzepte für die Mobilität von Morgen

Frank Geraets (Deutsche Bahn, DE), Karl-Heinz Krempels (RWTH Aachen, DE), Christoph Terwelp (RWTH Aachen, DE)

January 22 – 25, 2012 | Meeting | <http://www.dagstuhl.de/12042>

12043 – Treffen ZBMATH & DBLP

Marcel R. Ackermann (Schloss Dagstuhl, DE), Marc Herbstritt (Schloss Dagstuhl, DE)

January 25 – 26, 2012 | Meeting | <http://www.dagstuhl.de/12043>

12053 – Projekt-Implementierungs-Treffen

H. Dieter Rombach (Fraunhofer IESE – Kaiserslautern, DE)

February 1 – 2, 2012 | Meeting | <http://www.dagstuhl.de/12053>

12054 – Studententag “Personal Branding und neue Medien”

Enrico Lieblang (HTW – Saarbrücken, DE), Nicole Schwarz (HTW – Saarbrücken, DE)

February 3, 2012 | Meeting | <http://www.dagstuhl.de/12054>

12062 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)

Holger Giese (Hasso-Plattner-Institut – Potsdam, DE), Michaela Huhn (TU Clausthal, DE), Bernhard Schätz (fortiss GmbH – München, DE)

February 6 – 8, 2012 | Meeting | <http://www.dagstuhl.de/12062>

12072 – Erneuerbare Mobilität – Angebote für die Mobilität von Morgen

Karl-Heinz Krempels (RWTH Aachen, DE), Christoph Terwelp (RWTH Aachen, DE)

February 12 – 15, 2012 | Meeting | <http://www.dagstuhl.de/12072>

12073 – Lehrstuhltreffen AG Zeller/Hack

Andreas Zeller (Universität des Saarlandes, DE)

February 15 – 17, 2012 | Meeting | <http://www.dagstuhl.de/12073>

12074 – Lehrstuhltreffen Hanebeck

Uwe D. Hanebeck (KIT – Karlsruhe Institute of Technology, DE)

February 15 – 17, 2012 | Meeting | <http://www.dagstuhl.de/12074>**12082 – Workshop on “Quantum Computing”**

Jörg Hettel (FH Kaiserslautern-Zweibrücken, DE), Hans-Jürgen Steffens (FH Kaiserslautern-Zweibrücken, DE)

February 20 – 23, 2012 | Meeting | <http://www.dagstuhl.de/12082>**12083 – Klausurtagung: Software Systems Engineering – LST Schmid**

Klaus Schmid (Universität Hildesheim, DE)

February 21 – 24, 2012 | Meeting | <http://www.dagstuhl.de/12083>**12092 – Klausurtagung “Numerical Simulation, Optimization and High Performance Computing – Software Design for Numerical Libraries”**

Vincent Heuveline (KIT – Karlsruhe Institute of Technology, DE)

February 27 to March 2, 2012 | Meeting | <http://www.dagstuhl.de/12092>**12093 – Offsite Meeting Commercial Performance Management**

Thomas In der Rieden (T-Systems International GmbH, DE)

February 27 – 28, 2012 | Meeting | <http://www.dagstuhl.de/12093>**12102 – Deutsch-Japanisches Symposium**

Michael Beigl (KIT – Karlsruhe Institute of Technology, DE), Antonio Krüger (DFKI – Saarbrücken, DE)

March 4 – 7, 2012 | Meeting | <http://www.dagstuhl.de/12102>**12109 – Forschungsaufenthalt**

Albert Zündorf (Universität Kassel, DE)

March 5 – 9, 2012 | Meeting | <http://www.dagstuhl.de/12109>**12112 – Lehrstuhltreffen Fischlin**

Marc Fischlin (TU Darmstadt, DE)

March 11 – 14, 2012 | Meeting | <http://www.dagstuhl.de/12112>**12123 – Klausurtagung Telematik Karlsruhe**

Sören Finster (KIT – Karlsruhe Institute of Technology, DE), Martina Zitterbart (KIT – Karlsruhe Institute of Technology, DE)

March 21 – 23, 2012 | Meeting | <http://www.dagstuhl.de/12123>**12132 – Klausurtagung “LST Rannenberg”**

Kai Rannenberg (Goethe-Universität Frankfurt am Main, DE)

March 25 – 28, 2012 | Meeting | <http://www.dagstuhl.de/12132>**12134 – EUSJA Workshop: Information and Communication Technology in Europe**

Friederike Meyer zu Tittingdorf (Saarbrücken, DE)

March 29 – 30, 2012 | Meeting | <http://www.dagstuhl.de/12134>**12142 – Lehrstuhltreffen “Embedded Intelligence”**

Bernhard Sick (Universität Kassel, DE)

April 1 – 4, 2012 | Meeting | <http://www.dagstuhl.de/12142>**12143 – GIBU 2012: GI-Beirat der Universitätsprofessoren**

Gregor Snelting (KIT – Karlsruhe Institute of Technology, DE)

April 1 – 4, 2012 | Meeting | <http://www.dagstuhl.de/12143>**12144 – IMAGINE Project – Dagstuhl Meeting**

Raimund Bröchler (INTRASOFT International SA – Luxembourg, LU)

April 2 – 4, 2012 | Meeting | <http://www.dagstuhl.de/12144>**12162 – Projekttreffen “Schnittstellen in der medizinischen Versorgung und ihre Überwindung durch Einsatz innovativer Lehr-/Lern-Technologien”**

Christoph Igel (Universität des Saarlandes, DE)

April 16, 2012 | Meeting | <http://www.dagstuhl.de/12162>**12183 – Erneuerbare Mobilität – alternative Verkehrssysteme für die Mobilität von Morgen**

Frank Geraets (Deutsche Bahn, DE), Karl-Heinz Krempels (RWTH Aachen, DE), Niko Mehl (Stadtwerke Osnabrück, DE)

May 1 – 4, 2012 | Meeting | <http://www.dagstuhl.de/12183>

12184 – Treffen DBLP/GESIS

Oliver Hoffmann (Schloss Dagstuhl, DE)
 May 2, 2012 | Meeting | <http://www.dagstuhl.de/12184>

12192 – CELSTEC Retreat Meeting

Rob Koper (Open University – Heerlen, NL)
 May 6 – 11, 2012 | Meeting | <http://www.dagstuhl.de/12192>

12202 – Kolloquium zum GI Dissertationspreis 2011

Steffen Hölldobler (TU Dresden, DE)
 May 13 – 16, 2012 | Meeting | <http://www.dagstuhl.de/12202>

12203 – Doktorandenseminar

Bernhard Nebel (Universität Freiburg, DE)
 May 13 – 16, 2012 | Meeting | <http://www.dagstuhl.de/12203>

12233 – econnect Workshop – RWTH Aachen

Karl-Heinz Krempels (RWTH Aachen, DE), Christoph Terwelp (RWTH Aachen, DE)
 June 4 – 5, 2012 | Meeting | <http://www.dagstuhl.de/12233>

12242 – The Belgian School of Database Theory

Jan Paredaens (University of Antwerpen, BE), Jan Van den Bussche (Hasselt University, BE)
 June 11 – 12, 2012 | Meeting | <http://www.dagstuhl.de/12242>

12243 – IOI-Training

Christoph Weidenbach (MPI für Informatik – Saarbrücken, DE)
 June 12 – 15, 2012 | Meeting | <http://www.dagstuhl.de/12243>

12252 – Gemeinsamer Workshop der Graduiertenkollegs I: GK 1651 SOAMED et al.

Kristian Duske (TU Berlin, DE), Daniel Janusz (HU Berlin, DE), Wolfgang Reisig (HU Berlin, DE)
 June 17 – 20, 2012 | Meeting | <http://www.dagstuhl.de/12252>

12253 – Gemeinsamer Workshop der Graduiertenkollegs II: GK 1424 MUSAMA et al.

Alexander Gladisch (Universität Rostock, DE), Thomas Kirste (Universität Rostock, DE), Kristina Yordanova (Universität Rostock, DE)
 June 20 – 22, 2012 | Meeting | <http://www.dagstuhl.de/12253>

12254 – Deutsch-Pakistanischer Workshop

Karsten Berns (TU Kaiserslautern, DE)
 June 20 – 21, 2012 | Meeting | <http://www.dagstuhl.de/12254>

12255 – Treffen ZBMATH & DBLP

Marcel R. Ackermann (Schloss Dagstuhl, DE)
 June 18 – 19, 2012 | Meeting | <http://www.dagstuhl.de/12255>

12262 – Klausurtagung FG Security Engineering

Stefan Katzenbeisser (TU Darmstadt, DE), Martin Mink (TU Darmstadt, DE)
 June 24 – 27, 2012 | Meeting | <http://www.dagstuhl.de/12262>

12283 – Lehrstuhltreffen AG Wilhelm

Reinhard Wilhelm (Universität des Saarlandes, DE)
 July 12 – 13, 2012 | Meeting | <http://www.dagstuhl.de/12283>

12322 – Projekttreffen “BMBF-Antrag”

Hans Hagen (TU Kaiserslautern, DE)
 August 6 – 7, 2012 | Meeting | <http://www.dagstuhl.de/12322>

12333 – Klausurtagung LST Organic Computing und FG Softwaretechnik für Ubiquitous-Computing Anwendungen

Jörg Hähner (Universität Augsburg, DE), Sven Tomforde (Universität Augsburg, DE)
 August 13 – 16, 2012 | Meeting | <http://www.dagstuhl.de/12333>

12332 – Klausurtagung LST Stuckenschmidt

Heiner Stuckenschmidt (Universität Mannheim, DE)
 August 14 – 17, 2012 | Meeting | <http://www.dagstuhl.de/12332>

12353 – CEBug-Treffen

Bernd Becker (Universität Freiburg, DE)
 August 28 – 31, 2012 | Meeting | <http://www.dagstuhl.de/12353>

12389 – Forschungsaufenthalt

Stephan Diehl (Universität Trier, DE)

September 19 – 21, 2012 | Meeting | <http://www.dagstuhl.de/12389>**12392 – Klausurtagung “LST Freiling”**

Felix C. Freiling (Universität Erlangen-Nürnberg, DE)

September 23 – 27, 2012 | Meeting | <http://www.dagstuhl.de/12392>**12409 – Forschungsaufenthalt**

Egon Börger (University of Pisa, IT)

October 5 – 7, 2012 | Meeting | <http://www.dagstuhl.de/12409>**12412 – Klausurtagung Graduierten-Kolleg 1194**

Tobias Gädeke (KIT – Karlsruhe Institute of Technology, DE), Uwe D. Hanebeck (KIT – Karlsruhe Institute of Technology, DE)

October 8 – 10, 2012 | Meeting | <http://www.dagstuhl.de/12412>**12422 – Semantic Statistics for Social, Behavioural, and Economic Sciences: Leveraging the DDI Model for the Linked Data Web**

Richard Cyganiak (National University of Ireland – Galway, IE), Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (Population Center, University of Minnesota, US), Joachim Wackerow (GESIS – Mannheim, DE)

October 14 – 19, 2012 | Meeting | <http://www.dagstuhl.de/12422>**12432 – DDI Lifecycle: Moving Forward**

Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (Population Center, University of Minnesota, US), Mary Vardigan (University of Michigan – ICPSR, US), Joachim Wackerow (GESIS – Mannheim, DE)

October 21 – 26, 2012 | Meeting | <http://www.dagstuhl.de/12432>**12443 – Doktorandenseminar Prof. Halang**

Wolfgang A. Halang (FernUniversität in Hagen, DE)

October 31 to November 2, 2012 | Meeting | <http://www.dagstuhl.de/12443>**12474 – IT-Seminar LST Thomas Engel**

Thomas Engel (University of Luxembourg, LU)

November 22 – 23, 2012 | Meeting | <http://www.dagstuhl.de/12474>**12483 – Klausurtagung “LST Schmeck”**

Florian Allerdig (KIT – Karlsruhe Institute of Technology, DE)

November 28 – 30, 2012 | Meeting | <http://www.dagstuhl.de/12483>**12494 – Klausurtagung ICSY**

Paul Müller (TU Kaiserslautern, DE)

December 5 – 7, 2012 | Meeting | <http://www.dagstuhl.de/12494>**12504 – Vorbereitungstreffen SFB-Antrag**

Hans Hagen (TU Kaiserslautern, DE)

December 10 – 12, 2012 | Meeting | <http://www.dagstuhl.de/12504>



Mitglied der

Leibniz
Leibniz-Gemeinschaft



© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2012
<http://www.dagstuhl.de>