



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report

2018



SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2018

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Germany Telefon: +49 681 302 - 4396 Fax: +49 681 302 - 4397 E-Mail: service@dagstuhl.de
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen
Geschäftsführung	Prof. Raimund Seidel, Ph. D. Heike Meißner
Gesellschafter	Gesellschaft für Informatik e. V., Deutschland Universität des Saarlandes, Deutschland Technische Universität Kaiserslautern, Deutschland Karlsruher Institut für Technologie (KIT), Deutschland Technische Universität Darmstadt, Deutschland Universität Stuttgart, Deutschland Universität Trier, Deutschland Johann Wolfgang Goethe-Universität Frankfurt am Main, Deutschland Institut National de Recherche en Informatique et en Automatique (INRIA), Frankreich Centrum voor Wiskunde en Informatica (CWI), Niederlande Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Deutschland
Einbindung	Mitglied der Leibniz-Gemeinschaft Mitglied im Leibniz-Forschungsverbund Science 2.0
Verantwortlicher Redakteur	Prof. Raimund Seidel, Ph. D.
Redaktion	Michael Gerke, Dr. Shida Kunz, Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Heike Clemens, Dr. Michael Didas, Dr. Andreas Dolzmann, Jutka Gasiorowski, Michael Gerke, Dagmar Glaser, Oliver Hoffmann, Dr. Shida Kunz, Heike Meißner, Petra Meyer, Dr. Florian Reitz, Thomas Schillo, Raimund Seidel, Ph. D., Dr. Michael Wagner und die Autoren der Zusammenfassungen in Kapitel 6
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
Druck	COD Büroservice GmbH, Saarbrücken
Druckauflage	150
ISSN	2199-1995
Online-Version	https://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	April 2019
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	© Creative Commons BY-ND 3.0 DE

Vorwort

Foreword

Das Jahr 2018 war für Schloss Dagstuhl sehr erfolgreich. Bei den Dagstuhl Seminars und Perspective Workshops gab es mit 136 Anträgen einen neuen Rekord. 81 der Vorschläge wurden angenommen. Die dblp Datenbank sah ein weiteres Rekordjahr sowohl im Wachstum wie auch in der Benutzung. Ebenso erzielten wir weiteres Wachstum bei unseren open access Publikationen, insbesondere bei unserer Vorzeigeserie LIPIcs.

Vielleicht am wichtigsten für die Zukunft: Die letzten politischen und administrativen Hürden wurden genommen, so dass ab 1. Januar 2019 dblp offiziell Teil von Dagstuhl ist und von Dagstuhl vollständig betrieben wird.

Weiter Details und Zahlen können Sie auf den folgenden Seiten finden.

2018 was a successful year for Schloss Dagstuhl. We received a record number of 136 proposals for Dagstuhl Seminars or Perspectives Workshops of which 81 were accepted. The dblp data base saw further record growth and usage. Similarly we achieved further growth in our open access publication endeavor, in particular with our flagship conferences proceedings series LIPIcs.

Maybe most significant for the future: The final political and administrative hurdles were taken so that as of January 1st 2019 dblp is officially a part of and completely operated by Schloss Dagstuhl.

More details and numbers you can find in the following pages.

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph.D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

On behalf of the Managing Directors

Inhaltsverzeichnis

Contents

	Vorwort	
	<i>Foreword</i>	iii
1	Das Zentrum Schloss Dagstuhl	
	<i>Schloss Dagstuhl Center</i>	1
1.1	Dagstuhls Leitbild	
	<i>Dagstuhl's Mission</i>	2
1.2	Neuigkeiten in 2018	
	<i>News from 2018</i>	4
2	Seminare und Workshops	
	<i>Seminars and Workshops</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>	13
2.4	Seminar-Programm 2018	
	<i>Seminar-Programm in 2018</i>	14
2.5	Angaben zu Teilnehmern und Organisatoren	
	<i>Participant and Organizer Data</i>	15
2.6	Themen und Forschungsgebiete	
	<i>Topics and Research Areas</i>	16
2.7	Weitere Veranstaltungstypen	
	<i>Further Event Types</i>	18
2.8	Qualitätssicherung	
	<i>Quality Assurance</i>	19
2.9	Auslastung des Zentrums	
	<i>Utilization of the Center</i>	19
3	Bibliographiedatenbank dblp	
	<i>dblp computer science bibliography</i>	21
3.1	Offene Bibliographiedaten für die Informatik	
	<i>Open Bibliographic Data in Computer Science</i>	22
3.2	Schloss Dagstuhl und dblp	
	<i>Schloss Dagstuhl and dblp</i>	23
3.3	Statistiken der Datenakquise	
	<i>Data Acquisition Statistics</i>	24
3.4	Nutzungsstatistiken	
	<i>Usage Statistics</i>	25

4	Dagstuhl Publishing	
	<i>Dagstuhl Publishing</i>	27
4.1	Portfolio	
	<i>Portfolio</i>	28
4.2	Infrastruktur	
	<i>Infrastructure</i>	34
5	Resonanz	
	<i>Feedback</i>	37
5.1	Resonanz zu Seminaren und Workshops	
	<i>Feedback on Seminars and Workshops</i>	38
5.2	Resonanz zur Bibliographiedatenbank dblp	
	<i>Feedback on the dblp Computer Science Bibliography</i>	43
6	Die Seminare in 2018	
	<i>The 2018 Seminars</i>	45
6.1	Symmetric Cryptography	48
6.2	Personalized Multiobjective Optimization: An Analytics Perspective	50
6.3	Foundations of Data Visualization	52
6.4	Proof Complexity	54
6.5	Genetic Improvement of Software	60
6.6	Evidence About Programmers for Programming Language Design	64
6.7	Planning and Operations Research	65
6.8	Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization	66
6.9	Formal Methods for the Synthesis of Biomolecular Circuits	67
6.10	Data Consistency in Distributed Systems: Algorithms, Programs, and Databases	68
6.11	The Logical Execution Time Paradigm: New Perspectives for Multicore Systems	70
6.12	Scheduling	72
6.13	Dynamic Traffic Models in Transportation Science	73
6.14	Loop Optimization	74
6.15	Coding Theory for Inference, Learning and Optimization	76
6.16	Machine Learning and Model Checking Join Forces	77
6.17	Automatic Quality Assurance and Release	78
6.18	Program Equivalence	79
6.19	Blockchains, Smart Contracts and Future Applications	80
6.20	Visualization of Biological Data – Crossroads	82
6.21	Normative Multi-Agent Systems	84
6.22	Algebraic Effect Handlers go Mainstream	86
6.23	Towards Accountable Systems	88
6.24	Software Business, Platforms, and Ecosystems: Fundamentals of Software Production Research	90
6.25	Secure Compilation	92
6.26	Inter-Vehicular Communication Towards Cooperative Driving	94
6.27	Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance	96
6.28	On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces	98
6.29	The Constraint Satisfaction Problem: Complexity and Approximability	100
6.30	High-Performance Graph Algorithms	104
6.31	Secure Routing for the Internet	105
6.32	Database Architectures for Modern Hardware	106
6.33	Ubiquitous Gaze Sensing and Interaction	108
6.34	Discipline Convergence in Networked Systems	109
6.35	10 Years of Web Science: Closing The Loop	110
6.36	In Situ Visualization for Computational Science	111
6.37	Synergies between Adaptive Analysis of Algorithms, Parameterized Complexity, Compressed Data Structures and Compressed Indices	112
6.38	Extreme Classification	114
6.39	Web Application Security	116
6.40	Human-Computer Integration	118
6.41	Algorithmic Foundations of Programmable Matter	120
6.42	Blockchain Technology for Collaborative Information Systems	122

6.43	Formalization of Mathematics in Type Theory	123
6.44	Modeling for Sustainability	124
6.45	Measuring the Complexity of Computational Content: From Combinatorial Problems to Analysis	125
6.46	Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web	126
6.47	Quantum Programming Languages	128
6.48	Algebraic Methods in Computational Complexity	130
6.49	Automating Data Science	132
6.50	Progressive Data Analysis and Visualization	134
6.51	Encouraging Reproducibility in Scientific Research of the Internet	136
6.52	Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative	140
6.53	Shape Analysis: Euclidean, Discrete and Algebraic Geometric Methods	141
6.54	Computational Aspects of Fabrication	142
6.55	Data Physicalization	144
6.56	Visualization and Processing of Anisotropy in Imaging, Geometry, and Astronomy	148
6.57	Genomics, Pattern Avoidance, and Statistical Mechanics	150
6.58	Blockchain Security at Scale	151
6.59	Provenance and Logging for Sense Making	152
6.60	Next Generation Domain Specific Conceptual Modeling: Principles and Methods	154
6.61	Implementing FAIR Data Infrastructures	156
6.62	High Throughput Connectomics	158
6.63	Network Visualization in the Humanities	160
6.64	Multidirectional Transformations and Synchronisations	162
6.65	Algebraic Coding Theory for Networks, Storage, and Security	164
7	Öffentlichkeitsarbeit	
	Public Relations and Outreach	167
7.1	Pressemitteilungen und Medienarbeit <i>Press Releases and Media Work</i>	168
7.2	Fortbildung <i>Educational Training</i>	168
7.3	„Dagstuhler Gespräche“ <i>“Dagstuhler Gespräche”</i>	169
8	Einrichtungen	
	Facilities	171
8.1	Hauptstandort in Wadern <i>Main Site in Wadern</i>	172
8.2	Geschäftsstelle in Saarbrücken <i>Dagstuhl Office at Saarbrücken</i>	174
8.3	Dagstuhl an der Universität Trier <i>Dagstuhl at University of Trier</i>	175
9	Zentrale Dienste	
	Central Services	177
9.1	Bibliothek <i>Research Library</i>	178
9.2	IT-Service <i>IT Service</i>	180
10	Kunst	
	Art	181
10.1	Dagstuhl als Galerie <i>Dagstuhl as Art Gallery</i>	182
10.2	Kunstankauf durch Spenden <i>Art Sponsorship and Donations</i>	185
10.3	Dagstuhls permanente Kunstaussstellung <i>Dagstuhl's Permanent Art Exhibition</i>	185

11	Struktur der Gesellschaft	
	<i>Structure of the Company</i>	187
11.1	Gründung und Gesellschafter	
	<i>Formation and Shareholders</i>	188
11.2	Organe der Gesellschaft	
	<i>Organs of the Organization</i>	188
11.3	Gremien der Gesellschaft	
	<i>Bodies of the Organization</i>	190
12	Holger Hermanns, Erich Reindel: Förderverein „Freunde von Dagstuhl“	
	<i>Holger Hermanns, Erich Reindel: Association “Friends of Dagstuhl”</i>	195
13	Statistiken	
	<i>Statistics</i>	199
13.1	Statistiken zu Seminaren und Workshops	
	<i>Statistics on Seminars and Workshops</i>	200
13.2	Statistiken zur Bibliographiedatenbank dblp	
	<i>Statistics of the dblp computer science bibliography</i>	210
13.3	Statistiken zu Dagstuhl Publishing	
	<i>Statistics of Dagstuhl Publishing</i>	214
14	Veranstaltungen 2018	
	<i>Schedule of Events 2018</i>	217
14.1	Dagstuhl-Seminare	
	<i>Dagstuhl Seminars</i>	218
14.2	Dagstuhl-Perspektiven-Workshops	
	<i>Dagstuhl Perspectives Workshops</i>	222
14.3	GI-Dagstuhl-Seminare	
	<i>GI-Dagstuhl Seminars</i>	223
14.4	Lehrveranstaltungen	
	<i>Educational Events</i>	223
14.5	Forschungsgruppentreffen	
	<i>Research Group Meetings</i>	223

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

Dagstuhls Leitbild

1.1

Dagstuhl's Mission

Schloss Dagstuhl – Leibniz Zentrum für Informatik fördert die Informatikforschung auf internationalem Spitzenniveau durch die Bereitstellung von Infrastrukturen zur wissenschaftlichen Kommunikation und für den Austausch zwischen Forschenden. Ziel von Schloss Dagstuhl ist

- die Förderung der Grundlagenforschung und der anwendungsorientierten Forschung auf dem Gebiet der Informatik,
- die wissenschaftliche Fort- und Weiterbildung im Informatikbereich,
- der Wissenstransfer zwischen Forschung und Anwendung der Informatik,
- der Betrieb einer internationalen Begegnungs- und Forschungsstätte für die Informatik.

Die Förderung und Einbindung von Nachwuchswissenschaftlern ist dabei ein wichtiger Teil dieser Aufgabe; ebenso wie der Technologietransfer zwischen Forschung und Industrie.

■ Entwicklung des Zentrums

Die Idee zur Gründung eines Tagungszentrum für Informatik wurde Ende der 1980er Jahre geboren, zu einem Zeitpunkt, an dem die Informatikforschung – ursprünglich der Mathematik und den Ingenieurwissenschaften entsprungen – enormen Aufwind erfuhr. Die *Gesellschaft für Informatik* beobachtete damals die zunehmende Nachfrage von Informatikwissenschaftlern am weltbekannten *Mathematischen Forschungsinstitut Oberwolfach* und sah die Notwendigkeit, ein eigens auf die Informatik ausgerichtetes Zentrum einzurichten. Schloss Dagstuhl wurde schließlich 1990 gegründet und entwickelte sich rasch zu einem weltweit renommierten Treffpunkt in der Informatikforschung. Heute beherbergt die Begegnungsstätte (siehe Fig. 1.1) jährlich mehr als 3 000 internationale Gäste.

Seit 2005 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 95 Forschungsinstituten, Bibliotheken und Museen.¹ Schloss Dagstuhl wird seit 2006 durch eine Bund-Länder-Förderung finanziert.

Zu dem anfänglich alleinigen Schwerpunkt des Seminarprogramms haben sich in den vergangenen Jahren zwei weitere Geschäftsfelder hinzugesellt: Zum einen der Betrieb der offenen Bibliographiedatenbank dblp, zum anderen die Angebote als Open-Access-Verleger für die Informatikforschenden.

■ Seminar- und Workshop-Programm

Schwerpunkt des wissenschaftlichen Programms von Schloss Dagstuhl sind die Dagstuhl-Seminare und die Dagstuhl-Perspektiven-Workshops: Etwa 30 bzw. 45 internationale Forscher treffen sich eine halbe bis ganze Woche auf Schloss Dagstuhl, um im Rahmen eines Dagstuhl-Seminars

Schloss Dagstuhl – Leibniz Zentrum für Informatik (Leibniz Center for Informatics) pursues its mission of furthering world class research in computer science by facilitating communication and interaction between researchers. The objective of Schloss Dagstuhl is

- to promote basic and application-oriented research in the field of informatics,
- to support advanced, scientific vocational training and to further education in the field of informatics,
- to promote the transfer of knowledge between research into informatics and application of informatics,
- and to operate an international forum and research institute for informatics.

Including and thus promoting young talents is seen as an important part of our efforts, so is promoting the exchange of knowledge and findings between academia and industry.

■ History of the Center

The idea behind a seminar center for informatics came about during the late 1980s, when research in computer science grew rapidly worldwide as an offshoot of mathematics and engineering. At that time the German *Gesellschaft für Informatik* (German Informatics Society) became aware of the growing number of computer scientists at the world-famous *Mathematics Research Institute* in Oberwolfach, Germany, and recognized the need for a meeting venue specific to the informatics community. Schloss Dagstuhl was founded in 1990 and quickly became established as one of the world's premier centers for informatics research. Today, Schloss Dagstuhl (see Fig. 1.1) hosts over 3,000 research guests each year from countries across the globe.

Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association, a non-profit research consortium composed of 95 research institutes, libraries and museums throughout Germany.¹ Since 2006 the center is jointly funded by the German federal and state governments.

Since the very first days of Schloss Dagstuhl, the seminar and workshop meeting program has always been the focus of its programmatic work. In recent years, Schloss Dagstuhl has expanded its operation and also has significant efforts underway in operating the dblp computer science bibliography and in open access publishing for the computer science community.

■ Seminar and Workshop Program

The Dagstuhl Seminars and Dagstuhl Perspectives Workshops form the focus of the center's work. Whereas ca. 30 or 45 established and young researchers gather at the Dagstuhl Seminars to report on and discuss their current work, smaller groups of ca. 30 of the international elite of

¹ Stand Februar 2019.
As of Februar 2019.



Fig. 1.1
Aerial photography of Schloss Dagstuhl.

intensiv über ihre aktuelle Forschung zu diskutieren. Darüber hinaus trifft sich in Dagstuhl-Perspektiven Workshops eine kleinere Gruppe von ca. 30 Spitzenforschern, um über den aktuellen Stand und die zukünftigen Schwerpunkte eines ganzen Forschungsfeldes zu beraten.

Die Seminare und Perspektiven-Workshops werden jeweils von bis zu vier ausgewiesenen Wissenschaftlern im entsprechenden Gebiet beantragt. Anträge werden durch das wissenschaftliche Direktorium (siehe Kapitel 11.3) begutachtet. Stellenwert bei der Begutachtung haben neben dem eigentlichen Inhalt des Antrags auch die vorgeschlagene Gästeliste sowie die Antragsteller. Nach Annahme finden die entsprechenden Veranstaltungen dann durchschnittlich zwischen 6 und 18 Monaten später statt. Eine Teilnahme ist nur mit einer persönlichen Einladung durch das Zentrum möglich.

Das Seminarzentrum ist im und rund um das 1760 erbaute Schloss Dagstuhl beheimatet und 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 mit anderen Gästen in einer wohnlichen Atmosphäre fördert. Gemütliche Sitzecken, ansprechende Essräume, eine herausragenden Informatik-Fachbibliothek, sowie eine Vielzahl von zusätzlichen Arbeits- und Freizeiträumen bieten vielfältige Möglichkeiten, damit sich die Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

Nähere Informationen über Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops finden sich in Kapitel 2.

a field gather at the Dagstuhl Perspectives Workshops for the purpose of reflecting on the current status of research and potential development perspectives.

These seminars are characterized by the fact that they are subject to an exacting quality assurance process. A small group of up to four scientists of international standing submit a proposal for a seminar on a specific research topic. The proposal is reviewed by the center's Scientific Directorate (see Section 11.3) with regard to its content, the proposed guest list and those submitting the proposal. The seminars and workshops are held 6 to 18 months later in the seclusion of the center's facilities at Dagstuhl Castle. Participation in a seminar is possible only by way of personal invitation by the center.

Located in a 1760 build manor house 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.

More information on the Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Chapter 2.

■ Bibliographiedatenbank dblp

Bereits seit 2011 betreibt Schloss Dagstuhl in enger Zusammenarbeit mit der Universität Trier die Bibliographiedatenbank dblp. Seit November 2018 ist Schloss Dagstuhl in vollem Umfang alleine für den Betrieb der Datenbank verantwortlich.

Mit mittlerweile mehr als 4,4 Millionen Publikationseinträgen ist dblp die weltweit größte offene Sammlung bibliographischer Daten in der Informatik. Der dblp-Dienst ist darauf ausgerichtet, Forscher bei ihrer täglichen Arbeit zu unterstützen, etwa bei der Literaturrecherche oder beim Bezug von elektronisch verfügbaren Volltexten. Dabei gilt dblp in der Informatik insbesondere als die Referenzdatenbank für qualitätsgesicherte, normierte Bibliographiedaten. Aber auch Forschungsförderer und Entscheidungsträger unterstützt dblp, etwa durch das Pflegen und öffentlich Verfügbarmachen von personalisierten Publikationsnachweisen. Durch den Betrieb von dblp leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Details über dblp finden sich in Kapitel 3.

■ Dagstuhl Publishing

Die Förderung der Kommunikation zwischen den Wissenschaftlern in der Informatik gehört zu der zentralen Aufgabe von Schloss Dagstuhl. Wissenschaftliche Veröffentlichungen sind Teil der Forschungskultur, um qualitätsgesicherte Forschungsergebnisse zu diskutieren und zu kommunizieren. Mit seinen Open-Access-Verlagsangeboten unterstützt Schloss Dagstuhl die Forschungsgemeinde dabei, freien Zugang zu den wichtigsten und neuesten Forschungsergebnissen zu erlangen.

Neben Veröffentlichungen, die in engem Bezug zum wissenschaftlichen Programm stehen, verlegt Schloss Dagstuhl auch Konferenzbände und Zeitschriften. Herausragende Reihe ist dabei LIPIcs, in der die Publikationen erstklassiger Konferenzen erscheinen. Alle Angebote der Verlagsabteilung werden durch international besetzte Editorial Boards qualitätsgesichert.

Kapitel 4 stellt Dagstuhls Verlagswesen ausführlicher dar.

■ dblp computer science bibliography

Since 2011, Schloss Dagstuhl has been operating the dblp computer science bibliography in close cooperation with the University of Trier. In November 2018, Schloss Dagstuhl alone assumed full responsibility for the operation of the database.

Listing more than 4.4 million articles, dblp is the world's most comprehensive open data collection of computer science research articles. The goal of dblp is to support computer scientists in their daily work, for example when reviewing the literature of a given author or subject area, or when searching for online full-text versions of research articles. The dblp database is often considered to be the reference database for quality-assured and normalized bibliographic metadata in computer science. Additionally, dblp supports funding agencies and decision makers by providing and curating personalized author profiles. By operating dblp, Schloss Dagstuhl furthers its mission of promoting the identification, dissemination and implementation of new computer science developments at an internationally recognized level.

More information about the dblp computer science bibliography can be found in Chapter 3.

■ Dagstuhl Publishing

Enabling communication between researchers in computer science is part of Dagstuhl's central mission. Scholarly publications belong to the culture of discussing and communicating quality-controlled research results on a global level. Dagstuhl's open-access publishing services hence support the need of the research community to have access to the most important and most recent research results.

In addition to the open documentation of proceedings of its seminar and workshop program, Schloss Dagstuhl also publishes proceedings for computer science conferences and journals. The flagship product of Dagstuhl Publishing is the LIPIcs series, which publishes proceedings of outstanding computer science conferences. The scientific quality of all products is supervised by international editorial boards.

More information on Dagstuhl Publishing can be found in Chapter 4.

Neuigkeiten in 2018

1.2

■ Neue Außenstelle von Schloss Dagstuhl in Trier gegründet

Am 16. November 2018 hat die Gemeinsame Wissenschaftskonferenz von Bund und Ländern die Gründung einer offiziellen Außenstelle von Schloss Dagstuhl auf dem Campus der Universität Trier beschlossen. Die rheinland-pfälzische Ministerpräsidentin Malu Dreyer und Wissenschaftsminister Konrad Wolf bezeichneten die Entscheidung zur Einrichtung als großen Gewinn für den

News from 2018

■ New branch office of Schloss Dagstuhl in Trier established

On November 16, 2018, the Joint Science Conference of the Federal Government and the Federal States of Germany decided to establish an official Schloss Dagstuhl branch office on the campus of the University of Trier. The Prime Minister of Rhineland-Palatinate, Malu Dreyer, and the Minister of Science, Konrad Wolf, described the decision as a great benefit for Trier as a science location:

Wissenschaftsstandort Trier: „Der neue Standort des Leibniz-Zentrums auf dem Campus der Universität Trier wird die Forschungsaktivitäten im Bereich Informatik und Digitalisierung fördern. Das ist eine einmalige Chance.“

Neben dem Tagungszentrum in Wadern und der Geschäftsstelle in Saarbrücken (beides im Saarland) ist die neue Außenstelle in Trier der dritte offizielle Standort von Schloss Dagstuhl und der erste Standort im Zuwendungsgelände Rheinland-Pfalz. Die neue Außenstelle beherbergt dabei vornehmlich die Mitarbeiter des dblp-Teams und erlaubt so auch weiterhin eine enge Zusammenarbeit mit der Universität und dem Fach Informatikwissenschaften am Standort Trier bei Betrieb und Weiterentwicklung der dblp computer science bibliography.

■ Das Team

Am Ende des Jahres 2018 beschäftigte Schloss Dagstuhl insgesamt 40 Vollzeitäquivalente bzw. 54 Mitarbeiter. Schloss Dagstuhl beschäftigt seit August 2018 eine neue Auszubildende, Frau Jule Schneider, im Bereich Hauswirtschaft. Seit Januar 2018 ist Frau Shida Kunz Mitglied des Wissenschaftlichen Stabs. Die langjährige und sehr geschätzte Mitarbeiterin Frau Margit Brücker ging im März 2018 in den wohlverdienten Ruhestand. Im Dezember 2018 begann Schloss Dagstuhl, das Team von dblp im Rahmen der geplanten Erweiterung und Verstärkung von dblp als strategische Aufgabe von Schloss Dagstuhl aufzustocken.

Nahezu alle Mitarbeiter von Schloss Dagstuhl wurden 2018 über den Kernhaushalt des Zentrums bezahlt. Eine Ausnahme bildet ein Mitarbeiter im Rahmen eines Projektes, das aus Fördermitteln des Leibniz-Wettbewerbs finanziert wurde. Zudem unterstützte das Heidelberger Institut für theoretische Studien (HITS) mit einer Spende von 24 000 € letztmalig die Open-Access-Publikationsaktivitäten von Schloss Dagstuhl.

■ Seminare und Workshops

In 2018 wurden 136 Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops gestellt. Dies ist mehr als im Vorjahr, und liegt als neuer Höchstwert über der langfristigen Tendenz zu etwa 100 Anträgen pro Jahr. Durch die entspanntere Antragslage des Vorjahrs sank die Anzahl der Seminare und Workshops in 2018 auf 65. Damit fanden erstmals seit 2012 wieder weniger als 70 Seminare und Workshops statt.

Von den mehr als 3 200 Gästen, die sich in Dagstuhl trafen, nahmen etwa 2 300 an Seminaren teil. Etwa 44% aller Seminarteilnehmer war zum ersten Mal in Dagstuhl, und mehr als ein Drittel der Teilnehmer an unserer Gastumfrage ordnete sich selbst als Junior-Wissenschaftler ein. Mehr als drei Viertel aller Seminarteilnehmer waren außerhalb von Deutschland beschäftigt.

Etwa 65% aller in 2018 stattgefundenen Seminare hatten mindestens eine Frau im drei- bis fünfköpfigen Organisationssteam, und rund 22% aller Organisatoren waren

“Der neue Standort des Leibniz-Zentrums auf dem Campus der Universität Trier wird die Forschungsaktivitäten im Bereich Informatik und Digitalisierung fördern. Das ist eine einmalige Chance.”²

In addition to the seminar center in Wadern and the Saarbrücken office (both in Saarland), the new Trier office is the third official location of Schloss Dagstuhl and the first location in the co-funding state of Rhineland-Palatinate. The new office mainly accommodates the dblp team. Thus, it allows to continue the close cooperation with the university and the computer science department in Trier for the benefit of operation and development of the dblp computer science bibliography.

■ The Team

By the end of 2018, Schloss Dagstuhl had a total of 54 staff members corresponding to 40 full-time positions. Schloss Dagstuhl employs one new trainee, Jule Schneider, in housekeeping since August 2018. Since January 2018, Shida Kunz is a member of the Scientific Staff. Schloss Dagstuhl's longstanding and highly appreciated employee Margit Brücker took her well-earned retirement in March 2018. In December 2018, Schloss Dagstuhl started to expand the dblp team as part of the planned expansion and perpetuation of dblp as a strategic task of Schloss Dagstuhl.

Nearly all staff at Schloss Dagstuhl were funded from the center's core budget in 2018. An exception is one staff member working on a project that is financed by a project grant of the Leibniz Competition. Additionally, Dagstuhl's open-access publishing service was supported by a generous final donation in the amount of 24,000 € from the Heidelberg Institute for Theoretical Studies (HITS).

■ Seminars and Workshops

In 2018, 136 applications for Dagstuhl Seminars and Dagstuhl Perspectives Workshops were submitted. This is more than in the previous year, and is a new record above the long-term trend of about 100 applications per year. Due to the more calm proposal submission period in the previous year, the number of seminars and workshops declined to 65 in 2018. For the first time since 2012, fewer than 70 seminars and workshops took place.

Of the more than 3,200 guests who met in Dagstuhl, about 2,300 took part in seminars. About 44% of all seminar participants were in Dagstuhl for the first time, and more than a third of the participants in our guest survey classified themselves as junior scientists. More than three quarters of all seminar participants were employed outside Germany.

About 65% of all seminars held in 2018 had at least one woman in the team of three to five organizers, and about 22% of all organizers were women. Women accounted for

² engl.: “The new office of the Leibniz Center on the campus of the University of Trier will promote research activities in the field of computer science and digitization. This is a unique opportunity.”

Frauen. Der Frauenanteil unter allen Seminarteilnehmern betrug 19,5 %.

Mehr Details und Zahlen zum Seminarprogramm finden sich in Kapitel 2.

■ Bibliographiedatenbank dblp wächst auch nach 25 Jahren weiter stark

2018 feierte die dblp computer science bibliography bereits ihr 25-jähriges Bestehen. Aus diesem Anlass fand am 23. November 2018 an der Universität Trier ein Festkolloquium unter dem Motto „25 Jahre dblp – 2²² Einträge“ statt. Als Festredner gab Prof. Dr. Gerhard Weikum (Max-Planck-Institut für Informatik, Saarbrücken) zum Thema „Machine Knowledge: Encyclopedic, Scholarly, Commonsense“ einen Einblick in die Chancen und Herausforderungen von semantischen Wissensdatenbanken. Zudem zeigte Prof. Dr. Claude Kirchner (INRIA Frankreich) in seinem Vortrag mit dem Titel „Strategic Tools for Scientific Sovereignty“ Möglichkeiten auf, wie Forscherinnen und Forscher in einer sich drastisch wandelnden Wissenschaftslandschaft die Souveränität über ihre Arbeiten und den Nutzen für die Allgemeinheit aufrecht erhalten können.

Pünktlich zum Jubiläum erfolgte zudem auch die Stafelübergabe des Betriebes der Datenbank von der Universität Trier an Schloss Dagstuhl. Eine entsprechende Vereinbarung wurde von den Gremien beider Institutionen verabschiedet und wurde im Rahmen des Festkolloquiums unterzeichnet. Fig. 1.2 zeigt einige Eindrücke des Festkolloquiums und der Vertragsunterzeichnung.

In Jahr 2018 konnte dblp mit mehr als 420 000 Publikationen einen neuen Rekordwert bei der Anzahl an Neuaufnahmen verzeichnen. Diese Aufnahmequote entspricht mehr als 1 680 neuen Publikationen pro Arbeitstag. Die Nutzung des dblp-Dienstes blieb dabei auf hohem Niveau und konnte sogar gegenüber dem Vorjahr gesteigert werden. Jeden Monat verzeichneten die dblp-Webseiten etwa 25 Millionen Seitenzugriffe von etwa 675 000 verschiedenen Nutzern aus aller Welt. Dies entspricht ca. 9,7 Seitenzugriffen pro Sekunde; zu Spitzenzeiten erfolgen über 100 Zugriffe gleichzeitig. Im Durchschnitt beginnt etwa alle zwei Sekunden ein neuer Nutzer, mit dblp zu arbeiten.

Mehr Informationen zu dblp finden sich in Kapitel 3.

■ Dagstuhl Publishing

Wie in den Vorjahren haben die Open-Access-Publikationsaktivitäten auch in 2018 starken Zuspruch bekommen. So wurden allein in der Konferenzbandreihe LIPIcs erstmals über 1 300 Publikationen innerhalb eines Jahres veröffentlicht. Zudem gab es auch in 2018 wieder viele Bewerbungen von wissenschaftlichen Konferenzen zur Veröffentlichung des Konferenzbandes in der Serie LIPIcs.

Mehr Informationen zu den Open-Access-Aktivitäten von Schloss Dagstuhl finden sich in Kapitel 4.

■ Dagstuhler Gespräche

Auch in 2018 wurde die erfolgreiche Vortragsreihe „Dagstuhler Gespräche“ als gemeinsame Veranstaltung

19.5 % of all seminar participants.

See Chapter 2 for more details and statistics regarding the seminar program.

■ 25 years of the dblp computer science bibliography – still growing strongly

In 2018 the dblp computer science bibliography celebrated its silver anniversary. To mark this occasion, a festive colloquium was held at the University of Trier on November 23, 2018. Under the motto “25 years of dblp – 2²² entries” the database also celebrated the milestone of more than 4 million indexed computer science publications. In the first keynote of the evening, “Strategic Tools for Scientific Sovereignty”, Prof. Dr. Claude Kirchner (INRIA France) presented strategic tools that allow researchers to maintain sovereignty over their work and to guarantee the benefit for the common good in a drastically changing scientific landscape. As second keynote speaker, Prof. Dr. Gerhard Weikum (Max Planck Institute for Informatics, Saarbrücken) gave an insight into the opportunities and challenges of machine knowledge with his keynote “Machine Knowledge: Encyclopedic, Scholarly, Commonsense”.

Just in time for the anniversary, the transfer of the operation of the database from the University of Trier to Schloss Dagstuhl took place. A corresponding agreement was passed by the committees of both institutions and was signed at the festive colloquium. Fig. 1.2 shows some impressions of the festive colloquium and the signing of the agreement.

During 2018, the dblp database grew by more than 420,000 new publication records. This is the largest figure ever achieved in the history of dblp and corresponds to more than 1,680 new records for each working day of the year. Up to 25 million web pages are visited each month by about 675,000 researchers and computer science enthusiasts all over the world. On average, about 9.7 web pages are requested from the dblp web servers in every second; at peak times, as many as 100 requests are made concurrently. On average, every two seconds, a new user session is started.

More information about dblp can be found in Chapter 3.

■ Dagstuhl Publishing

As in the previous years, Schloss Dagstuhl’s open-access publishing services experienced an on-going strong increase in demand from the community in 2018. For the first time, more than 1,300 articles have been published within one year in the conference proceedings series LIPIcs. Furthermore, LIPIcs again received and accepted proposals from several major scientific conferences.

More information about the Open Access activities of Schloss Dagstuhl can be found in Chapter 4.

■ “Dagstuhler Gespräche”

The successful series of “Dagstuhler Gespräche” (Dagstuhl Talks) was continued as a joint event of Schloss



Fig. 1.2
Celebrating 25 years of the dblp computer science bibliography.

von Schloss Dagstuhl und der Stadt Wadern fortgeführt. Ziel dieser Gespräche ist es, der interessierten Öffentlichkeit die breite Vielfalt der Informatik und deren praktische Anwendungen im Alltag oder in wirtschaftlichen Prozessen nahezubringen und in einen gemeinsamen Dialog einzusteigen.

Für den 18. Mai 2018 konnte der ehemalige Präsident der Gesellschaft für Informatik e.V. und frühere Leiter des Fraunhofer Instituts für Rechnerarchitektur und Softwaretechnik (FIRST), Prof. Dr. Stefan Jähnichen, als Vortragender gewonnen werden. Unter dem Titel „Digitalisierung – (k)ein Ende in Sicht?“ erklärte er aus seinem Erfahrungsschatz die Digitalisierung und hob spezielle Problemstellungen zu Sicherheit und Akzeptanz ihres Einsatzes hervor, die – wenn nicht zufriedenstellend gelöst – dem Trend der zunehmenden Digitalisierung entgegenstehen können. Speziell erörterte er, wie man sich in der digitalen Welt ausweisen kann, wie also digitale Identitäten gestaltet werden können, um einen vertrauensvollen Umgang mit Daten zu ermöglichen.

Der Vortrag fand regen Anklang und löste angeregte Diskussionen aus. Die Reihe wird im kommenden Jahr gewiss fortgesetzt werden.

■ Lehrerfortbildung

In Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisierte Schloss Dagstuhl 2018 zum 28. Mal eine Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet.

■ Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch im Jahr 2018 gab es wieder eine Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum³ (HLF). Diese Veranstaltung bringt herausragende Mathematiker und Informatiker, nämlich Gewinner des ACM Turing Award, des Abelpreises, der Fields-Medaille, und des Nevanlinna-Preises, mit außergewöhnlich begabten jungen Wissenschaftlern aus aller Welt zusammen. Drei ausgewählte Teilnehmer des HLF 2018 erhielten in der Woche nach der sechsten Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Automating Data Science“ (18401).

Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2019 zugestimmt.

■ Spender und Förderer

Schloss Dagstuhl ist den wissenschaftlichen Gästen, Institutionen und Firmen dankbar, die durch großzügige

Dagstuhl and the city of Wadern. These talks aim at giving the interested public an understanding of the broad range of computer science and its practical applications in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

The talk on May 18, 2018 was given by the former President of the Gesellschaft für Informatik e.V. (German Informatics Society) and former head of the Fraunhofer FIRST institute, Prof. Dr. Stefan Jähnichen. Under the title “Digitalisierung – (k)ein Ende in Sicht?” (Digitalization – (no) end in sight?), he explained the digital transformation from his rich experience and highlighted specific problems concerning security and acceptance, which could impede the growing trend of digitalization if no satisfying solutions are found. He discussed how one can prove one’s identity in the digital world and how digital identities can be designed such as to foster trust in handling of data.

The talk was well received, and the discussion was lively. The Dagstuhler Gespräche will certainly see a continuation in the next year.

■ Teacher Training Program

In 2018, Schloss Dagstuhl hosted its teacher training course for the 28th time. This workshop is specifically designed for computer science and mathematics teachers teaching grades 11 and 12 in Saarland and Rhineland-Palatinate. It is organized in collaboration with the Landesinstitut für Pädagogik und Medien Saarland LPM (Saarland State Institute for Education and Media) and the Pädagogisches Landesinstitut Rheinland-Pfalz PL (Rhineland-Palatinate State Institute for Education).

■ Joint Outreach with the Heidelberg Laureate Forum

2018 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum³ (HLF). The HLF brings winners of the ACM Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize together with exceptionally talented young scientists from all over the world. Three participants were selected and invited to participate in the Dagstuhl Seminar “Automating Data Science” (18401), taking place during the week after the sixth edition of the forum.

Satisfied with the outstanding success of the initiative, both partners agreed to continue the cooperation in 2019.

■ Sponsors and Donors

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generous donations for the

³ <http://www.heidelberg-laureate-forum.org>

Spenden das Zentrum unterstützen.

2018 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 661 Bände als Spende, darunter 611 Monographien des Springer-Verlags im Wert von 47 597 €.

Dagstuhl Publishing wurde in 2018 mit 24 000 € durch das Heidelberger Institut für Theoretische Studien (HITS) unterstützt.

■ NSF Förderung von Nachwuchswissenschaftlern

Seit 2013 stehen Mittel zur Unterstützung von Nachwuchswissenschaftlern aus den USA bei der Teilnahme an Dagstuhl Seminaren zur Verfügung. Diese Fördermöglichkeit wird durch die National Science Foundation (NSF) finanziert⁴.

Im Berichtsjahr konnte durch die Förderung 41 Forschern aus den USA eine Teilnahme an insgesamt 32 Seminaren ermöglicht werden. Insgesamt wurden dafür 68 860,44 \$ Fördermittel ausgegeben. Die Förderung wurde um ein weiteres Jahr bis September 2019 verlängert.

■ Baumaßnahmen und Renovierung

In 2018 wurde der sogenannte Weinkeller umfangreich renoviert (siehe Fig. 1.3). Die Verbesserung der Raumakustik hatte bei der Renovierung höchste Priorität: große Leuchten mit Stoffbezug, ein schalldämmender Bodenbelag, Polstermöbel und Wände mit Stoffbezug ermöglichen nun eine deutlich verbesserte Kommunikation in den Abendstunden.

Ein ebenfalls schallschutz-, aber auch wärmetechnisches Problem stellte sich im Restaurantbereich. In den Betriebsferien hat Schloss Dagstuhl daraufhin die Decke komplett isolieren und eine Schallschutzdecke einbauen lassen. Es ist seitdem im Sommer deutlich kühler, im Winter deutlich wärmer und die Unterhaltungen beim Mittagessen verlaufen aufgrund der Lärmreduzierung sichtbar entspannter.

Zudem wurde Ende 2018 mit dem Umbau der Räumlichkeiten unseres Personals begonnen. Ziel war hier, sowohl die Richtlinien des Arbeitsschutzes zu erfüllen als auch eine separate Nutzung von Umkleieräumlichkeiten und Personaltoiletten für beide Geschlechter zu ermöglichen. Im Zuge dieser Maßnahme wurde außerdem der Pausenraum isoliert und modernisiert.

support of its center.

The center's research library received a large number of book donations from several publishing houses. The number of donated volumes totaled 611, including monographs at the total value of 47,597 € donated by Springer Science+Business Media publishing house.

The Heidelberg Institute for Theoretical Studies (HITS) supported Dagstuhl Publishing with 24,000 €.

■ NSF Grant for Junior Researchers

Since 2013, Dagstuhl helps junior researchers based in the USA to participate in Dagstuhl seminars with funds provided through the grant opportunity⁴ financed by the National Science Foundation (NSF).

In 2018, 41 US-based scientists were supported with a total amount of 68,860.44 \$ and hence were able to participate in 32 Dagstuhl Seminars overall. The grant was extended for one more year until September 2019.

■ Construction Work and Renovation

In 2018, the so-called wine cellar underwent major refurbishment (see Fig. 1.3). The highest priority was the improvement of the acoustics: large lamps with textile shades, an acoustically insulating flooring, upholstery, and walls with textile covers now enable a significantly improved communication during the evening hours.

The acoustics also used to be problematic in the restaurant area, which was compounded with an insufficient thermal insulation. Thus, during the summer break, the ceiling was completely insulated and a noise-insulating ceiling was installed. This proved to lower inside temperatures in summer, increase them in winter, and chatting over lunch became clearly easier due to the reduction in noise.

Additionally, the refurbishment of the staff rooms was started towards the end of 2018. This is meant to improve occupational safety and facilitate a gender separated use of changing rooms and staff toilets. In the course of this works, the staff's break room was thermally insulated and modernized, as well.

⁴ Grant CNS-1257011: „Schloss Dagstuhl –NSF Support Grant for Junior Researchers“.



Fig. 1.3

Refurbished wine cellar.

■ Ausstattung

In 2018 wurden fast alle Gästezimmer sowie die Bibliothek mit neuen und bequemen Freischwinger-Stühlen ausgestattet.

Im Raum Trier wurden die bereits in 2017 angeschafften höhenverstellbaren Tische um eine leichte, rollbare Konferenzbestuhlung ergänzt. Insgesamt soll diese Maßnahme die extrem flexible Nutzung dieses Raumes unterstützen.

Im Flur des „Neubaus“ wurden zur Schallisolierung Deckensegel und ein breiter Teppich angebracht.

Die Schreibtischplatten sämtlicher Gästezimmer im „Anbau“ wurden deutlich gekürzt, so dass die Räume jetzt großzügiger erscheinen, ohne an Nutzbarkeit eingebüßt zu haben. Außerdem wurden für alle Gästezimmer neue Matratzen angeschafft.

■ Facilities

In 2018, almost all guest rooms and the library were furnished with new comfortable cantilever chairs.

In room Trier, the height adjustable desks acquired in 2017 were complemented with light, rollable chairs. Together, this equipment is meant to support the very versatile use of this room.

In the hallway of the newer building, ceiling sails and a broad carpet were installed as acoustic insulation.

The desktops of all guests rooms in the new part of the castle were shortened considerably, so the rooms now have more free space without losing usability. Moreover, new mattresses have been purchased for all guest rooms.

2 **Seminare und Workshops** *Seminars and Workshops*

Dagstuhl-Seminare

2.1

Dagstuhl Seminars

Die Dagstuhl-Seminare haben als wesentliches Instrument der Forschungsförderung Priorität bei der Gestaltung des Jahresprogramms. Hauptziel der Seminare ist die Unterstützung der Kommunikation und des Dialogs zwischen Wissenschaftlern, die an den Forschungsfronten 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 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 25 Preisträger des Turing-Awards, der höchsten Auszeichnung, die im Bereich der Informatik auf internationaler Ebene 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 und sehr breite Übersicht über die aktuellen Forschungsgebiete der Informatik.

Jedes Dagstuhl Seminar wird gebeten, einen kurze Dokumentation zu erstellen, die eine Zusammenfassung des Seminarverlaufs, eine Kurzübersicht über die gehaltenen Vorträge und eine Zusammenfassung grundsätzlicher Ergebnisse enthält. Diese Berichte, die in der Zeitschrift *Dagstuhl Reports* veröffentlicht werden, gewährleisten eine hohe Sichtbarkeit und eine zeitnahe Kommunikation der Ergebnisse. *Dagstuhl Reports* wird jährlich in einem Band mit 12 Ausgaben veröffentlicht. Jede Ausgabe dokumentiert jeweils die Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops eines Monats. Die *Dagstuhl Reports* sind über die Dagstuhl-Website frei zugänglich.⁵

Kapitel 6 enthält Zusammenfassungen der Dagstuhl-Seminare und Perspektiven-Workshops. Im Kapitel 14 sind alle Veranstaltungen, die 2018 stattfanden, aufgelistet. Auf der Dagstuhl-Website ist das Programm der kommenden 24 Monate verfügbar.

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

Participation in these events – which generally last one week – is possible only by way of personal invitation from Schloss Dagstuhl. The center assumes part of the associated costs in order to enable the world's most qualified scientists, including young researchers and doctoral students, to participate. Among Dagstuhl's guests have been 25 winners of the ACM Turing Award, 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 themselves. Many of the topics addressed in-depth at Dagstuhl have subsequently developed 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 and broad overview of the areas currently under discussion in the 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's program, talks, and results 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 openly accessible and can be downloaded at the Dagstuhl website.⁵

Chapter 6 contains a collection of the summaries of the 2018 Seminars and Perspectives Workshops. Chapter 14 provides a comprehensive list of all events that took place during the year under review, and a seminar program covering the coming 24 months is available on the Dagstuhl website.

⁵ <https://www.dagstuhl.de/dagrep/>

Dagstuhl-Perspektiven-Workshops

2.2

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

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

Die Dagstuhl-Perspektiven-Workshops, die 2018 statt fanden, sind in Fig. 2.1 aufgelistet.

Die Ergebnisse der intensiven Diskussionen werden in einem Manifest zusammengefasst, welches die offenen Probleme und die möglichen Forschungsperspektiven für die nächsten 5–10 Jahre aufzeigt. Dagstuhl koordiniert die gezielte Weitergabe dieses Manifests, um forschungsspezifische Impulse an deutsche und europäische Institutionen der Forschungsförderung zu geben (EU, BMBF, DFG, etc.). Kurzfassungen der Manifeste werden regelmäßig im Forum des *Informatik Spektrum* (Springer-Verlag) vorgestellt. Die vollständigen Manifeste werden in unserer Fachzeitschrift *Dagstuhl Manifestos*⁶ veröffentlicht.

Eine Liste der vergangenen und kommenden Dagstuhl-Perspektiven-Workshops ist auf der Dagstuhl-Website verfügbar.⁷

Dagstuhl Perspectives Workshops

2

In addition to the traditional Dagstuhl Seminars, the center organizes Dagstuhl Perspectives Workshops. A Perspectives Workshop involves 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
- show research directions
- trigger innovation processes

Dagstuhl Perspectives Workshops held in 2018 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*⁶.

A list of past and upcoming Dagstuhl Perspectives Workshop can be found on our web site.⁷

Einreichung der Anträge und Begutachtungsverfahren

2.3

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch

Proposal Submission and Review Process

Schloss Dagstuhl maintains the high quality of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop

⁶ <https://www.dagstuhl.de/dagman>

⁷ <https://www.dagstuhl.de/pw-list>

10 Years of Web Science: Closing The Loop
<http://www.dagstuhl.de/18262>

Implementing FAIR Data Infrastructures
<http://www.dagstuhl.de/18472>

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2018.

Auswahl der Anträge gewährleistet, die aus Sicht von Schloss Dagstuhl das größte Potential haben, abseits etablierter Konferenzen neue und wichtige Forschungsprobleme mit Wissenschaftlern aus oft unterschiedlichen Gebieten zu identifizieren und zeitgleich mögliche Methoden und Lösungsansätze zu diskutieren.

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 werden dann vom Wissenschaftlichen Direktorium (siehe Kapitel 11.3) begutachtet und abschließend bei zweitägigen Sitzungen auf Schloss Dagstuhl intensiv diskutiert und über sie entschieden.

Es wird sicher gestellt, dass jedes Dagstuhl-Seminar durch ein starkes Organisatorenteam betreut 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. Zudem wird auf eine ausgeglichene Repräsentation wissenschaftlicher Gemeinden, geographischer Regionen und besonders auf das Miteinbeziehen junger und weiblicher Wissenschaftler geachtet.

Die Informatikforscher zeigten 2018 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 136 Anträgen in den Antragsrunden im Januar und Juni 2018. Der hohen Qualität der Anträge entsprechend, wurden etwa 60 % der eingereichten Anträge genehmigt. In den vergangenen 7 Jahren variierte die Rate der angenommenen Anträge zwischen 60 % und 76 % (siehe Fig. 2.2).

Unter den 81 in 2018 neu genehmigten Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops gab es wie in den vergangenen Jahren wieder verschiedene Konstellationen bzgl. Dauer und Größe (vgl. Fig. 2.3). Von diesen konnten 9 Seminare bereits 2018 ausgerichtet werden, der Großteil wurde jedoch für das Seminar-Programm in 2019 eingeplant (hier und im Folgenden wird, sofern nicht anders angegeben, das Wort "Seminar" sowohl für Dagstuhl-Seminare als auch für Dagstuhl-Perspektiven-Workshops verwendet). Insgesamt 10 der 2018 genehmigten Seminare werden – aufgrund des Überhangs an großen Seminaren – in 2020 stattfinden.

series by identifying those proposals that promise a high potential to engage researchers – often from different disciplines – in scientific discussions on new and important research problems and their most promising solutions, outside of the existing conferences.

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 Dagstuhl's Scientific Directorate (see Section 11.3) and finally discussed and decided during a two-day meeting at Schloss Dagstuhl.

This process ensures that every Dagstuhl Seminar and Dagstuhl Perspectives Workshop 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 balance of research communities and geographical regions, and especially the inclusion of junior and female researchers, are also taken into account during the review process.

The international scientific community expressed a lively interest in organizing seminars and workshops at Schloss Dagstuhl in 2018, submitting 136 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2018 and June 2018 submission rounds. The quality of the proposals was excellent, resulting in a 60 % acceptance rate by Dagstuhl's Scientific Directorate. In the previous seven years, proposal acceptance rates have tended to range between 60 % and 76 % (see Fig. 2.2).

Among the 81 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2018 there is – as in the past years – a wide variation with regard to length and size (see Fig. 2.3). Most of these seminars are part of the 2019 seminar program, although it was possible to schedule 9 of them already in 2018 (here and in the following, the word "seminar" is meant to include both Dagstuhl Seminars and Dagstuhl Perspectives Workshops, if not specified otherwise). A total of 10 seminars approved in 2018 will be held in 2020, as there are a lot of large seminars.

Seminar-Programm 2018

2.4

The Seminar Program in 2018

In 43 von 48 Wochen, in denen das Tagungszentrum 2018 geöffnet war, fand mindestens ein Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop statt. In 22 Wochen waren es sogar zwei. In fünf Wochen war das Zentrum nur durch andere Veranstaltungen belegt.

Seit 2012 ist es aufgrund des damals fertiggestellten Gästehauses möglich, zwei Seminare parallel in einer Woche zu veranstalten. Dadurch ist, verglichen mit den

At least one Dagstuhl Seminar or Dagstuhl Perspectives Workshop was held in 43 of the 48 weeks the center was open in 2018. In 22 of those weeks, there were in fact two seminars in parallel. In the five remaining weeks, there were exclusively other events scheduled.

Since the guest house opened in 2012, it has been possible for the center to schedule two parallel seminars in any given week. Thus, there was an increase of




Fig. 2.2 Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2012–2018.




Fig. 2.3 Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2012–2018. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

Jahren zuvor, seit 2012 die Gesamtanzahl an Seminaren pro Jahr gestiegen. 2018 fanden insgesamt 65 Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops statt. In Fig. 2.4 ist die Entwicklung der vergangenen Jahre dargestellt.

seminars held since 2012 compared with the years before. Altogether, there were 65 Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2018. Fig. 2.4 shows the evolution in recent years.

Angaben zu Teilnehmern und Organisatoren

2.5

Participant and Organizer Data

Viele der internationalen Teilnehmer der Seminare waren schon öfter in Dagstuhl. Dennoch zieht das Zentrum jedes Jahr auch neue Gesichter an, was den ständigen Wandel in der Forschung widerspiegelt. So nahmen – wie in den Vorjahren – auch in 2018 knapp die Hälfte (46%, 1 018 von 2 230) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 17% der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten, weitere 9% an zweien. Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass etwa 41% der Antwortenden 2018 das erste Mal, 16% zum zweiten Mal und weitere 11% zum dritten Mal (siehe Fig. 2.5a) teilgenommen haben.

Ein beträchtlicher Anteil der Gäste besteht aus jungen Wissenschaftlern, die am Anfang ihrer Karriere stehen, und für die der Aufenthalt in Dagstuhl oftmals prägend ist für den weiteren Verlauf ihres Lebenswegs. Etwa 35% der Gäste der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops in 2018, die an unserer Umfrage

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 informatics research across the globe. As in the previous years, in 2018, a bit less than half (1,018 of 2,230, or 46%) of the researchers were first-time visitors to Dagstuhl. About an additional 17% of the participating researchers had already attended one previous seminar in the years before, and another 9% had already attended two. Slightly different numbers are obtained from our guest survey: About 41% of the responders were first-time visitors, an additional 16% state their second visit, and yet another 11% their third (see Figure 2.5a).

A substantial number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 35% of 2018 Dagstuhl Seminar and Dagstuhl Perspectives Workshop survey respondents self-classified as junior (see Fig. 2.5b). This proportion of junior to senior




Fig. 2.4

Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2012–2018.

Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.5b). Diese ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern ist im Laufe der Jahre relativ konstant geblieben, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlern und ihren erfahrenen Kollegen zeigt.

Mit 79 % war der Anteil von Seminarteilnehmern aus dem Ausland 2018 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2018 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details können Kapitel 13 entnommen werden.

In 2018 waren etwa 65 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 19,5 % höher als in den Jahren zuvor, was auch 2017 schon der Fall war (siehe Fig. 2.6b).

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 around 79 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2018. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2018. For a detailed breakdown please refer to Chapter 13.

In 2018, 65 % of all organizer teams in our scientific seminar program were mixed with respect to gender (see Fig. 2.6a). The percentage of female seminar participants was higher than in previous years at 19.5 %, continuing the trend from 2017 (see Fig. 2.6b).

Themen und Forschungsgebiete

2.6

Topics and Research Areas

Die thematischen Schwerpunkte der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden von den internationalen Antragstellern identifiziert und dem wissenschaftlichen Direktorium zur Durchführung vorgeschlagen. Hierdurch wird die internationale Forschungsgemeinde aktiv in die Programmgestaltung eingebunden – zugleich ist gewährleistet, dass aufgrund der Expertise der Antragsteller in ihren jeweiligen Forschungsgebieten immer brandaktuelle Themenschwerpunkte gesetzt werden.


Im Folgenden sind beispielhaft einige thematische Schwerpunkte und dazugehörige Seminare aufgeführt. Die Aufzählung der Themen und Seminare hat keinen Anspruch auf Vollständigkeit und ist lediglich ein Versuch, einen kurzen Einblick in das umfangreiche Seminar-Programm zu geben. Kapitel 6 bietet mit den Kurzzusammenfassungen der Seminare und Perspektiven-Workshops einen vollständigen Überblick über das wissenschaftliche Seminar-Programm des Jahres 2018.

Unter den Seminaren, die sich Themen aus dem Bereich der theoretischen Informatik gewidmet haben war ein

The topics of Dagstuhl Seminars and Dagstuhl Perspectives Workshops are identified by researchers from all over the world, who pass on this information to the Schloss Dagstuhl Scientific Directorate in their submitted proposals. The international research community is thus actively involved in shaping Dagstuhl’s scientific seminar program, and their expertise ensures that the most important cutting edge topics are emphasized.


The following overview gives some topical focal points and a few respective seminars from 2018. Neither the list of focal points nor the list of seminars is exhaustive. It merely attempts to offer a brief insight into the multifarious scientific seminar program of 2018. Chapter 6, with the summary of the Seminars and Perspectives Workshops, provides a full overview of the 2018 scientific seminar program.

Among the seminars which addressed topics from theoretical computer science, there was an emphasis on complexity theory, with topics like *Proof Complexity* (18051), and *Algebraic Methods in Computational Complexity* (18391). Formal Methods were also well represented with




(a) Distribution of the number of previous attendances of participants, according to survey data.

(b) Percentage of junior researchers, according to survey data.



(c) Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2012–2018.



(a) Female organizers and mixed-gender organizer teams.

(b) Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2012–2018.

Schwerpunkt die Komplexitätstheorie, etwa mit *Proof Complexity* (18051), und *Algebraic Methods in Computational Complexity* (18391). Auch formale Methoden waren gut vertreten, etwa mit *Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance* (18211), auch in anderen Disziplinen, von der Mathematik mit *Formalization of Mathematics in Type Theory* (18341) bis zur Biologie mit *Formal Methods for the Synthesis of Biomolecular Circuits* (18082).

Aber auch die Algorithmik mit *Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization* (18081) und *High-Performance Graph Algorithms* (18241) und die Programmierung mit *Genetic Improvement of Software* (18052) und *Evidence About Programmers for Programming Language Design* (18061) kamen nicht zu kurz, und fundamental andersartige Felder für ihre Anwendung wurden mit *Algorithmic Foundations of Pro-*

topics like *Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance* (18211). This extended to other disciplines from mathematics with *Formalization of Mathematics in Type Theory* (18341) to biology with *Formal Methods for the Synthesis of Biomolecular Circuits* (18082).

But algorithmics with topics like *Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization* (18081) and *High-Performance Graph Algorithms* (18241) and programming with topics like *Genetic Improvement of Software* (18052) and *Evidence About Programmers for Programming Language Design* (18061) were not forgotten. They were also applied to fundamentally different fields with topics like *Algorithmic Foundations of Programmable Matter* (18331) and *Quantum Programming Languages* (18381).

grammable Matter (18331) und *Quantum Programming Languages* (18381) erschlossen.

Ein weiterer Schwerpunkt war der Umgang mit Daten und ihre Visualisierung, von *Foundations of Data Visualization* (18041), *In Situ Visualization for Computational Science* (18271) und *Progressive Data Analysis and Visualization* (18411) über *Data Consistency in Distributed Systems: Algorithms, Programs, and Databases* (18091) und *Automating Data Science* (18401) bis zu Anwendungen in der Biologie mit *Visualization of Biological Data – Crossroads* (18161) und den Geisteswissenschaften mit *Network Visualization in the Humanities* (18482).

Der Umgang mit großen Datenmengen und großer Komplexität unter Zeitdruck war ebenfalls ein Thema, von der Gehirnanalyse bei *High Throughput Connectomics* (18481) bis zu DevOps bei *Automatic Quality Assurance and Release* (18122).

Aktuell stark debattierte Themen waren natürlich deutlich vertreten, zum Beispiel Machine Learning etwa mit *Extreme Classification* (18291) und *Machine Learning and Model Checking Join Forces* (18121), Sicherheit mit *Symmetric Cryptography* (18021), *Secure Compilation* (18201), *Secure Routing for the Internet* (18242) und *Web Application Security* (18321), Mobilität mit *Dynamic Traffic Models in Transportation Science* (18102) und *Inter-Vehicular Communication Towards Cooperative Driving* (18202) und Mensch Maschine Interaktion mit *On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces* (18212), *Ubiquitous Gaze Sensing and Interaction* (18252) und gar *Human-Computer Integration* (18322).

Nicht zuletzt wurden mit *Encouraging Reproducibility in Scientific Research of the Internet* (18412) auch die Grundlagen des Wissenschaftlichen Arbeitens thematisiert.

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2018 veranstalteten Seminare wichtige Themen adressiert hat, die von den beteiligten Wissenschaftler mit großem Engagement diskutiert wurden und so die weitere Entwicklung in den einzelnen Gebieten wieder ein gutes Stück weitergebracht hat.

Another emphasis was on working with data and visualizing it, from *Foundations of Data Visualization* (18041), *In Situ Visualization for Computational Science* (18271), and *Progressive Data Analysis and Visualization* (18411) via *Data Consistency in Distributed Systems: Algorithms, Programs, and Databases* (18091), and *Automating Data Science* (18401) to applications in biology with *Visualization of Biological Data – Crossroads* (18161) and the humanities with *Network Visualization in the Humanities* (18482).

Handling a lot of data and complexity on the fly was a topic as well, from analysing brains with *High Throughput Connectomics* (18481) to DevOps with *Automatic Quality Assurance and Release* (18122).

Topics that are heavily debated at the moment were of course represented considerably, for example machine learning with *Extreme Classification* (18291) and *Machine Learning and Model Checking Join Forces* (18121), security with *Symmetric Cryptography* (18021), *Secure Compilation* (18201), *Secure Routing for the Internet* (18242), and *Web Application Security* (18321), mobility with *Dynamic Traffic Models in Transportation Science* (18102) and *Inter-Vehicular Communication Towards Cooperative Driving* (18202), and human machine interaction with *On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces* (18212), *Ubiquitous Gaze Sensing and Interaction* (18252), and even *Human-Computer Integration* (18322).

Not least, *Encouraging Reproducibility in Scientific Research of the Internet* (18412) had a look at the foundations of scientific working.

This brief selection of seminars should not draw attention from the fact that each of the 2018 seminars addressed important topics which were discussed by the involved researchers with great commitment and hence pushed forward the development in the individual areas.

Weitere Veranstaltungstypen

2.7

Further Event Types

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. Sie werden in Kooperation mit der GI durchgeführt und von dieser sowie von Dagstuhl gefördert. Anträge auf GI-Dagstuhl Seminare werden vom Vorstand der GIBU (GI Beirat der Universitätsprofessoren) und vom Wissenschaftlichen Direktor von Schloss Dagstuhl begutachtet.
- Weiterbildungsveranstaltungen wie Sommerschulen und Lehrerfortbildungen.

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

- GI-Dagstuhl Seminars bring young scholars together to discuss and learn about a specific topic. They are run and sponsored by the German Informatics Society (GI) in association with Schloss Dagstuhl. Proposals for GI-Dagstuhl Seminars are reviewed by the managing board of the GIBU (GI advisory board of computer science professors) and the Scientific Director of Schloss Dagstuhl.
- continuing education courses including summer schools and vocational training for teachers and instructors.

- Forschungsgruppentreffen wie Klausurtagungen von Graduiertenkollegs, GI-Fachgruppen und anderen akademischen Arbeitsgruppen.
- Forschungsaufenthalte von Einzelpersonen, die sich für eine oder mehrere Wochen für intensive Studien nach Dagstuhl in Klausur zurückziehen.

- research group meetings including conferences of graduate research training groups, GI specialist groups, and other academic working groups.
- research stays of scientists who wish to use the center as a retreat for several weeks in order to devote themselves to their studies undisturbed.

Qualitätssicherung

2.8

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

Fig. 2.7 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2018 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1385 Fragebögen, welche die Meinung von etwa 60 % der 2320 Teilnehmer repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

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

Quality Assurance

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

Fig. 2.7 shows the satisfaction of responding participants in 2018 with regard to selected aspects of their stay. The results were compiled from 1,385 questionnaires, representing the responses of about 60 % of all 2,320 participants. These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

Since 2013, Schloss Dagstuhl has also been offering all organizers a more transparent invitation process by giving them direct access to the status of invitee replies via a dedicated webpage. The page is updated daily and has met with very positive feedback from the organizers.

Auslastung des Zentrums

2.9

Auch 2018 konnte Schloss Dagstuhl die hohe Auslastung weitgehend halten. Es gab 2018 insgesamt 12 673 Gasttage, wobei 10 452 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Letztere Zahl bedeutet einen leichten Rückgang verglichen mit 2017. Es gab insgesamt etwas weniger Gasttage als in 2017. Es fanden im Berichtsjahr 109 Veranstaltungen mit insgesamt 3 203 Gästen statt. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2018 ebenso unbelegt wie eine Woche zum Jahresanfang, zwei Wochen im Juli/August und eine Woche am Jahresende. Diese wurden zu Instandhaltungs- und Verwaltungsarbeiten benötigt.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2018 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-

Utilization of the Center

Schloss Dagstuhl was able to uphold the high capacity utilization again in 2018. There were 12,673 overnight stays in total, with 10,452 overnight stays in Dagstuhl Seminars and Dagstuhl Perspectives Workshops. The latter number was a bit lower than in 2017. There were fewer overnight stays in total in 2018 compared to stays in 2017. The center hosted a total of 109 events with 3,203 guests in 2018. See Chapter 13 for further details.

Weekends were kept free in 2018, as well as a week at the beginning of the year, two weeks in July/August, and a week at the end of the year, this time being required for maintenance work to building facilities and administrative work.

A comprehensive listing of all events at Schloss Dagstuhl in 2018, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and




Fig. 2.7 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2018. According to survey results.

Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite ist ein Kalender⁸ verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

host-only events such as meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar⁸ of upcoming events and further information and materials on all events past, present and future.

⁸ https://www.dagstuhl.de/no_cache/programm/kalender/

3

Bibliographiedatenbank dblp

dblp computer science bibliography

Offene Bibliographiedaten für die Informatik

3.1

Open Bibliographic Data in Computer Science

Moderne Informatik-Forschung benötigt den unmittelbaren und umfassenden Zugriff auf aktuelle Publikationen, um den Bedürfnissen in einer sich immer schneller entwickelnden und immer komplexer werdenden Forschungslandschaft gerecht zu werden. Doch nicht nur im Forscheralltag, auch bei der Einschätzung von Forschungsleistung ist die Verfügbarkeit verlässlicher Publikationsdaten unverzichtbar. Hoch qualitative und vollständige Metadaten sind in der Regel jedoch nur sehr schwer zu erhalten. Freie Suchmaschinen wie etwa Google erlauben einen weiten Einblick in das Internet, besitzen aber keinerlei Qualitätsgarantien oder semantische Organisation. Kommerzielle Datenbanken verkaufen Metadaten als teure Dienstleistung, weisen aber in vielen Fachdisziplinen (wie etwa in der Informatik) nur eine mangelhafte Abdeckung und eine oft ungenügende Datenqualität auf. Insbesondere die einzigartige Publikationskultur der Informatik mit ihrem Schwerpunkt auf Konferenzpublikationen bleibt dabei unberücksichtigt, da für kommerzielle Anbieter hier die Breite des Marktes zu fehlen scheint. Universitäten und außeruniversitäre Forschungseinrichtungen bemühen sich oftmals mit immensen personellen und finanziellen Aufwand und unter Belastung der einzelnen forschenden Akteure, eigene Daten zu erheben. Diese Datensätze weisen jedoch zwangsläufig einen lokalen Einschlag auf und vermögen es nicht, ein detailliertes Bild einer Forschungsdisziplin als Ganzes zu zeichnen.

Die „dblp computer science bibliography“ leistet auf diesem Gebiet nun bereits seit über 25 Jahren einen substanziellen Beitrag durch die offene Bereitstellung qualitätsgeprüfter und aufbereiteter Publikationsdaten für die gesamte Informatik. Dabei unterstützt dblp die Informatik-Forschung auf gleich mehreren Ebenen, etwa durch:

- Unterstützung der täglichen Forschungsarbeit, etwa bei der Literaturrecherche und dem Bezug von verfügbaren Volltexten
- Unterstützung des wissenschaftlichen Publikationsprozesses durch die Bereitstellung normierter bibliographischer Referenzdaten
- Unterstützung von Forschern und Institutionen bei der Berichtspflicht durch die Sammlung und Aufbereitung von qualitätsgesicherten Publikationslisten
- Unterstützung von Forschungsförderern und Entscheidungsträgern durch das öffentliche Verfügbarmachen von nach Daten-Facetten aufgeschlüsselten Publikationsnachweisen

Darüber hinaus ist der dblp-Datensatz selbst Untersuchungsgegenstand mehrerer tausend Fachartikel.⁹ Insgesamt ist dblp daher für die Informatik sowohl als Recherche-Tool, aber auch als Forschungsdatensatz unverzichtbar geworden.

Modern computer science research requires the immediate and comprehensive access to current publications to meet the needs of an ever faster evolving and ever more complex research landscape. Not only in the everyday work of a researcher but also in the assessment of research performance, the availability of reliable bibliographic metadata has become indispensable. However, high-quality and complete metadata is very difficult to obtain. Free search engines like Google allow a broad insight into the Internet but have neither guarantees of quality nor any semantic organization. Commercial databases sell metadata as an expensive service, but in many disciplines (such as in computer science), their coverage is insufficient and the data quality is quite poor. In particular, the unique publication culture of computer science with its emphasis on conference publications remains disregarded, as for commercial providers the width of the market seems to be missing here. Most universities and non-university research institutions endeavor to collect their own data, yet often consume enormous human and financial resources and impose a burden on the individual researchers. However, these local data sets do inevitably have a local bias and are not suited to draw a detailed picture of a research discipline as a whole.

For over 25 years now, the “dblp computer science bibliography” has substantially contributed to solving this dilemma in the field of computer science by providing open, quality-checked, and curated bibliographic metadata. The dblp web service supports the computer science research community on several levels, for example by:

- supporting researchers in their daily work, e.g., when reviewing the literature or searching for full-text research articles
- supporting the scientific publication process by providing standardized bibliographic reference data
- supporting researchers and institutions in their reporting duties by collecting and editing quality-assured bibliographies
- supporting research funders and decision-makers, e.g., by providing publicly available and explorable bibliographic references

In addition, the dblp data set itself is object of study of several thousand research articles.¹⁰ Hence, dblp has become indispensable to the computer science community as both a research tool and a research data set.

⁹ Google Scholar liefert zum Suchbegriff „dblp“ über 34 200 Treffer; im Einzelnen weisen SpringerLink ca. 3 400 Artikel, Elsevier ScienceDirect über 720 Artikel, die ACM Digital Library ca. 2 100 Artikel und IEEE Xplore über 2 300 Artikel nach.

¹⁰ The search term “dblp” results in 34,200 hits at Google Scholar; in particular, SpringerLink lists about 3,400 articles, Elsevier ScienceDirect lists more than 720 articles, the ACM Digital Library lists 2,100 articles, and IEEE Xplore lists more than 2,300 articles.

Schloss Dagstuhl und dblp

3.2

Schloss Dagstuhl and dblp

3

Bereits seit Ende 2010 engagiert sich Schloss Dagstuhl für die ursprünglich an der Universität Trier entwickelten Bibliographiedatenbank dblp. Zunächst durch ein Projekt im Leibniz-Wettbewerb gefördert, wurde die Datenbank seit Juni 2013 von Schloss Dagstuhl direkt mitfinanziert. Im Zuge der Konsolidierung der Zusammenarbeit mit der Universität Trier wurden unter dem Dach von Schloss Dagstuhl Mitarbeiterstellen im wissenschaftlichen Stab geschaffen, die hauptamtlich für die Betreuung und Weiterentwicklung von dblp beauftragt sind. Der unter dem Dach von Schloss Dagstuhl gegründete dblp-Beirat (siehe Fig. 3.1) leistet seit November 2011 die wissenschaftliche Aufsicht und unterstützt das dblp-Team mit seiner Expertise.

Pünktlich zum 25-jährigen Jubiläum von dblp im November 2018 erfolgte die endgültige Staffelübergabe des Betriebes der Datenbank von der Universität Trier an Schloss Dagstuhl. Damit einhergehend wurden durch die Zuwendungsgeber weitere Mittel für den Betrieb von dblp bereit gestellt und eine eigens neu eingerichtete Außenstelle von Schloss Dagstuhl auf dem Campus II der Universität Trier angesiedelt. Betrieb und die Erforschung der Datenbank erfolgen dabei weiterhin in enger Kooperation mit dem Fach Informatikwissenschaften der Universität sowie dem Trierer Center for Informatics Research and Technology (CIRT).¹¹


The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – has existed since late 2010. The commitment of Schloss Dagstuhl to dblp, initially funded by a project of the Leibniz Competition, has been funded directly by Schloss Dagstuhl since June 2013. As part of the consolidation of this cooperation, scientific staff positions – assigned full-time to the support and development of dblp – were created. The dblp advisory board (c.f. Figure 3.1), established in November 2011 at Schloss Dagstuhl, provides scientific supervision and supports dblp with its expertise.

In November 2018, the transfer of the database from the University of Trier to the Leibniz Center for Informatics in Schloss Dagstuhl took place just in time for dblp's 25th anniversary. At the same time, Dagstuhl's funding had been increased to support the operation of dblp and a new Schloss Dagstuhl branch office for the dblp team has been established on Campus II of the University of Trier. The database will continue to be operated and researched in close cooperation with the University's Department of Computer Sciences and the Trierer Center for Informatics Research and Technology (CIRT).¹¹


¹¹ <https://cirt.uni-trier.de/>

dblp-Beirat dblp Advisory Board
Prof. Dr. Hannah Bast University of Freiburg, Germany <i>Chair</i>
Prof. Dr. Andreas Butz Ludwig Maximilians University Munich, Germany
Prof. Dr.-Ing. Rüdiger Dillmann Karlsruhe Institute of Technology, Germany
Prof. Dr. Hans-Peter Lenhof Saarland University, Germany
Prof. Dr. Mila Majster-Cederbaum Ludwig Maximilian University of Munich, Germany
Prof. Dr. Andreas Oberweis Karlsruhe Institute of Technology, Germany
Prof. Dr. Rüdiger Reischuk University of Lübeck, Germany
Prof. Dr. Dietmar Saupe University of Konstanz, Germany
Prof. Dr. Dr. h.c. Otto Spaniol RWTH Aachen, Germany
Prof. Dr.-Ing. Jürgen Teich University of Erlangen-Nuremberg, Germany
Prof. Dr. Dr. h.c. Reinhard Wilhelm Saarland University, Germany

Fig. 3.1
dblp Advisory Board.



(a) Total number of records by year and type



(b) New records by year and type

Fig. 3.2

Development of the dblp data stock.**Statistiken der Datenakquise****3.3****Data Acquisition Statistics**

Die Bibliophiedatenbank dblp indexiert Publikationen anhand vollständiger Inhaltsverzeichnisse von Konferenzbänden oder Journalausgaben. Mit Hilfe einer eigens entwickelten Software zur Datenextraktion werden Metadaten von Verlagswebseiten ausgelesen und zur weiteren Bearbeitung vorbereitet. Die Metadaten werden anschließend vom dblp-Team redaktionell bearbeitet: Eventuelle Fehler werden korrigiert, mehrdeutige und ungenaue Angaben werden verbessert. Diese Datenpflege wird zwar von Hilfssoftware unterstützt, erfolgt aber vornehmlich händisch durch den jeweiligen Mitarbeiter.

Ende Dezember 2018 verzeichnete dblp bereits mehr als 4,4 Millionen Publikationen. Im Laufe des Jahres wurden somit eine erneute Rekordanzahl an 420 000 neuen Publikationseinträgen aufgenommen. Dies entspricht mehr als 1 680 neuen Publikationen pro Arbeitstag. Die neu aufgenommenen Einträge verteilen sich zu 51,3% auf Konferenzbeiträge, zu 35,5% auf Journalartikel, zu 10,1% auf Preprints und „graue“ Literatur, sowie zu 3,1% auf andere Publikationstypen.

Ein Überblick über die Entwicklung der Datenakquise kann Fig. 3.2a und Fig. 3.2b entnommen werden.

The dblp computer science bibliography indexes conferences and journals on a per-volume basis. Using dblp's own web harvesting software, bibliographic metadata of journal or proceedings volumes are extracted from the publisher's website. This metadata 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 the end of December 2018, more than 4.4 million publications have been listed in dblp. During the course of 2018, the dblp database grew by more than 420,000 publication records. This is the largest figure ever achieved in the history of dblp and corresponds to more than 1,680 new records for each working day of the year. This year's new records consist of 51.3% conference papers, 35.5% journal articles, 10.1% preprints and "grey" literature, and 3.1% further publications.

The development of the dblp dataset is summarized in Figure 3.2a and Figure 3.2b.

	Trier 1		Trier 2		Dagstuhl		Total	
	2017	2018	2017	2018	2017	2018	2017	2018
user sessions (visits) per day	27 931	31 530	2 836	3 233	5 366	11 483	36 133	46 247
page views per day	466 989	618 067	35 140	20 208	85 537	202 301	587 668	840 577
page views per user session	16,7	19,6	12,4	6,2	15,9	17,6	16,3	18,2
distinct users (IPs) per month	390 886	451 769	58 975	27 448	86 985	197 270	536 847	676 489
data served per month	1 082,3 GB	1 535,0 GB	82,8 GB	72,6 GB	235,0 GB	469,7 GB	1 400,1 GB	2 077,3 GB

Fig. 3.3

Average usage of the three dblp servers. Trier 1 = dblp.uni-trier.de, Trier 2 = dblp2.uni-trier.de, Dagstuhl = dblp.dagstuhl.de

Nutzungsstatistiken

3.4

Usage Statistics

Im Jahr 2018 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

- Server Trier 1: dblp.uni-trier.de
- Server Trier 2: dblp2.uni-trier.de
- Server Dagstuhl: dblp.dagstuhl.de

Die allgemeine Adresse dblp.org ist dabei ein Alias für den dblp-Server in Dagstuhl.

Seit Mitte 2014 stehen vergleichbare Nutzerstatistiken von allen drei dblp-Servern zur Verfügung. Dabei war Server Trier 1 in der Vergangenheit aufgrund seiner prominenten Sichtbarkeit in den Google-Suchergebnissen die mit Abstand bekannteste Adresse. Im Laufe des Jahres 2018 konnte Server Dagstuhl jedoch zu Trier 1 aufschließen. Bereits zu Beginn von 2019 sind Server Dagstuhl und Server Trier 1 bezüglich der Anzahl der Nutzer sowie der Platzierung bei Google gleich auf.

Fig. 3.3 fasst die durchschnittliche Nutzung aller drei dblp-Server zusammen. Diese Statistiken ignorieren die Zugriffe, die durch bekannte Bot- und Crawler-Software verursacht wurden.

In 2018, three official dblp web servers were updated and synchronized on a daily basis:

- server Trier 1: dblp.uni-trier.de
- server Trier 2: dblp2.uni-trier.de
- server Dagstuhl: dblp.dagstuhl.de

The main domain dblp.org is used as an alias for dblp server Dagstuhl.

Starting in mid-2014, usage data have been collected on all three mirror sites. In the past, Trier 1 had been the most widely known server due to its high visibility and prominence in the Google search engine. However, during the course of 2018, server Dagstuhl has become increasingly more visible. Starting in 2019, servers Dagstuhl and Trier 1 are already on the same level with respect to number of users and Google search ranking.

Figure 3.3 shows the average usage of all three servers in 2018. These figures ignore the traffic caused by known bots and crawlers.

4 **Dagstuhl Publishing** *Dagstuhl Publishing*

Portfolio

4.1

Portfolio

Die Open-Access-Verlagsdienstleistungen von Schloss Dagstuhl werden in der Wissenschaftsgemeinde gut aufgenommen. Im Portfolio des Angebots gibt es zum einen Publikationsserien, die sich auf Veranstaltungen beziehen, die auf Schloss Dagstuhl abgehalten wurden (*Dagstuhl Reports*, *Dagstuhl Manifestos*, *Dagstuhl Follow-Ups*), zum anderen Serien, die Konferenzen und Workshops außerhalb von Schloss Dagstuhl bedienen (*LIPICs* und *OASICs*). Ergänzt wird das Portfolio seit 2013 um die wissenschaftliche Zeitschrift *LITES* und seit 2015 um die Serie *DARTS*, in der Forschungsartefakte veröffentlicht werden.

■ Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*¹² dokumentiert um eine Zitation der Seminare im wissenschaftlichen Kontext zu ermöglichen. Zudem erlaubt es auch den Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Dagstuhl-Seminaren und -Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 11.4) agiert als Herausbergremium für die Reihe. Um umfassende Zusammenstellungen von begutachteten Artikeln auf Basis eines Dagstuhl-Seminars oder -Perspektiven-Workshops zu ermöglichen, wurde die Buchreihe *Dagstuhl Follow-Ups* (siehe unten) gegründet.

In 2018 wurde für 78 Dagstuhl-Seminare und -Perspektiven-Workshops ein Bericht in der Reihe *Dagstuhl Reports* veröffentlicht. An dieser Stelle bedanken wir uns ganz herzlich bei den Organisatoren und Kollektoren für die erfolgreiche Zusammenarbeit.

■ Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*¹³ die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Dagstuhl-Perspektiven-Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 11.4) fungiert hier ebenfalls als Herausbergremium. In 2018 wurde eine Ausgabe mit einem Manifesto veröffentlicht (siehe Fig. 4.1).

¹² <http://drops.dagstuhl.de/dagrep>

¹³ <http://drops.dagstuhl.de/dagman>

The scientific community appreciates the Open Access publishing services offered by Schloss Dagstuhl. The portfolio covers series related to events at Schloss Dagstuhl (*Dagstuhl Reports*, *Dagstuhl Manifestos*, *Dagstuhl Follow-Ups*) and series for conferences and workshops held outside of Schloss Dagstuhl (*OASICs* and *LIPICs*). The portfolio is supplemented by the scholarly journal *LITES* since 2013 and by the *DARTS* series which aims at publishing research artifacts since 2015.

■ Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*¹² which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

The periodical started with the first seminars of January 2011 and publishes, in monthly issues, reports on Dagstuhl Seminars and Perspectives Workshops that took place in a given month. The content is not peer-reviewed. The Scientific Directorate (see Fig. 11.4) acts as editorial board. For comprehensive collections of peer-reviewed articles developed on the basis of a Dagstuhl Seminar or Perspectives Workshop, we offer seminar organizers the possibility of publishing a volume in our book series *Dagstuhl Follow-Ups* (see below).

In 2018, 78 reports of Dagstuhl Seminars and Dagstuhl Perspectives Workshops have been published. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

■ Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*¹³ in an Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. In 2018 one volume with one manifesto was published (see Fig. 4.1).

Engineering Academic Software (Dagstuhl Perspectives Workshop 16252)
<https://dx.doi.org/10.4230/DagMan.6.1.1>
 based on Dagstuhl Perspectives Workshop 16252 <https://www.dagstuhl.de/16252>

Fig. 4.1

Manifestos published in the 2018 volume of the journal *Dagstuhl Manifestos*.

■ Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*¹⁴ ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert. Für jedes Buch ist ein Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausbergremium verantwortlich ist) begutachtet und freigegeben werden muss. In 2018 wurde kein Buch in der Reihe veröffentlicht.

■ OASlcs: OpenAccess Series in Informatics

Die *OASlcs*-Reihe¹⁵ veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausbergremium (Fig. 4.2), diskutiert sorgfältig alle Anträge, um ausschließlich qualitativ hochwertige sowie professionell durchgeführte Veranstaltungen in die Reihe aufzunehmen und um gegebenenfalls Empfehlungen zur Verbesserung der Veranstaltungsstruktur zu geben.

In 2018 wurden 7 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.3.

■ Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*¹⁴ book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Perspectives Workshop. Each book requires a proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). In 2018, no volume was published in the series.

■ OASlcs: OpenAccess Series in Informatics

The *OASlcs* series¹⁵ aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 4.2, discusses carefully all submitted proposals to ensure that only significant and professionally organized events are added to the series and that – if applicable – suggestions are given for improving the structure of the event.

In 2018, Dagstuhl published 7 *OASlcs* volumes covering the proceedings of typically widespread workshops and conferences; see Fig. 4.3.

¹⁴ <http://drops.dagstuhl.de/dfu>

¹⁵ <http://drops.dagstuhl.de/oasics>

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 Chair

Fig. 4.2
OASlcs Editorial Board.

Vol. 58 Technical Communications of the 32nd International Conference on Logic Programming (ICLP 2017) https://www.dagstuhl.de/dagpub/978-3-95977-058-3
Vol. 60 2017 Imperial College Computing Student Workshop (ICCSW 2017) https://www.dagstuhl.de/dagpub/978-3-95977-059-0
Vol. 61 1st Symposium on Simplicity in Algorithms (SOSA 2018) https://www.dagstuhl.de/dagpub/978-3-95977-064-4
Vol. 62 7th Symposium on Languages, Applications and Technologies (SLATE 2018) https://www.dagstuhl.de/dagpub/978-3-95977-072-9
Vol. 63 18th International Workshop on Worst-Case Execution Time Analysis (WCET 2018) https://www.dagstuhl.de/dagpub/978-3-95977-073-6
Vol. 64 Technical Communications of the 34th International Conference on Logic Programming (ICLP 2018) https://www.dagstuhl.de/dagpub/978-3-95977-090-3
Vol. 65 18th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-096-5

Fig. 4.3
OASlcs volumes published in 2018.

■ LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs-Reihe*¹⁶ veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausgebergremium (siehe Fig. 4.4) besteht aus einschlägig bekannten Wissenschaftlern und wird seit Oktober 2017 von Luca Aceto als Hauptherausgeber geleitet.

Die Amtszeiten von Chris Hankin und Deepak Kapur sind 2018 ausgelaufen. Beide haben sowohl als langjährige Mitglieder des Herausgebergremiums eine wichtige Rolle in der Entwicklung der Serie gespielt. Für diese Verdienste möchten wir uns an dieser Stelle herzlich bedanken.

Christel Baier und Javier Esparza wurden in einem anonymen Wahlverfahren innerhalb des Herausgebergremiums neu in das Gremium gewählt. Siehe auch Fig. 4.4.

In 2018 wurden Tagungsbände von 32 Konferenzen mit insgesamt 1387 Artikeln veröffentlicht, so viel wie noch nie zuvor (in 2017 waren es 25 Bände mit 1127 Artikeln); siehe Fig. 4.5 und 4.6.

In 2018 gab es erneut viele Anträge bei LIPIcs, womit die große Nachfrage aus den Vorjahren fortgesetzt wurde. Die große Anzahl an Anträgen sind die erfreulichen Ergebnisse unserer langjährigen Bemühungen, einige der wichtigsten Konferenzen an LIPIcs zu binden. In Fig. 4.7 sind alle Konferenzen aufgelistet, deren Anträge 2018 bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischerweise 5 Jahre) eingegangen wurde. Drei dieser Konferenzen haben erstmals einen Antrag bei LIPIcs gestellt. Die anderen Konferenzen haben bereits vorher mit LIPIcs kooperiert.

■ LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*¹⁷ veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 4.8), begutachtet alle eingereichten Arbeiten. Die Zeitschrift wird gemeinsam mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*¹⁸ der Fachgesellschaft *European Design and Automation Association (EDAA)*¹⁹ herausgegeben. Die Fachgruppe ist dabei für die Besetzung des Herausgebergremiums verantwortlich, während Schloss Dagstuhl die administrativen Aufgaben der Herausgeberschaft übernimmt.

Im Gegensatz zu anderen Zeitschriften im Bereich eingebetteter Systeme, steht bei *LITES* eine moderate Veröffentlichungsgebühr (article-processing charge, APC) sowie ein schnelles Begutachtungsverfahren (innerhalb eines Jahres ab Einreichung) im Vordergrund.

In 2018 wurde eine Ausgabe von *LITES* mit insgesamt 4 Artikeln veröffentlicht.

■ LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs series*¹⁶ publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers (see Fig. 4.4) supervises the conferences that are accepted for LIPIcs and is headed since October 2017 by Luca Aceto.

The terms of Chris Hankin and Deepak Kapur ended in 2018. Both served as members of the editorial board for several years and played an important role for the development of the series. We would like to take this opportunity to thank them for their extraordinary dedication.

In 2018, Christel Baier and Javier Esparza were voted in an anonymous voting within the editorial board as new members of the editorial board. See also Fig. 4.4.

The series published the proceedings of 32 major conferences with more than 1387 articles in total in 2018, marking again a record high since the series was started (in 2017 there were 25 volumes with 1127 articles in total); see Fig. 4.5 and 4.6.

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2018 has again seen several applications for LIPIcs, continuing the high interest from the previous years. Fig. 4.7 lists all conferences that have been accepted in 2018 for a cooperation covering several years (typically 5 years). Three of these conferences have submitted a proposal to LIPIcs for the first time. The other conferences have already cooperated with LIPIcs in the past.

■ LITES: Leibniz Transactions on Embedded Systems

The *LITES*¹⁷ journal publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was established in 2012 and started operating in early 2013. A broad team of experienced researchers, acting as editorial board (see Fig. 4.8), reviews all submitted contributions. The journal is jointly published with the *EMbedded Systems Special Interest Group (EMSIG)*¹⁸ of the *European Design and Automation Association (EDAA)*¹⁹. The special interest group is responsible for appointing the editorial board, while Schloss Dagstuhl takes over the administrative tasks of the publication.

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.

In 2018, one issue of *LITES* containing 4 articles in total was published.

¹⁶ <http://drops.dagstuhl.de/lipics>

¹⁷ <http://drops.dagstuhl.de/lites>

Prof. Dr. Luca Aceto Gran Sasso Science Institute, Italy and Reykjavik University, Iceland Chair
Prof. Dr. Susanne Albers Technical University Munich, Germany
Prof. Dr. Christel Baier Technische Universität Dresden, Germany tenure started in June 2018
Prof. Dr. Chris Hankin Imperial College London, United Kingdom tenure ended in May 2018
Prof. Dr. Javier Esparza Technical University Munich, Germany tenure started in June 2018
Prof. Deepak Kapur, Ph. D. University of New Mexico, US tenure ended in May 2018
Prof. Michael Mitzenmacher, Ph. D. Harvard University, US
Prof. Madhavan Mukund, Ph. D. Chennai Mathematical Institute, India
Prof. Dr. Anca Muscholl LaBRI and University Bordeaux, France
Dr. Catuscia Palamidessi INRIA, France
Prof. Dr. Thomas Schwentick TU Dortmund, Germany
Prof. Raimund Seidel, Ph. D. Saarland University, Germany
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Saarland University, Germany

Fig. 4.4

LIPICs Editorial Board.

Vol. 69 21st International Conference on Types for Proofs and Programs (TYPES 2015) https://www.dagstuhl.de/dagpub/978-3-95977-030-9
Vol. 73 12th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2017) https://www.dagstuhl.de/dagpub/978-3-95977-034-7
Vol. 89 12th International Symposium on Parameterized and Exact Computation (IPEC 2017) https://www.dagstuhl.de/dagpub/978-3-95977-051-4
Vol. 93 37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2017) https://www.dagstuhl.de/dagpub/978-3-95977-055-2
Vol. 94 9th Innovations in Theoretical Computer Science Conference (ITCS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-060-6
Vol. 95 21st International Conference on Principles of Distributed Systems (OPDIS 2017) https://www.dagstuhl.de/dagpub/978-3-95977-061-3
Vol. 96 35th Symposium on Theoretical Aspects of Computer Science (STACS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-062-0
Vol. 97 22nd International Conference on Types for Proofs and Programs (TYPES 2016) https://www.dagstuhl.de/dagpub/978-3-95977-065-1
Vol. 98 21st International Conference on Database Theory (ICDT 2018) https://www.dagstuhl.de/dagpub/978-3-95977-063-7
Vol. 99 34th International Symposium on Computational Geometry (SoCG 2018) https://www.dagstuhl.de/dagpub/978-3-95977-066-8

Fig. 4.5

LIPICs volumes published in 2018 – Part 1.

Vol. 100 9th International Conference on Fun with Algorithms (FUN 2018) https://www.dagstuhl.de/dagpub/978-3-95977-067-5
Vol. 101 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018) https://www.dagstuhl.de/dagpub/978-3-95977-068-2
Vol. 102 33rd Computational Complexity Conference (CCC 2018) https://www.dagstuhl.de/dagpub/978-3-95977-069-9
Vol. 103 17th International Symposium on Experimental Algorithms (SEA 2018) https://www.dagstuhl.de/dagpub/978-3-95977-070-5
Vol. 105 29th Annual Symposium on Combinatorial Pattern Matching (CPM 2018) https://www.dagstuhl.de/dagpub/978-3-95977-074-3
Vol. 106 30th Euromicro Conference on Real-Time Systems (ECRTS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-075-0
Vol. 107 45th International Colloquium on Automata, Languages, and Programming (ICALP 2018) https://www.dagstuhl.de/dagpub/978-3-95977-076-7
Vol. 108 3rd International Conference on Formal Structures for Computation and Deduction (FSCD 2018) https://www.dagstuhl.de/dagpub/978-3-95977-077-4
Vol. 109 32nd European Conference on Object-Oriented Programming (ECOOP 2018) https://www.dagstuhl.de/dagpub/978-3-95977-079-8
Vol. 110 29th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AoFA 2018) https://www.dagstuhl.de/dagpub/978-3-95977-078-1
Vol. 111 13th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2018) https://www.dagstuhl.de/dagpub/978-3-95977-080-4
Vol. 112 26th Annual European Symposium on Algorithms (ESA 2018) https://www.dagstuhl.de/dagpub/978-3-95977-081-1
Vol. 113 18th International Workshop on Algorithms in Bioinformatics (WABI 2018) https://www.dagstuhl.de/dagpub/978-3-95977-082-8
Vol. 114 10th International Conference on Geographic Information Science (GIScience 2018) https://www.dagstuhl.de/dagpub/978-3-95977-083-5
Vol. 116 Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018) https://www.dagstuhl.de/dagpub/978-3-95977-085-9
Vol. 117 43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-086-6
Vol. 118 29th International Conference on Concurrency Theory (CONCUR 2018) https://www.dagstuhl.de/dagpub/978-3-95977-087-3
Vol. 119 27th EACSL Annual Conference on Computer Science Logic (CSL 2018) https://www.dagstuhl.de/dagpub/978-3-95977-088-0
Vol. 120 25th International Symposium on Temporal Representation and Reasoning (TIME 2018) https://www.dagstuhl.de/dagpub/978-3-95977-089-7
Vol. 121 32nd International Symposium on Distributed Computing (DISC 2018) https://www.dagstuhl.de/dagpub/978-3-95977-092-7
Vol. 122 38th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2018) https://www.dagstuhl.de/dagpub/978-3-95977-093-4
Vol. 123 29th International Symposium on Algorithms and Computation (ISAAC 2018) https://www.dagstuhl.de/dagpub/978-3-95977-094-1

Fig. 4.6

LIPIcs volumes published in 2018 – Part 2.

AofA	International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms accepted for 2018–2022
DISC	International Symposium on Distributed Computing accepted for 2017–2021
FSTTCS	Foundations of Software Technology and Theoretical Computer Science accepted for 2019–2023 (Re-evaluation)
ITCS	Innovations in Theoretical Computer Science Conference accepted for 2017–2021
STACS	Symposium on Theoretical Aspects of Computer Science accepted for 2019–2023 (Re-evaluation)

Fig. 4.7

Conferences accepted in 2018 for publication in LIPIcs.

Prof. Alan Burns, DPhil University of York, UK Editor-in-Chief
Prof. Sang Lyul Min, Ph. D. 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, Ph. D. Lund University, Sweden Subject area: Control
Prof. Steve Goddard, Ph. D. University of Nebraska-Lincoln, US Subject area: Cyber-physical systems
Prof. Dr. Axel Jantsch Technical University of Vienna, Austria Subject area: Distributed embedded systems and networks
Prof. Bashir Al Hashimi University of Southampton, UK Subject area: Energy-efficiency
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, Ph. D. University of North Carolina at Chapel Hill, US Subject area: Scheduling
Prof. Dr. Florence Maraninchi University of Grenoble, France and Verimag Lab, France Subject area: Verification, formal methods, model-based design

Fig. 4.8

LITES Editorial Board.

■ DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*²⁰ werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2018 wurde die vierte Ausgabe mit drei Heften und insgesamt 19 Artefakten veröffentlicht.

Die Veröffentlichung und Bereitstellung von Forschungsdaten und -artefakten ist aktuell ein wichtiges Thema in den wissenschaftlichen Disziplinen und bei den Forschungsfördereinrichtungen. Im Bereich der Informatik wird dieses Thema ebenfalls diskutiert. In 2015 gab es zum Beispiel einen Perspektiven-Workshop mit dem Titel „Artifact Evaluation for Publications“²¹, der in 2016 durch zwei Seminare ergänzt wurde: „Reproducibility of Data-Oriented Experiments in e-Science“²² und „Rethinking Experimental Methods in Computing“²³.

Schloss Dagstuhl unterstützt mit DARTS die Wissenschaftsgemeinde in der Informatik bei dem Wunsch, Forschungsdaten und -artefakte in einer geeigneten Reihe zu veröffentlichen. Hierbei berücksichtigt DARTS insbesondere auch die Publikationskultur in der Informatik mit ihrem Schwerpunkt auf Konferenzbandveröffentlichungen.

■ DARTS: Dagstuhl Artifacts Series

The *DARTS* series²⁰ publishes evaluated research data and artifacts. It is organized as a periodical. In 2018, one volume containing three issues with 19 artifacts in total was published.

The publishing of research data and artifacts is currently in the general focus of the scientific community and funding agencies. In the area of computer science, this topic is also under discussion. For example, in 2015 a Perspectives Workshop on “Artifact Evaluation for Publications”²¹ took place which was complemented with two seminars in 2016: “Reproducibility of Data-Oriented Experiments in e-Science”²² and “Rethinking Experimental Methods in Computing”²³.

With DARTS, Schloss Dagstuhl is aiming to support the computing research community with a publishing venue dedicated to research data and artifacts. Especially, DARTS takes into account the publication culture in computer science which focuses on conference proceedings publications.

Infrastruktur

4.2

Infrastructure

■ Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 4.9. Die Bände aus den Reihen *LIPICs* und *OASICs* werden zudem bei Scopus²⁴ eingereicht, wo sie regelmäßig indiziert werden. Die Reihen *LIPICs* und *OASICs* sowie die Zeitschrift *LITES* sind zudem im Directory of Open Access Journals (DOAJ) gelistet, siehe Fig. 4.9.

Zudem unterstützen die technischen Schnittstellen die Datenakquisition durch Google Scholar, so dass die Publikationen sichtbar und besser recherchierbar sind.

■ LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*²⁵ ein Online-Repository ins Leben gerufen, um Open Access-Veröffentlichungen von Leibniz-Instituten und deren Wissenschaftlern zu unterstützen und sichtbar zu machen. Schloss Dagstuhl liefert alle Artikel aus den Reihen *Dagstuhl Reports* und *Dagstuhl Manifestos* an das Repository und stärkt dadurch Forschungsergebnisse aus der Informatik innerhalb dieses multidisziplinären Repositoriums.

■ Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 4.9. The *LIPICs* and *OASICs* volumes are submitted to Scopus²⁴ where they are regularly indexed. The *LIPICs* and *OASICs* series as well as the journal *LITES* are also listed in the Directory of Open Access Journals (DOAJ), see Fig. 4.9.

The technical interface of our publication server enables harvesting according to the Google Scholar guidelines. Google Scholar regularly retrieves metadata and full-texts from our server.

■ LeibnizOpen

The Leibniz Association has established the *Leibniz-Open*²⁵ repository 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.

¹⁸ <http://www.emsig.net/>

¹⁹ <https://www.edaa.com/>

²⁰ <https://www.dagstuhl.de/darts>

²¹ <https://www.dagstuhl.de/15452>

²² <https://www.dagstuhl.de/16041>

²³ <https://www.dagstuhl.de/16111>

²⁴ <https://www.scopus.com>

²⁵ <http://www.leibnizopen.de/>

■ AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Predatory Publishing“²⁶ mit organisiert, welcher bereits der fünfte Workshop in Folge seit 2013 ist. Der Workshop findet am 17. und 28. Januar 2019 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt.

■ Technisches Back-end: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)²⁷ werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative²⁸ berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer grossen Leserschaft zur Verfügung. Als technische Grundlage dient eine adaptierte Version des OPUS-Systems.²⁹

■ Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Predatory Publishing”²⁶ was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop will take place at the Leibniz Association headquarters in Berlin on January 17 and 18, 2019.

■ Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)²⁷. The general guidelines of the Dublin Core initiative²⁸ applicable to online publications are adhered to, meaning that all the requisite metadata of each publication is stored, thus ensuring availability in the long term. This enables the online publications to be cited by and accessible to a wide readership. The technical basis for this is an adapted version of the OPUS system.²⁹

²⁶ <https://www.dagstuhl.de/fileadmin/dagpub/journalmanagement-leibniz/2019-01-workshop/>

²⁷ <https://www.dagstuhl.de/drops>

²⁸ <http://dublincore.org/>

²⁹ <https://opus4.kobv.de/>

dblp
Dagstuhl Reports https://dblp.org/db/journals/dagstuhl-reports/
Dagstuhl Manifestos https://dblp.org/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups https://dblp.org/db/series/dfu/
OASlcs https://dblp.org/db/series/oasics/
LIPlcs https://dblp.org/db/series/lipics/
LITES https://dblp.org/db/journals/lites/
DARTS https://dblp.org/db/journals/darts/
DOAJ
OASlcs https://doaj.org/toc/2190-6807
LIPlcs https://doaj.org/toc/1868-8969
LITES https://doaj.org/toc/2199-2002

Fig. 4.9

Indexing of Dagstuhl Publishing series in dblp and DOAJ.

■ Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)³⁰ zur (digitalen) Langzeitarchivierung eingereicht.

■ Mirroring

Um dem Verlust von Daten vorzubeugen, werden seit 2010 zwei Kooperationen zur Spiegelung (Mirroring) von Inhalten des Publikationsservers DROPS gepflegt:

- emis.de: Das unter Leitung des FIZ Karlsruhe, Leibniz-Institut für Informationsinfrastruktur, organisierte Mathematik-Publikations-Portal European Mathematical Information Service (EMIS) spiegelt alle Bände der LIPIcs-Reihe.³¹
- SunSite Central Europe: Der Sun-Server-Park, der an der RWTH Aachen unter Leitung von Prof. Matthias Jarke betrieben wird, bietet eine Heimat für zahlreiche Software-Archive als auch Publikationen. Der gesamte DROPS-Bestand wird nun in regelmäßigen Abständen auf der SunSite Aachen gespiegelt.³²

■ Long-term Archiving

All publications are submitted to the German National Library (D-NB)³⁰ for (digital) long-term archiving.

■ Mirroring

In order to prevent data loss, two cooperative ventures were initiated in 2010 for mirroring the content of the DROPS publication server:

- emis.de: The portal for electronic math resources European Mathematical Information Service (EMIS), organized under the auspices of FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, mirrors all volumes of the LIPIcs series³¹.
- SunSite Central Europe: The Sun server park, located at the Aachen University of Technology and operated under the guidance of Prof. Matthias Jarke, is home to numerous software archives and publications. All the DROPS assets are now mirrored at regular intervals on the Aachen SunSite.³²

³⁰ http://www.dnb.de/DE/Netzpublikationen/Langzeitarchivierung/langzeitarchivierung_node.html

³¹ <https://subs.emis.de/LIPIcs/>

³² <http://vesta.informatik.rwth-aachen.de/Dagstuhl/>

5 Resonanz *Feedback*

Resonanz zu Seminaren und Workshops

5.1

Feedback on Seminars and Workshops

■ Resonanz von Teilnehmern

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

17221 – Geometric Modelling, Interoperability and New Challenges | Dagstuhl Seminar | <https://www.dagstuhl.de/17221>

I have been attending the meetings on CAGD (under various meeting titles) at Dagstuhl for decades, and they have always been a high point, meeting the right people and addressing pertinent questions, which have developed between each meeting and the next.[...]

You probably don't need to be told how good the support environment at Dagstuhl is, and how easy it is to feel comfortable. Food, accommodation, facilities - all excellent. Thank you.

Malcolm Sabin

Karl-Heinz Küfer

18031 – Personalized Multiobjective Optimization: An Analytics Perspective | Dagstuhl Seminar | <https://www.dagstuhl.de/18031>

The Dagstuhl seminars, the discussions with colleagues and the talks have influenced my way of thinking and finally contributed heavily to the success of my own research efforts. Besides this meanwhile two younger researchers that visited the Dagstuhl seminars together with me applied at Fraunhofer ITWM and were appointed with postdoc positions. Ideas for several publications and projects have been originated at the Dagstuhl seminars. I am looking forward to the next announcement of a multicriteria decision making seminar at Dagstuhl and I will happily join if I am invited.

■ Resonanz unserer Organisatoren

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

■ Feedback from Organizers

The success of Schloss Dagstuhl depends to a large extent on our outstanding seminar organizers, who continually enrich the scientific program with a range of interesting and new topics. We are very glad to be able to provide services and an environment that organizers appreciate. The following comments from organizers are excerpted from the Dagstuhl Report or personal emails to us. We cite them with their kindly permission.

Organizers of Dagstuhl Seminar 18482

18482 – Network Visualization in the Humanities | Dagstuhl Seminar | <https://www.dagstuhl.de/18482>

As an organizing committee of the seminar we would like to thank the scientific and administration staff of Schloss Dagstuhl for the excellent support they provided, both in the preparation phase and during the seminar. On behalf of all participants, we would also like to thank Dagstuhl for the high quality facilities provided, for excellent rooms for work and socializing, for the tasty meals, and of course also for the excellent wine cellar.

■ Resonanz in Sozialen Netzwerken

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

■ Feedback in Social Media

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

Lewis Chuang (LMU München, DE)

18252 – Ubiquitous Gaze Sensing and Interaction | Dagstuhl Seminar | <https://twitter.com/LewisChuang/status/1009666814266298370>

Can't believe #ugsi got so much done in only 2 days. And still found the time to play games. @dagstuhl brings out the best in us. Thanks!

Blair D. Sullivan (North Carolina State University, Raleigh, US)

18241 – High-Performance Graph Algorithms | Dagstuhl Seminar | <https://twitter.com/BlairDSullivan/status/1015308093931089922>

Just wanted to post a (belated) thanks to @dagstuhl for enabling me to attend with my toddler (w/ a shout-out to @MooreFound #MooreData for supporting caregiver travel). Can't say enough nice things about the Dagstuhl facilities/staff. #TCSwithatoddler #womenincomputing @cscncsu

Juan F. Sequeda (Capsenta Inc., Austin, US)

18371 – Knowledge Graphs | Dagstuhl Seminar | <https://twitter.com/juansequeda/status/1040879998725443585>

Stayed an extra night at @dagstuhl to geek out at the library and process all the conversations and ideas I got through the week. I love this place so much. Luckily I get to come back in a year.

Martin Roetteler (Microsoft Corporation, Redmond, US)

18381 – Quantum Programming Languages | Dagstuhl Seminar | <https://twitter.com/MartinQuantum/status/1043756282811944960>

Amazing workshop! I've been to many @Dagstuhl seminars, but have to really think hard if I ever had such a productive week! Felt proud to be part of this group. [...]

@Uri_Hasson

Dagstuhl Seminar | https://twitter.com/Uri_Hasson/status/1063460725883117569

Very impressed by the childcare solutions provided by @dagstuhl (Leibniz Center for Informatics; Schloss Dagstuhl Germany) for its seminar attendees. Really raises the bar. <https://bit.ly/2DoQemo>

■ Resonanz im Fragebogen

Jeder Teilnehmer erhält von uns einen Fragebogen zur Evaluation des vom Teilnehmer besuchten Dagstuhl-Seminars oder Dagstuhl-Perspektiven-Workshops. Durch diese anonymen Befragung erhalten wir ebenfalls eine Menge positiver Kommentare. Im Folgenden zitieren wir hier einige von diesen.

■ Survey Feedback

Every participant has the opportunity to fill out a questionnaire about the Dagstuhl Seminar or Dagstuhl Perspectives Workshop they attended for evaluation purposes. Below are some excerpts from the many positive comments we received through this anonymous survey.

18041 – Foundations of Data Visualization | Dagstuhl Seminar | <https://www.dagstuhl.de/18041>

Dagstuhl seminars are incredible place to use. Deep respect for the staff to make it possible.

18071 – Planning and Operations Research | Dagstuhl Seminar | <https://www.dagstuhl.de/18071>

Continue to provide a secluded space for researchers to naturally interact. The strength of the location stems from the fact that all activities are done in a limited space so everyone has a chance to mingle / interact. This is something a normal conference venue does not offer.

18081 – Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization | Dagstuhl Seminar | <https://www.dagstuhl.de/18081>

The best: friendly and very comfortable work environment, first-class library facilities.

18111 – Loop Optimization | Dagstuhl Seminar | <https://www.dagstuhl.de/18111>

I always appreciate more vegetarian protein options, but in general I felt the cooking staff put in much appreciated effort to accommodate everyone. As someone with dietary restrictions not very well served by average European food, I felt much better served than at almost all other retreat-style settings I have experienced.

18151 – Program Equivalence | Dagstuhl Seminar | <https://www.dagstuhl.de/18151>

Dagstuhl is an outstanding venue for furthering computing science research. It has always been a high point of my academic life to attend Dagstuhl seminars.

18161 – Visualization of Biological Data - Crossroads | Dagstuhl Seminar | <https://www.dagstuhl.de/18161>

Best aspect is the close and highly creative interaction with excellent researchers from within and outside of the field. Worst aspect is that it only lasts a week; I could work very productively in such an environment for the rest of my life!

18161 – Visualization of Biological Data - Crossroads | Dagstuhl Seminar | <https://www.dagstuhl.de/18161>

The staff was very friendly and responsive. They went out of their way to be helpful.

18161 – Visualization of Biological Data - Crossroads | Dagstuhl Seminar | <https://www.dagstuhl.de/18161>

The venue is outstanding and provides a wonderful environment for creative thinking.

18172 – Algebraic Effect Handlers go Mainstream | Dagstuhl Seminar | <https://www.dagstuhl.de/18172>

I hope you can continue to maintain the excellent library. It's always nice to see the collection of books authored or edited by participants.

18172 – Algebraic Effect Handlers go Mainstream | Dagstuhl Seminar | <https://www.dagstuhl.de/18172>

I think Dagstuhl is amazing. There is no other place where you can spend 5 days and leave so inspired. Actually, it seems to be getting better. The facilities are in top condition and seemed a better than last time. Even the food is better.

18202 – Inter-Vehicular Communication Towards Cooperative Driving | Dagstuhl Seminar | <https://www.dagstuhl.de/18202>

Many areas/rooms seem to be freshly renovated and everything is very clean and tidy. I like this a lot.

18211 – Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance | Dagstuhl Seminar | <https://www.dagstuhl.de/18211>

Thanks so much for keeping Dagstuhl running! It's such an important contribution to the community.

18231 – The Constraint Satisfaction Problem: Complexity and Approximability | Dagstuhl Seminar | <https://www.dagstuhl.de/18231>

One of such a Dagstuhl Seminar is worth more than 10 traditional conferences with respect to effectiveness in spreading ideas.

18241 – High-Performance Graph Algorithms | Dagstuhl Seminar | <https://www.dagstuhl.de/18241>

The facilities and support for bringing a child were exceptional. We especially appreciated the special toddler-friendly meals and availability of toys, crib, and highchair. The family apartment is great though we had a few issues with running out of hot water for showers.

18251 – Database Architectures for Modern Hardware | Dagstuhl Seminar | <https://www.dagstuhl.de/18251>

Thank you for maintaining such a great facility. The staff were incredibly organized, and I felt well served in all aspects.

18251 – Database Architectures for Modern Hardware | Dagstuhl Seminar | <https://www.dagstuhl.de/18251>

The Dagstuhl staff are excellent. Kitchen/dining room staff were unfailingly pleasant and helpful. Staff at reception were extremely flexible, especially in accommodating Friday departure plans.

18252 – Ubiquitous Gaze Sensing and Interaction | Dagstuhl Seminar | <https://www.dagstuhl.de/18252>

This has been the best use of my time in a conference/seminar setting in the past decade. I feel very fortunate to have participated in this.

18252 – Ubiquitous Gaze Sensing and Interaction | Dagstuhl Seminar | <https://www.dagstuhl.de/18252>

Dagstuhl provides a great opportunity for making new collaborations and learning from experts and for mentoring young researchers.

18261 – Discipline Convergence in Networked Systems | Dagstuhl Seminar | <https://www.dagstuhl.de/18261>

The basic tenets of Dagstuhl - moderate isolation and face-to-face discussion, are arguably more important today in the age of social media than they were when the center was set up. It is a special place.

18271 – In Situ Visualization for Computational Science | Dagstuhl Seminar | <https://www.dagstuhl.de/18271>

Wow, I think Dagstuhl is a model for producing quality R&D amongst communities! I hope other communities and countries see what Dagstuhl has done and use it as a model to produce and support robust communities of the various science.

18351 – Modeling for Sustainability | Dagstuhl Seminar | <https://www.dagstuhl.de/18351>

A detail: It is nice that extra consumption (drinks etc.) is based on trust, as if we would all share a minibar. This not only saves bureaucracy, it also has a positive effect on the overall atmosphere.

18411 – Progressive Data Analysis and Visualization | Dagstuhl Seminar | <https://www.dagstuhl.de/18411>

The food is just amazing, in particular the variety of dishes the kitchen prepares!

18441 – Data Physicalization | Dagstuhl Seminar | <https://www.dagstuhl.de/18441>

While I have been attending academic events since 1995, this has been perhaps the most impactful research event of my career. It was a wonderful, transformative event, which will not only substantially evolve some of my own perspectives and approaches, but also significantly enhance the prospects for success and impact by at least 2 (and perhaps 3) of my present doctoral students.

18441 – Data Physicalization | Dagstuhl Seminar | <https://www.dagstuhl.de/18441>

I have never been in a conference venue that was so well prepared, so well-equipped, and so seamlessly run. It meant that, as researchers, all we had to worry about was making connections and getting work done. I actually truly enjoyed and appreciated the fact that dinner was (a) at a specific time (no worrying about when anyone was going to choose to have dinner) (b) assigned seating (no worrying about who to sit next to) and (c) one menu for everyone (no worrying about whether you could find something to eat, or if you would like it – all of the food was lovely!)

18441 – Data Physicalization | Dagstuhl Seminar | <https://www.dagstuhl.de/18441>

I loved the music room – very much so appreciate having it there!

18441 – Data Physicalization | Dagstuhl Seminar | <https://www.dagstuhl.de/18441>

I did not attempt access to the electronic library services. However, I was quite impressed with the physical library services, and especially the tradition of acquiring books by session participants. I thought this was an outstanding idea; surely including the signature dimension. I am now finishing a book myself (with colleagues including an organizer of another upcoming Dagstuhl, to which I have agreed), and will surely endeavor to send a signed copy as soon as it exists.

18462 – Provenance and Logging for Sense Making | Dagstuhl Seminar | <https://www.dagstuhl.de/18462>

Two junior (PhD student) researchers at this seminar remarked to me about how much they appreciated being invited, and how they expect it to advance their careers because they've been able to network all week with senior members of our research community, an opportunity they otherwise wouldn't have had.

18462 – Provenance and Logging for Sense Making | Dagstuhl Seminar | <https://www.dagstuhl.de/18462>

I really appreciate what Schloss Dagstuhl does for the Computer Science community, and particularly the visualization community. It is extremely helpful to have a venue like this where a small group can come together for an extended period and define a research agenda. The previous seminar I attended (Data Storytelling) resulted in a book, plus many new contacts and research collaborations. It gave me an entirely new (much broader) perspective on the topic of data storytelling and helped me identify several worthy research directions. I am hopeful that the current seminar will have similar positive outcomes.

18511 – Algebraic Coding Theory for Networks, Storage, and Security | Dagstuhl Seminar | <https://www.dagstuhl.de/18511>

Best: free and open atmosphere.

18511 – Algebraic Coding Theory for Networks, Storage, and Security | Dagstuhl Seminar | <https://www.dagstuhl.de/18511>

Many thanks to the permanent personal of Dagstuhl for their work, it is always a pleasure to come here, everything is very well organised.

Resonanz zur Bibliographiedatenbank dblp

5.2

Die Bibliographiedatenbank dblp wird von zahlreichen internationalen Wissenschaftlern hoch geschätzt und erhält viel Lob. Feedback erhalten wir per Mail, durch Gespräche mit Forschern vor Ort in Dagstuhl, oder durch die sozialen Medien.

Feedback on the dblp Computer Science Bibliography

5

The dblp computer science bibliography is internationally well known and appreciated. We receive a lot of feedback via mail, through discussions with researchers at Schloss Dagstuhl, and via social media.

Jinseok Kim (University of Michigan, Ann Arbor, MI, USA)

„Evaluating author name disambiguation for digital libraries: a case of DBLP“, *Scientometrics* 116(3): 1867-1886 (2018) | <https://doi.org/10.1007/s11192-018-2824-5>

In conclusion, the evaluation results reported in this paper suggest that scholars can regard DBLP data as highly accurate in disambiguating author names. But a caveat to keep in mind is that some homonym cases (distinct authors with the same names) may not be properly distinguished.

Andreas Halkjær From (Technical University of Denmark, Copenhagen)

Twitter | <https://twitter.com/andreasfrom/status/971447624455901185>

I'm not very excited about the article but I'm quite excited about being on dblp now

Mark J. Nelson (American University, Washington, DC, USA)

Twitter | https://twitter.com/mm_jj_nn/status/979508034484494336

Thing I like abt @dblp_org vs Google Scholar (besides open data & diligently curated): easy to see *where* someone normally publishes & filter by venue. "Which confs do they publish at?" is often a good proxy for research areas in comp sci.

Anish Singh Shekhawat (San José State University, CA, USA)

Twitter | https://twitter.com/anish_shekhawat/status/994810026354200576

+1 for DBLP. Has been a huge help throughout my graduate studies.

David Maus (Herzog August Library, Wolfenbüttel, Germany)

Twitter | https://twitter.com/_dmaus/status/1043020348902391809

Just discovered the "computer science bibliography" (@dblp_org) provided by @dagstuhl and @TrierUni. Overwhelming.

Mario Gleirscher (University of York, UK)

Twitter | <https://twitter.com/MarioGleirscher/status/1070711361880145921>

Thank you #DBLP for extraordinarily comprehensive, accurate, conveniently searchable #bibliographical data, having been so useful to me over the years... great to hear that DBLP will be receiving additional financial support!

Johannes Hölzl (VU Amsterdam, The Netherlands)

Twitter | <https://twitter.com/johannes2007/status/1070780368092045313>

@dblp_org is one of the most important web tools I use for research!

Michael Marek (Tübingen, Germany)

Twitter | <https://twitter.com/michaelmarek10/status/1076435778023157761>

Super. Keine lange Liste im Vergleich mit wirklich aktiven WissenschaftlerInnen ;-)) aber nachdem mein Namensvetter und ich sortiert wurden, passt die Liste bei der »dblp: computer science bibliography« Vielen Dank an das Team von @dblp_org!³³

³³ engl.: Awesome. Not a long list compared to really active scientists ;-)) but after my namesake and I have been sorted, the list at the "dblp: computer science bibliography" is correct. Many thanks to the @dblp_org team!



Fig. 5.1
Visit of the Saarland Minister President Tobias Hans during his summer tour 2018.
Photo courtesy of Saarland/mn.

6

Die Seminare in 2018

The 2018 Seminars

■ Applications, Interdisciplinary Work

- 10 Years of Web Science: Closing The Loop (18262)
- Data Physicalization (18441)
- Formal Methods for the Synthesis of Biomolecular Circuits (18082)
- Genomics, Pattern Avoidance, and Statistical Mechanics (18451)
- Human-Computer Integration (18322)
- Modeling for Sustainability (18351)
- On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces (18212)
- Progressive Data Analysis and Visualization (18411)
- Provenance and Logging for Sense Making (18462)
- Ubiquitous Gaze Sensing and Interaction (18252)
- Visualization of Biological Data – Crossroads (18161)
- Towards Accountable Systems (18181)

■ Artificial Intelligence, Computational Linguistics

- Coding Theory for Inference, Learning and Optimization (18112)
- Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (18371)
- Machine Learning and Model Checking Join Forces (18121)
- Normative Multi-Agent Systems (18171)
- Planning and Operations Research (18071)

■ Cryptography, Security, Privacy

- Blockchain Security at Scale (18461)
- Blockchain Technology for Collaborative Information Systems (18332)
- Blockchains, Smart Contracts and Future Applications (18152)
- Secure Routing for the Internet (18242)
- Symmetric Cryptography (18021)
- Web Application Security (18321)

■ Databases, Information Retrieval, Machine Learning, Data Mining

- Automating Data Science (18401)
- Data Consistency in Distributed Systems: Algorithms, Programs, and Databases (18091)
- Database Architectures for Modern Hardware (18251)
- Implementing FAIR Data Infrastructures (18472)
- Multidirectional Transformations and Synchronisations (18491)

■ Data Structures, Algorithms, Complexity

- Algebraic Methods in Computational Complexity (18391)
- Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative (18421)
- Algorithmic Foundations of Programmable Matter (18331)
- Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization (18081)
- Dynamic Traffic Models in Transportation Science (18102)
- High-Performance Graph Algorithms (18241)
- High Throughput Connectomics (18481)
- Measuring the Complexity of Computational Content: From Combinatorial Problems to Analysis (18361)
- Personalized Multiobjective Optimization: An Analytics Perspective (18031)
- Proof Complexity (18051)
- Scheduling (18101)
- Synergies between Adaptive Analysis of Algorithms, Parameterized Complexity, Compressed Data Structures and Compressed Indices (18281)
- The Constraint Satisfaction Problem: Complexity and Approximability (18231)

■ Distributed Computation, Networks, Architecture, Systems

- Algebraic Coding Theory for Networks, Storage, and Security (18511)
- Discipline Convergence in Networked Systems (18261)
- Encouraging Reproducibility in Scientific Research of the Internet (18412)
- Inter-Vehicular Communication Towards Cooperative Driving (18202)

■ Geometry, Image Processing, Graphics, Visualization

- Computational Aspects of Fabrication (18431)
- Extreme Classification (18291)
- Foundations of Data Visualization (18041)
- In Situ Visualization for Computational Science (18271)
- Network Visualization in the Humanities (18482)
- Shape Analysis: Euclidean, Discrete and Algebraic Geometric Methods (18422)
- Visualization and Processing of Anisotropy in Imaging, Geometry, and Astronomy (18442)

■ Software Technology, Programming Languages

- Algebraic Effect Handlers go Mainstream (18172)
- Automatic Quality Assurance and Release (18122)
- Evidence About Programmers for Programming Language Design (18061)
- Genetic Improvement of Software (18052)
- Loop Optimization (18111)
- Quantum Programming Languages (18381)
- Secure Compilation (18201)
- Software Business, Platforms, and Ecosystems: Fundamentals of Software Production Research (18182)

■ Verification, Logic, Formal Methods, Semantics

- Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance (18211)
- Formalization of Mathematics in Type Theory (18341)
- Next Generation Domain Specific Conceptual Modeling: Principles and Methods (18471)
- Program Equivalence (18151)
- The Logical Execution Time Paradigm: New Perspectives for Multicore Systems (18092)

6.1 Symmetric Cryptography

Organizers: Joan Daemen, Tetsu Iwata, Nils Gregor Leander, and Kaisa Nyberg

Seminar No. 18021

Date: January 7–12, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.1.1

© Creative Commons BY 3.0 Unported license

© Nils Gregor Leander, Joan Daemen, Tetsu Iwata, and Kaisa Nyberg



Participants: Frederik Armknecht, Tomer Ashur, Christof Beierle, Daniel J. Bernstein, Eli Biham, Alex Biryukov, Anne Canteaut, Joan Daemen, Itai Dinur, Christoph Dobraunig, Orr Dunkelman, Maria Eichlseder, Henri Gilbert, Tetsu Iwata, J r my Jean, Dmitry Khovratovich, Stefan K lbl, Virginie Lallemand, Tanja Lange, Nils Gregor Leander, Ga tan Leurent, Stefan Lucks, Willi Meier, Bart Mennink, Vasily Mikhalev, Kazuhiko Minematsu, Nicky Mouha, Mridul Nandi, Maria Naya-Plasencia, Kaisa Nyberg, Stav Perle, L o Paul Perrin, Thomas Peyrin, Christian Rechberger, Arnab Roy, Yu Sasaki, Yannick Seurin, Adi Shamir, Marc Stevens, Stefano Tessaro, Yosuke Todo, Gilles Van Assche, Damian Viz r, Meiqin Wang, Kan Yasuda

IT Security plays an increasingly vital role in everyday life and business. When talking on a mobile phone, when withdrawing money from an ATM or when buying goods over the internet, security plays a crucial role in both protecting the user and in maintaining public confidence in the system. Especially after the disclosure of the NSA’s world-spanning spying activities and in the context of the Internet of Things, IT Security and privacy protection is a vital topic of the 21st century. In the Internet of Things (IoT) era, everything will be connected. Intel estimates that 200 billion objects will be connected by 2020. The objects include for instance smart devices for healthcare, industrial control systems, automotive, and smart homes. Virtually all modern security solutions rely on cryptography.

Symmetric cryptography deals with the case that both the sender and the receiver of a message are using the same key. This differentiates symmetric cryptography from its asymmetric counterpart, where senders or verifiers use a “public key” and receivers or signers use a corresponding but different “private key”. As asymmetric primitives are typically orders of magnitude less efficient than symmetric cryptographic schemes, symmetric cryptosystems remain the main workhorses of cryptography and highly relevant not only for academia, but also for industrial research and applications. While great progress has been made in designing and analyzing ciphers, fundamental aspects of these ciphers are still not fully understood. Moreover, as we have learned from the Snowden revelations, cryptography in general and symmetric cryptography in particular faces new fascinating challenges.

Current Topics and Challenges We identified the following three areas as among the most important topics for future research.

Cryptography for the IoT. Motivated by the upcoming IoT, one of the strong research trends in symmetric cryptography

is about lightweight cryptography. Here, lightweight cryptography refers to strong cryptography, that can be executed on heavily resource constrained devices. Those efforts resulted in a wide variety of block cipher designs suitable for IoT applications. For instance, PRESENT designed in 2007 is one of the early designs with strong implementation advantages on hardware, and there have been other innovative follow-up block cipher designs. Some of them are standardized as the international standard, and used in thousands of devices in our daily lives. However, a block cipher is not the solution to all cryptographic purposes. For instance, to encrypt a certain amount of data, the block cipher has to be integrated into a suitable mode of operation. In most practical use cases, confidentiality is not the only concern, as many scenarios require data authenticity as well. Here a message authentication code (MAC) can be used to ensure authenticity. Authenticated encryption (AE) is used for protecting both confidentiality and authenticity.

The first MAC, called Chaskey, that specifically targets applications for lightweight cryptography was proposed only recently in 2014. The CAESAR project, an international competition for AE initiated at Dagstuhl, attracted several submissions that were designed for the purposes for lightweight cryptography. There is also a recent attempt to design a lightweight tweakable block cipher, an advanced primitive of a block cipher that allows more flexible usage, which can be efficiently integrated into highly secure encryption and/or authentication mechanisms. However, this research just started and many primitives and modes of operations suitable for lightweight crypto remain to be explored.

Statistical Attacks. Statistical attacks have been deployed widely and providing strong resistance against them has resulted in several important design criteria for contemporary symmetric primitives. The first type of statistical attacks that is applicable to a large set of block ciphers is differential cryptanalysis, introduced by Biham and Shamir. Since its

invention in the early nineties several variants, tweaks and generalizations have been proposed and applied to many block ciphers. The second generally applicable attack on block ciphers is Matsui's linear cryptanalysis. Similarly to differential attacks, since its introduction, many extensions and improvements have been made. One main issue that has become apparent only recently is the accuracy of the underlying statistical models that researchers are using. Typically, those models are presented under some simplifying assumptions, whose validity remains an open question. It is an important challenge to settle these unsatisfactory simplifications. This becomes even more important when the attacks are hard or impossible to verify experimentally due to the large computational costs involved. Moreover, to allow comparison between different attacks the researchers must agree on common attack models and parameters that measure the performance of the attack.

Symmetric Cryptography and Real-World Needs.

The symmetric cryptography community has many very talented people and the state of the area has moved from its infancy in the seventies to a mature field today. However, we should ensure that the world's population does benefit of this progress. In particular, the Snowden leaks have painfully illustrated that citizen privacy and anonymity is next to non-existent nowadays. Secret services and IT corporations massively spy on people's communication and data storage for motives such as profit and surveillance. They don't seem to be hindered significantly in this at all by the pervasive deployment of cryptography (TLS, GSM, WPA, etc.). Cynically, monopolistic corporations like Google use encryption to protect the data of their users from prying eyes of other players such as network providers. It appears that much of the cryptography deployed today is there to protect the powers that be rather than protect human rights. With the roll-out of smart grid and internet-of-things surveillance will become quasi universal with all imaginable devices reporting on our behavior to big corporations. This situation has been addressed in several invited talks by Bart Preneel and Adi Shamir and they rightfully say that we as a cryptographic community should attempt to improve this. Along the same lines, Phil Rogaway gave a highly acclaimed invited talk at Asiacrypt 2015 on the moral aspects on cryptographic research. He invites us to do some introspection and ask the question: are we doing the right thing?

We believe these questions are important also for the symmetric crypto community. While the problem is certainly not restricted to symmetric cryptography and probably cannot be solved by symmetric cryptography alone, we should consider it our moral duty to improve the situation.

Seminar Program The seminar program consists of presentations about the above topics, and relevant areas of symmetric cryptography, including new cryptanalytic techniques and new designs. Furthermore, there were discussion sessions. In "Discussion on CAESAR with focus on robustness", we discussed about the meaning and relevance of the term robustness in general and for the CAESAR competition in particular. In "Discussion on Mass Surveillance", a number of questions related to the real-world relevance of the symmetric crypto community and its research were discussed. For both discussions we provide summary of the questions and results.

6.2 Personalized Multiobjective Optimization: An Analytics Perspective

Organizers: Kathrin Klamroth, Joshua D. Knowles, Günter Rudolph, and Margaret M. Wiecek
Seminar No. 18031

Date: January 14–19, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.1.33

© Creative Commons BY 3.0 Unported license

© Kathrin Klamroth, Joshua D. Knowles, Günter Rudolph, and Margaret M. Wiecek



Participants: Richard Allmendinger, Mickaël Binois, Jürgen Branke, Dimo Brockhoff, Roberto Calandra, Carlos A. Coello Coello, Kerstin Dächert, Kalyanmoy Deb, Matthias Ehrgott, Gabriele Eichfelder, Michael Emmerich, Alexander Engau, Georges Fadel, José Rui Figueira, Carlos M. Fonseca, Abhinav Gaur, Salvatore Greco, Jussi Hakanen, Johannes Jahn, Andrzej Jaszkievicz, Milosz Kadzinski, Kathrin Klamroth, Karl Heinz Küfer, Christoph Lofi, Manuel López-Ibáñez, Kaisa Miettinen, Sanaz Mostaghim, Boris Naujoks, Frank Neumann, Luís Paquete, Robin Purshouse, Patrick M. Reed, Günter Rudolph, Stefan Ruzika, Serpil Sayin, Pradyumn Kumar Shukla, Roman Slowinski, Ralph E. Steuer, Theodor J. Stewart, Michael Stiglmayr, Lothar Thiele, Selvakumar Ulaganathan, Daniel Vanderpooten, Margaret M. Wiecek

The topic of the seminar, Personalization in Multiobjective Optimization, was motivated by ongoing changes in many areas of human activity. In particular, personalization, mass customization, and mass data have become essential in current business and engineering operations creating new challenges for academic and research communities. In the seminar, the EMO and MCDM communities, including junior and senior academic researchers as well as industry representatives, took an effort to jointly address the ongoing changes in the real-world with multiobjective optimization.

The purpose of multiobjective optimization is to develop methods that can solve problems having a number of (conflicting) optimization criteria and constraints, providing a multitude of solution alternatives, rather than pursuing only one “optimal” solution. In this aim the field has been highly successful: its methods have a track record of improving decision making across a broad swath of applications, indeed wherever there are conflicting goals or objectives. Yet, multiobjective optimization has so far focused almost exclusively on serving a single “decision maker”, providing solutions merely as potential (not actual) alternatives. In order to fulfill the demanding aims of mass-customization, product/service variation and personalization we see today in areas such as engineering, planning, operations, investment, media and Web services, and healthcare, new and innovative approaches are needed. This seminar took the first steps towards this goal by bringing together leading specialists in EMO and MCDM.

Personalization in multiobjective optimization as the main theme of the seminar has focused around three **application challenges** which are highly characteristic for real-world decision making and represent different ways that personalization is needed or delivered in an optimization setting. These were (i) Platform design and product lines, (ii) Responsive and online personalization, and (iii) Complex networks of decision makers.

These three application challenges were crosslinked with three **research domains** that constitute the methodological core of multiobjective optimization and have been the foundation for the discussions at the previous Dagstuhl seminars. These were (1) Model building, (2) Preference modelling, and (3) Algorithm design and efficiency.

During the seminar, we formed five multi-disciplinary working groups (WGs) to implement the crosslinking between these application challenges and research domains, see Table 6.1. Each working group was focused on an application challenge (a row in Table 6.1; WGs 2, 3 and 4) or a research domain (a column in Table 6.1; WGs 1 and 5), all taking specific perspectives on the respective topics.

Table 6.1

Working groups (WGs) crosslinking application challenges (rows) with research domains (columns). WG 1: Preference uncertainty quantification; WG 2: Personalization and customization of decision support; WG 3: Invariant rule extraction; WG 4: Complex networks and MCDA; WG 5: Metamodelling for interactive optimization.

	Modelling	Preferences	Algorithms
Platform design and product lines	WG3, WG5	WG1, WG3	WG3
Responsive and online personalization	WG2, WG5	WG1, WG2	WG2
Complex networks of decision makers	WG4, WG5	WG1, WG4	WG4

The program was updated on a daily basis to maintain flexibility in balancing time slots for talks, discussions, and working groups. The working groups were established on the first day in an open and highly interactive discussion. The program included several opportunities to report back from the working groups in order to establish further links and allow for adaptations and feedback. Some of the working groups split into subgroups and rejoined later in order to focus more strongly on

different aspects of the topics considered. Abstracts of the talks and extended abstracts of the working groups can be found in subsequent chapters of this report. Further notable events during the week included: (i) a hike on Wednesday afternoon with some sunshine (despite the quite terrible weather during the rest of the week), (ii) an announcements session allowing us to share details of upcoming events in our research community, and (iii) a wine and cheese party made possible by the support of the ITWM Kaiserslautern, represented by Karl-Heinz Küfer.

Outcomes Fourteen topical presentations were complemented by discussions in five working groups, covering the main themes of the seminar. The outcomes of each of the working groups can be seen in the sequel. Extended versions of their findings will be submitted to a Special Issue on “Personalization in Multiobjective Optimization: An Analytics Perspective” of the *Journal of Multicriteria Decision Analysis*, edited by Theo Stewart, that is guest edited by the organizers of this seminar. The submission deadline is July 31, 2018, and several working groups plan to submit extended versions of their reports to this special issue.

The seminar was highly productive, very lively and full of discussions, and has thus further strengthened the interaction between the EMO and MCDM communities. We expect that the seminar will initiate a new research domain interrelating multiobjective optimization and personalization, as it similarly has happened after the previous seminars in this series.

Acknowledgments A huge thank you to the Dagstuhl office and its very helpful and patient staff; many thanks to the organizers of the previous seminars in the series for the initiative and continuing advice; and many thanks to all the participants, who contributed in so many different ways to make this week a success. In the appendix, we also give special thanks to Joshua Knowles as he steps down from the organizer role.

6.3 Foundations of Data Visualization

Organizers: Helwig Hauser, Penny Rheingans, and Gerik Scheuermann

Seminar No. 18041

Date: January 21–26, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.1.100

© Creative Commons BY 3.0 Unported license

© Penny Rheingans



Participants: James Ahrens, Johanna Beyer, Michael Böttinger, Stefan Bruckner, Roxana Bujack, Hamish Carr, Min Chen, Leila De Florian, Christoph Garth, Eduard Gröller, Hans Hagen, Charles D. Hansen, Helwig Hauser, Hans-Christian Hege, Nathalie Henry Riche, Mario Hlawitschka, Ingrid Hotz, Christopher R. Johnson, Alark Joshi, Gordon Kindlmann, Helen-Nicole Kostis, Barbara Kozlíková, David H. Laidlaw, Heike Leitte, Ross Maciejewski, Georgeta Elisabeta Marai, Kresimir Matkovic, Laura A. McNamara, Silvia Miksch, Torsten Möller, Daniela Oelke, Kristi Potter, Bernhard Preim, Penny Rheingans, Gerik Scheuermann, Marc Streit, Holger Theisel, Jarke J. van Wijk, Amitabh Varshney, Maria Velez-Rojas, Anna Vilanova, Ivan Viola, Daniel Weiskopf, Ross Whitaker, Thomas Wischgoll, Anders Ynnerman, Caroline Ziemkiewicz

Data Visualization is the transformation of data, derived from observation or simulation, and models into interactive images. It has become an indispensable part of the knowledge discovery process in many fields of contemporary endeavor. Since its inception about three decades ago, the techniques of data visualization have aided scientists, engineers, medical practitioners, analysts, and others in the study of a wide variety of data, including numerical simulation based on high-performance computing, measured data from modern scanners (CT, MR, seismic imaging, satellite imaging), and survey and sampled data, and metadata about data confidence or provenance. One of the powerful strengths of data visualization is the effective and efficient utilization of the broad bandwidth of the human sensory system in interpreting and steering complex processes involving spatiotemporal data across a diverse set of application disciplines. Since vision dominates our sensory input, strong efforts have been made to bring the mathematical abstraction and modeling to our eyes through the mediation of computer graphics. The interplay between these multidisciplinary foundations of visualization and currently emerging, new research challenges in data visualization constitute the basis of this seminar.

The rapid advances in data visualization have resulted in a large collection of visual designs, algorithms, software tools, and development kits. There is also a substantial body of work on mathematical approaches in visualizations such as topological methods, feature extraction approaches, and information theoretical considerations. However, a unified description of theoretical and perceptual aspects of visualization would allow visualization practitioners to derive even better solutions using a sound theoretical basis. There are promising ideas but they need further discussion. Currently, we employ user studies to decide if a visual design is more effective, but a comprehensive theory would allow visualization researchers to answer why one visual design is more effective than another and how the visual

design can be optimized. Furthermore, we usually have an understanding of the role of a specific visualization in a specific analytic workflow, but we would like to formalize the general role of visualization in the analytic workflow. This would also allow for more quantitative measures of visualization quality. In addition, the community needs a deeper, general understanding of the most informative way to conduct perceptual and usability studies involving domain experts.

For this seminar, we chose to take a focused consideration of the foundations of visualization in order to establish an integrated discussion on the fundamental understanding and generic methodologies of data visualization, including theories, models and workflows of data visualization, evaluation metrics, and perceptual and usability studies. We included experts from all areas of visualization such as scientific visualization, information visualization, and visual analytics to allow for an in-depth discussion of our shared research foundations based on a broad expertise.

With the experience of delivering technical advances over the past three decades, it is timely for the visualization community to address these fundamental questions with a concerted effort. Such an effort will be critical to the long-term development of the subject, especially in building a theoretical foundation for the subject. The community needs to develop suitable models for the whole visualization process from cleaning and filtering the data, analysis processing, mapping to graphical scenes, to the interpretation by the human visual system. While there are some methods of evaluation based on user studies and findings in applications, a complete theoretical foundation for evaluations is missing. Modern visualization includes advanced numerical and combinatorial data processing, so the correctness of this processing including a critical look at its assumptions with respect to the application at hand is needed. Only then, visualization can establish strong correlations between visualization algorithms

and questions in the application domains. In addition, uncertainty has received attention from the visualization community in recent years, but a full analysis of uncertainty at all stages of the established visualization pipeline is still not available. Theoretical foundations of uncertainty in visualization need to look at uncertainty in the data, errors due to numerical processing, errors due to visual depiction and, finally, uncertainty in the results based on human misinterpretation of interactive visual depictions.

This workshop addressed five important topics:

Theory of overall visualization process. A theory of the whole visualization process needs to cover all parts of the visualization pipeline and should be applicable to broad classes of application domains. Of course, it is the ultimate foundation, but there are a few formulation attempts and the seminar discussed them. Such a theory should allow to find optimal visualizations and to quantify the value of visualizations. In addition, it is strongly believed by most experts that such a theory needs to cover the challenge of uncertainty in the data, the processing including visual mapping and potential misinterpretation by human observers.

Foundations of evaluation. Evaluation allows designers and analysts to select visualization approaches from among different options for a specific problem. One evaluation method is a user study, usually with a larger group of subjects. Here, it is often a challenge that there is only a very small set of experts available that understand the scientific questions behind the data. Guidelines for user study design in these situations are necessary. In addition, evaluation needs to look at limits of the human visual system. In advanced analytic applications, it is also very important to study the relation between user interest and visualization. There are many open questions in this area that will be discussed in the seminar.

Collaboration with domain experts. Many visualizations address questions and needs from expert researchers, engineers, analysts, or decision makers. Therefore, visualization nearly always involves people outside the visualization community. The seminar included some representatives from large applied research centers so that the discussion about relations between visual data analysis and application semantics was not carried out without domain experts. These participants also commented on methodologies for defining domain requirements and realistic roles of application researchers in evaluation.

Visualization for broad audiences. Visualizations developed for broad audiences involve context and constraints different from those developed for expert domain collaborators. Such visualizations include those for personal information, school use, science centers and other public settings, and communication with a broad general public. Issues with developing visualizations for broad audiences include a higher need for intuitive metaphors and conventions, a larger imperative for drawing participants into interaction, and more requirements for robust interfaces and systems.

Mathematical foundations of visual data analysis. There is a rich tradition of mathematical/computational methods used in visualization, such as topological approaches, mathematical descriptions of feature extraction, numerical sampling and reconstruction methods, integration, differential operators, filtering, dimension reduction, and applications of information theory. In addition, we have seen promising attempts to incorporate uncertainty in these mathematical approaches. While all these methods have a solid mathematical foundation, a careful look at the relation between theories in applications and these mathematical approaches in visual data analysis was taken in this seminar.

The format of the seminar incorporated several elements: overview talks on each topic, clusters of short talks on a single topic followed by a joint panel discussion, and breakout groups

on each of the five topics. Unlike the typical arrangement, all presentations in each session were given in sequence without a short Q&A session at the end of each talk. Instead, all speakers of a session were invited to sit on the stage after the presentations, and answer questions in a manner similar to panel discussions. This format successfully brought senior and junior researchers onto the same platform, and enabled researchers to seek a generic and deep understanding through their questions and answers. It also stimulated very long, intense, and fruitful discussions that were embraced by all participants. The breakout groups focused on the general themes and are reported in a later section.

6.4 Proof Complexity

Organizers: Albert Atserias, Jakob Nordström, Pavel Pudlák, and Rahul Santhanam
Seminar No. 18051

Date: January 28–February 2, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.1.124

© Creative Commons BY 3.0 Unported license

© Albert Atserias, Jakob Nordström, Pavel Pudlák, and Rahul Santhanam



Participants: Amirhossein Akbar Tabatabaei, Albert Atserias, Paul Beame, Arnold Beckmann, Olaf Beyersdorff, Ilario Bonacina, Igor Carboni Oliveira, Marco Carmosino, Leroy Chew, Stefan Dantchev, Yuval Filmus, Noah Fleming, Michael A. Forbes, Nicola Galesi, Michal Garlik, Joshua A. Grochow, Tuomas Hakoniemi, Johan Hastad, Edward A. Hirsch, Pavel Hrubes, Dmitry Itsykson, Emil Jerabek, Jan Johannsen, Raheleh Jalali Keshavarz, Leszek Kolodziejczyk, Antonina Kolokolova, Oliver Kullmann, Massimo Lauria, Meena Mahajan, Barnaby Martin, Moritz Müller, Jakob Nordström, Joanna Ochremiak, Jan Pich, Aaron Potechin, Pavel Pudlák, Ninad Rajgopal, Kilian Risse, Robert Robere, Rahul Santhanam, Dmitry Sokolov, Neil Thapen, Jacobo Torán, Iddo Zameret, Alasdair Urquhart, Marc Vinyals

This workshop brought together the whole proof complexity community spanning from Frege proof systems and circuit-inspired lower bounds via geometric and algebraic proof systems all the way to bounded arithmetic. In this executive summary, we first give an overview of proof complexity, and then describe the goals of the seminar week. Finally, we discuss the relation to previous workshops and conferences.

■ Topic of the Seminar

Ever since the groundbreaking NP-completeness paper of Cook [18], the problem of deciding whether a given propositional logic formula is satisfiable or not has been on centre stage in theoretical computer science. During the last two decades, SATISFIABILITY has also developed from a problem of mainly theoretical interest into a practical approach for solving applied problems. Although all known Boolean satisfiability solvers (SAT solvers) have exponential running time in the worst case, enormous progress in performance has led to satisfiability algorithms becoming a standard tool for solving large-scale problems in, for example, hardware and software verification, artificial intelligence, bioinformatics, operations research, and sometimes even pure mathematics.

The study of proof complexity originated with the seminal paper of Cook and Reckhow [19]. In its most general form, a proof system for a formal language L is a predicate $P(x, \pi)$, computable in time polynomial in the sizes $|x|$ and $|\pi|$ of the input, and having the property that for all $x \in L$ there exists a string π (a *proof*) for which $P(x, \pi)$ evaluates to true, whereas for any $x \notin L$ it should hold for all strings π that $P(x, \pi)$ evaluates to false. A proof system is said to be polynomially bounded if for every $x \in L$ there exists a proof π_x for x that has size at most polynomial in $|x|$. A *propositional proof system* is a proof system for the language of tautologies in propositional logic, i.e., for formulas

that always evaluate to true no matter how the values true and false are assigned to variables in the formula.

From a theoretical point of view, one important motivation for proof complexity is the intimate connection with the fundamental problem of P versus NP. Since NP is exactly the set of languages with polynomially bounded proof systems, and since TAUTOLOGY can be seen to be the dual problem of SATISFIABILITY, we have the famous theorem of [19] that $NP = \text{coNP}$ if and only if there exists a polynomially bounded propositional proof system. Thus, if it could be shown that there are no polynomially bounded proof systems for tautologies, $P \neq NP$ would follow as a corollary since P is closed under complement. One way of approaching this problem is to study stronger and stronger proof systems and try to prove superpolynomial lower bounds on proof size. However, although great progress has been made in the last couple of decades for a variety of proof systems, this goal still appears very distant.

A second theoretical motivation is that simple propositional proof systems provide analogues of subsystems of Peano Arithmetic where the power of mathematical reasoning is restricted. Of particular interest here are various bounded arithmetic systems, which in some sense are intended to capture feasible/polynomial-time reasoning. Proving strong lower bounds on propositional logic encodings of some combinatorial principle, say, in a propositional proof system can in this way show that establishing the validity of this principle requires more powerful mathematics than what is provided by the corresponding subsystem of Peano Arithmetic. One can thus quantify how “deep” different mathematical truths are, as well as shed light on the limits of our (human, rather than automated) proof techniques. At the same time, since it is an empirically verified fact that low-complexity proofs generalize better and are often more constructive, classifying which truths have feasible proofs is also a way to approach the classification of algorithmic problems

by their computational complexity. The precise sense in which this can be formalized into a tool for the complexity theorist is one of the goals of bounded arithmetic.

A third prominent motivation for the study of proof complexity is also algorithmic but of a more practical nature. As was mentioned above, designing efficient algorithms for proving tautologies—or, equivalently, testing satisfiability—is a very important problem not only in the theory of computation but also in applied research and industry. All SAT solvers, regardless of whether they produce a written proof or not, explicitly or implicitly define a system in which proofs are searched for and rules which determine what proofs in this system look like. Proof complexity analyses what it takes to simply write down and verify the proofs that such a solver might find, ignoring the computational effort needed to actually find them. Thus, a lower bound for a proof system tells us that any algorithm, even an optimal (non-deterministic) one magically making all the right choices, must necessarily use at least the amount of a certain resource specified by this bound. In the other direction, theoretical upper bounds on some proof complexity measure give us hope of finding good proof search algorithms with respect to this measure, provided that we can design algorithms that search for proofs in the system in an efficient manner.

The field of proof complexity also has rich connections to algorithmic analysis, combinatorial optimization, cryptography, artificial intelligence, and mathematical logic. A few good sources providing more details are [6, 17, 47].

■ A Very Selective Survey of Proof Complexity

Any propositional logic formula can be converted to a formula in conjunctive normal form (CNF) that is only linearly larger and is unsatisfiable if and only if the original formula is a tautology. Therefore, any sound and complete system that certifies the unsatisfiability of CNF formulas can be considered as a general propositional proof system.

The extensively studied *resolution* proof system, which appeared in [9] and began to be investigated in connection with automated theorem proving in the 1960s [21, 22, 48], is such a system where one derives new disjunctive clauses from an unsatisfiable CNF formula until an explicit contradiction is reached. Despite the apparent simplicity of resolution, the first superpolynomial lower bounds on proof size were obtained only after decades of study in 1985 [33], after which truly exponential size lower bounds soon followed in [15, 52]. It was shown in [8] that these lower bounds can be established by instead studying the *width* of proofs, i.e., the maximal size of clauses in the proofs, and arguing that any resolution proof for a certain formula must contain a large clause. It then follows by a generic argument that any such proof must also consist of very many clauses. Later research has led to a well-developed machinery for showing width lower bounds, and hence also size lower bounds, for resolution.

The more general proof system *polynomial calculus* (PC), introduced in [1, 16],³⁴ instead uses algebraic geometry to reason about SAT. In polynomial calculus clauses are translated to multilinear polynomials over some (fixed) field, and a CNF formula F is shown to be unsatisfiable by proving that there is no common root for the polynomials corresponding to all the clauses, or equivalently that the multiplicative identity 1 lies in the ideal generated by these polynomials. Here the size of a proof is measured as the number of monomials in a proof when all polynomials are expanded out as linear combinations of monomials, and the width of a clause corresponds to the (total)

degree of the polynomial representing the clause. It can be shown that PC is at least as strong as resolution with respect to both size and width/degree, and there are families of formulas for which PC is exponentially stronger.

In the work [36], which served, interestingly enough, as a precursor to [8], it was shown that strong lower bounds on the degree of polynomial calculus proofs are sufficient to establish strong size lower bounds. In contrast to the situation for resolution after [8], however, this has not been followed by a corresponding development of a generally applicable machinery for proving degree lower bounds. For fields of characteristic distinct from 2 it is sometimes possible to obtain lower bounds by doing an affine transformation from $\{0, 1\}$ to the “Fourier basis” $\{-1, +1\}$, an idea that seems to have appeared first in [13, 28]. For fields of arbitrary characteristic a powerful technique for general systems of polynomial equations was developed in [2], which when restricted to CNF formulas F yields that polynomial calculus proofs require high degree if the corresponding clause-variable incidence graphs $G(F)$ are good enough bipartite expander graphs. There are several provably hard formula families for which this criterion fails to apply, however, and even more formulas that are believed to be hard for both resolution and PC, but where lower bounds are only known for the former proof system and not the latter.

Another proof system that has been the focus of much research is *cutting planes* (CP), which was introduced in [20] as a way of formalizing the integer linear programming algorithm in [14, 27]. Here the disjunctive clauses in a CNF formula are translated to linear inequalities, and these linear inequalities are then manipulated to derive a contradiction. Thus, questions about the satisfiability of Boolean formulas are reduced to the geometry of polytopes over the real numbers. Cutting planes is easily seen to be at least as strong as resolution, since a CP proof can mimic any resolution proof line by line. An intriguing fact is that encodings of the *pigeonhole principle*, which are known to be hard to prove for resolution [33] and many other proof systems, are very easy to prove in cutting planes. It follows from this that not only is cutting planes never worse than resolution, but it can be exponentially stronger.

Exponential lower bounds on proof length for cutting planes were first proven in [10] for the restricted subsystem CP*, where all coefficients in the linear inequalities can be at most polynomial in the formula size, and were later extended to general CP in [34, 44]. The proof technique in [44] is very specific, however, in that it works by *interpolating* monotone Boolean circuits for certain problems from CP proofs of related formulas with a very particular structure, and then appealing to lower bounds in circuit complexity. A longstanding open problem is to develop techniques that would apply to other formula families. For example, establishing that randomly sampled k -CNF formulas are hard to refute for CP, or that CP cannot efficiently prove the fact that the sum of all vertex degrees in an undirected graph is even (encoded in so-called *Tseitin formulas*), would constitute major breakthroughs.

We remark that there are also other proof systems inspired by linear and semidefinite programming, e.g., in [38, 39, 50], which are somewhat similar to but incomparable with cutting planes, and a deeper understanding of which appear even more challenging. Some notable early papers in proof complexity investigating these so-called *semialgebraic proof systems* were published around the turn of the millennium in [30, 31, 45], but then this area of research seems to have gone dormant. In the last few years, these proof

³⁴ Expert readers will note that we do not distinguish between PC [16] and PCR [1] below due to space constraints.

systems have made an exciting reemergence in the context of hardness of approximation, revealing unexpected and intriguing connections between approximation and proof complexity. A precursor to this is the work by Schoenebeck [49], which gave strong integrality gaps in the so-called Lasserre SDP hierarchy using results from proof complexity. These results were later realized to be a rediscovery of results by Grigoriev [29] proving degree lower bounds for what he called the *Positivstellensatz Calculus* [31]. More recently we have the work of Barak et al. [4], which was the first to explicitly point out this intriguing connection between approximability and proof complexity. Following this paper, several papers have appeared that continue the fruitful exploration of the interplay between approximability and proof complexity. Results from this area also appeared in the invited talk of Boaz Barak at the International Congress of Mathematicians in 2014 (see [5]).

The paper [19] initiated research in proof complexity focused on a more general and powerful family of propositional proof systems called *Frege systems*. Such systems consist of a finite implicationally complete set of axioms and inference rules (let us say over connectives AND, OR, and NOT for concreteness), where new formulas are derived by substitution into the axioms and inference rules. Various forms of Frege systems (also called *Hilbert systems*) typically appear in logic textbooks, and typically the exact definitions vary. Such distinctions do not matter for our purposes, however—it was shown in [19] that all such systems are equivalent up to an at most polynomial blow-up in the proof size.

Frege systems are well beyond what we can prove nontrivial lower bounds for; the situation is similar to the problem of proving lower bound on the size of Boolean circuits. Therefore restricted versions of Frege systems have been studied. One natural restriction is to allow unbounded fan-in AND-OR formulas (where negations appear only in front of atomic variables) but to require that all formulas appearing in a proof have bounded depth (i.e., a bounded number of alternations between AND and OR). Such a model is an analogue of the bounded-depth circuits studied in circuit complexity, but first arose in the context of bounded first-order arithmetic in logic [12, 41]. For such *bounded-depth Frege systems* exponential lower bounds on proof size were obtained in [37, 42], but these lower bounds only work for depth smaller than $\log \log n$. This depth lower bound was very recently improved to $\sqrt{\log n}$ in [43], but in terms of the size lower bound this recent result is much weaker. By comparison, for the corresponding class in circuit complexity strong size lower bounds are known all the way up to depth $\log n / \log \log n$. Also, if one extends the set of connectives with exclusive or (also called parity) to obtain *bounded-depth Frege with parity gates*, then again no lower bounds are known, although strong lower bounds have been shown for the analogous class in circuit complexity [46, 51].

The quest for lower bounds for bounded-depth Frege systems and beyond are mainly motivated by the P vs. NP problem. Regarding connections to SAT solving, it is mostly weaker proof systems such as resolution, polynomial calculus, and cutting planes that are of interest, whereas the variants of Frege systems discussed above do not seem to be suitable foundations for SAT solvers. The issue here is that not only do we want our proof system to be as powerful as possible, i.e., having short proofs for the formulas under consideration, but we also want to be able to *find these proofs efficiently*.

We quantify this theoretically by saying that a proof system is *automatizable* if there is an algorithm that finds proofs in this system in time polynomial in the length of an optimal proof. This seems to be the right notion: If there is no short proof of a formula in the system, then we cannot expect any algorithm to

find a proof quickly, but if there is a short proof to be found we want an algorithm that is competitive with respect to the length of such a proof. Unfortunately, there seems to be a trade-off here in the sense that if a proof system is sufficiently powerful, then it is not automatizable. For instance, bounded-depth Frege systems are not automatizable under plausible computational complexity assumptions [11]. However, analogous results have later been shown also for resolution [3], and yet proof search is implemented successfully in this proof system in practice. This raises intriguing questions that seem to merit further study.

■ Goals of the Seminar

There is a rich selection of open problems that could be discussed at a workshop focused on proof complexity. Below we just give a few samples of such problems that came up during the workshop—it should be emphasized that this list is very far from exhaustive and is only intended to serve as an illustration.

For starters, there are a number of NP-complete problems for which we would like to understand the hardness with respect to polynomial calculus and other algebraic proof systems. For the problem of cliques of constant size k in graphs, there is an obvious polynomial-time algorithm (since only $\binom{n}{k} \leq n^k$ possible candidate cliques need to be checked). Whether this brute-force algorithm is optimal or not is a deep question with connections to fixed-parameter tractability and parameterized proof complexity. This is completely open for polynomial calculus, and even for resolution. The ultimate goal here would be to prove average-case lower bounds for k -clique formulas over Erdős–Rényi random graphs $G(n, p)$ with edge probability just below the threshold $p = n^{-2/(k-1)}$ for the appearance of k -cliques.

In contrast to the clique problem, graph colouring is NP-complete already for a constant number 3 of colours. If we believe that $P \neq NP$, then, in particular, it seems reasonable to expect that this problem should be hard for polynomial calculus. No such results have been known, however. On the contrary, in the papers [23–25] recognized with the *INFORMS Computing Society Prize 2010*, the authors report that they used algebraic methods formalizable in polynomial calculus that “successfully solved graph problem instances having thousands of nodes and tens of thousands of edges” and that they could not find hard instances for these algorithms. This is very surprising. For resolution, it was shown in [7] that random graphs with the right edge density are exponentially hard to deal with, and it seems likely that the same should hold also for polynomial calculus. This appears to be a very challenging problem, however, but we hope that techniques from [2, 40] can be brought to bear on it.

For cutting planes, a longstanding open problem is to prove lower bounds for random k -CNF formulas or Tseitin formulas over expander graphs. An interesting direction in the last few years has been the development of new techniques for size-space trade-offs, showing that if short cutting planes proofs do exist, such proofs must at least have high space complexity in that they require a lot of memory to be verified. Such results were first obtained via a somewhat unexpected connection to communication complexity in [35], and have more recently been strengthened in [26, 32].

Admittedly, proving lower bounds for bounded-depth Frege systems and beyond is another formidable challenge, and it only seems prudent to say that this is a high-risk proposal. However, the very recent, and exciting, progress in [43] give hope that new techniques might be developed to attack also this problem.

■ Relation to Previous Dagstuhl Seminars

The area of proof complexity has a large intersection with computational complexity theory, and are two recurring workshops at Dagstuhl dedicated to complexity theory broadly construed, namely *Computational Complexity of Discrete Problems* and *Algebraic Methods in Computational Complexity*. However, these two workshops have had very limited coverage of topics related to proof complexity in the past.

On the more applied side, there have been two workshops *SAT and Interactions* and *Theory and Practice of SAT Solving* that have explored the connections between computational complexity

and more applied satisfiability algorithms as used in industry (so-called SAT solvers). These workshops have focused on very weak proof systems, however, which are the ones that are of interest in connection to SAT solving, but have not made any connections to stronger proof systems or to bounded arithmetic.

Although proof complexity has turned out to have deep connections to both complexity theory and SAT solving, proof complexity is an interesting and vibrant enough area to merit a seminar week in its own right. This workshop at Dagstuhl provided a unique opportunity for the community to meet during a full week focusing on the latest news in various subareas and major challenges going forward.

■ References

- 1 Michael Alekhnovich, Eli Ben-Sasson, Alexander A. Razborov, and Avi Wigderson. Space complexity in propositional calculus. *SIAM Journal on Computing*, 31(4):1184–1211, 2002. Preliminary version in *STOC '00*.
- 2 Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at <http://people.cs.uchicago.edu/~razborov/files/misha.pdf>. Preliminary version in *FOCS '01*.
- 3 Michael Alekhnovich and Alexander A. Razborov. Resolution is not automatizable unless $W[P]$ is tractable. *SIAM Journal on Computing*, 38(4):1347–1363, October 2008. Preliminary version in *FOCS '01*.
- 4 Boaz Barak, Fernando G. S. L. Brandão, Aram Wettroth Harrow, Jonathan A. Kelner, David Steurer, and Yuan Zhou. Hypercontractivity, sum-of-squares proofs, and their applications. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 307–326, May 2012.
- 5 Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of International Congress of Mathematicians (ICM)*, 2014.
- 6 Paul Beame. Proof complexity. In Steven Rudich and Avi Wigderson, editors, *Computational Complexity Theory*, volume 10 of *IAS/Park City Mathematics Series*, pages 199–246. American Mathematical Society, 2004.
- 7 Paul Beame, Joseph C. Culberson, David G. Mitchell, and Christopher Moore. The resolution complexity of random graph k -colorability. *Discrete Applied Mathematics*, 153(1-3):25–47, December 2005.
- 8 Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in *STOC '99*.
- 9 Archie Blake. *Canonical Expressions in Boolean Algebra*. PhD thesis, University of Chicago, 1937.
- 10 María Bonet, Toniann Pitassi, and Ran Raz. Lower bounds for cutting planes proofs with small coefficients. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing (STOC '95)*, pages 575–584, May 1995.
- 11 María Luisa Bonet, Carlos Domingo, Ricard Gavaldà, Alexis Maciel, and Toniann Pitassi. Non-automatizability of bounded-depth Frege proofs. In *Proceedings of the 14th Annual IEEE Conference on Computational Complexity (CCC '99)*, pages 15–23, May 1999.
- 12 Samuel R. Buss. *Bounded Arithmetic*. Bibliopolis, Naples, 1986. Revision of PhD thesis.
- 13 Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *Journal of Computer and System Sciences*, 62(2):267–289, March 2001. Preliminary version in *CCC '99*.
- 14 Vašek Chvátal. Edmonds polytopes and a hierarchy of combinatorial problems. *Discrete Mathematics*, 4(1):305–337, 1973.
- 15 Vašek Chvátal and Endre Szemerédi. Many hard examples for resolution. *Journal of the ACM*, 35(4):759–768, October 1988.
- 16 Matthew Clegg, Jeffery Edmonds, and Russell Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pages 174–183, May 1996.
- 17 Peter Clote and Evangelos Kranakis. *Boolean Functions and Computation Models*. Springer, 2002.
- 18 Stephen A. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing (STOC '71)*, pages 151–158, 1971.
- 19 Stephen A. Cook and Robert Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979.
- 20 William Cook, Collette Rene Coullard, and György Turán. On the complexity of cutting-plane proofs. *Discrete Applied Mathematics*, 18(1):25–38, November 1987.
- 21 Martin Davis, George Logemann, and Donald Loveland. A machine program for theorem proving. *Communications of the ACM*, 5(7):394–397, July 1962.
- 22 Martin Davis and Hilary Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7(3):201–215, 1960.
- 23 Jesús A. De Loera, Jon Lee, Peter N. Malkin, and Susan Margulies. Hilbert's Nullstellensatz and an algorithm for proving combinatorial infeasibility. In *Proceedings of the 21st International Symposium on Symbolic and Algebraic Computation (ISSAC '08)*, pages 197–206, July 2008.

- 24 Jesús A. De Loera, Jon Lee, Peter N. Malkin, and Susan Margulies. Computing infeasibility certificates for combinatorial problems through Hilbert's Nullstellensatz. *Journal of Symbolic Computation*, 46(11):1260–1283, November 2011.
- 25 Jesús A. De Loera, Jon Lee, Susan Margulies, and Shmuel Onn. Expressing combinatorial problems by systems of polynomial equations and Hilbert's Nullstellensatz. *Combinatorics, Probability and Computing*, 18(04):551–582, July 2009.
- 26 Susanna F. de Rezende, Jakob Nordström, and Marc Vinyals. How limited interaction hinders real communication (and what it means for proof and circuit complexity). In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS '16)*, pages 295–304, October 2016.
- 27 Ralph E. Gomory. An algorithm for integer solutions of linear programs. In R.L. Graves and P. Wolfe, editors, *Recent Advances in Mathematical Programming*, pages 269–302. McGraw-Hill, New York, 1963.
- 28 Dima Grigoriev. Tseitin's tautologies and lower bounds for Nullstellensatz proofs. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS '98)*, pages 648–652, November 1998.
- 29 Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1–2):613–622, May 2001.
- 30 Dima Grigoriev, Edward A. Hirsch, and Dmitrii V. Pasechnik. Complexity of semialgebraic proofs. *Moscow Mathematical Journal*, 2(4):647–679, 2002. Preliminary version in *STACS '02*.
- 31 Dima Grigoriev and Nicolai Vorobjov. Complexity of Null- and Positivstellensatz proofs. *Annals of Pure and Applied Logic*, 113(1–3):153–160, December 2001.
- 32 Mika Göös and Toniann Pitassi. Communication lower bounds via critical block sensitivity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC '14)*, pages 847–856, May 2014.
- 33 Armin Haken. The intractability of resolution. *Theoretical Computer Science*, 39(2-3):297–308, August 1985.
- 34 Armin Haken and Stephen A. Cook. An exponential lower bound for the size of monotone real circuits. *Journal of Computer and System Sciences*, 58(2):326–335, 1999.
- 35 Trinh Huynh and Jakob Nordström. On the virtue of succinct proofs: Amplifying communication complexity hardness to time-space trade-offs in proof complexity (Extended abstract). In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC '12)*, pages 233–248, May 2012.
- 36 Russell Impagliazzo, Pavel Pudlák, and Jiří Sgall. Lower bounds for the polynomial calculus and the Gröbner basis algorithm. *Computational Complexity*, 8(2):127–144, 1999.
- 37 Jan Krajíček, Pavel Pudlák, and Alan R. Woods. An exponential lower bound to the size of bounded depth Frege proofs of the pigeonhole principle. *Random Structures and Algorithms*, 7(1):15–40, 1995. Preliminary version in *STOC '92*.
- 38 Jean B. Lasserre. An explicit exact SDP relaxation for nonlinear 0-1 programs. In *Proceedings of the 8th International Conference on Integer Programming and Combinatorial Optimization (IPCO '01)*, volume 2081 of *Lecture Notes in Computer Science*, pages 293–303. Springer, June 2001.
- 39 László Lovász and Alexander Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM Journal on Optimization*, 1(2):166–190, 1991.
- 40 Mladen Mikša and Jakob Nordström. A generalized method for proving polynomial calculus degree lower bounds. In *Proceedings of the 30th Annual Computational Complexity Conference (CCC '15)*, volume 33 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 467–487, June 2015.
- 41 Jeff B. Paris and Alex J. Wilkie. Counting problems in bounded arithmetic. In *Methods in Mathematical Logic: Proceedings of the 6th Latin American Symposium on Mathematical Logic*, volume 1130 of *Lecture Notes in Mathematics*, pages 317–340. Springer, 1985.
- 42 Toniann Pitassi, Paul Beame, and Russell Impagliazzo. Exponential lower bounds for the pigeonhole principle. *Computational Complexity*, 3:97–140, 1993. Preliminary version in *STOC '92*.
- 43 Toniann Pitassi, Benjamin Rossman, Rocco Servedio, and Li-Yang Tan. Poly-logarithmic Frege depth lower bounds. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC '16)*, pages 644–657, June 2016.
- 44 Pavel Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *Journal of Symbolic Logic*, 62(3):981–998, September 1997.
- 45 Pavel Pudlák. On the complexity of propositional calculus. In S. Barry Cooper and John K. Truss, editors, *Sets and Proofs*, volume 258 of *London Mathematical Society Lecture Note Series*, pages 197–218. Cambridge University Press, 1999.
- 46 Alexander A. Razborov. Lower bounds on the size of bounded depth networks over a complete basis with logical addition. *Matematicheskie Zametki*, 41(4):598–607, 1987. English Translation in *Mathematical Notes of the Academy of Sciences of the USSR*, 41(4):333–338, 1987.
- 47 Alexander A. Razborov. Proof complexity and beyond. *ACM SIGACT News*, 47(2):66–86, June 2016.
- 48 John Alan Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.
- 49 Grant Schoenebeck. Linear level Lasserre lower bounds for certain k -CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 593–602, October 2008.
- 50 Hanif D. Sherali and Warren P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3:411–430, 1990.
- 51 Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC '87)*, pages 77–82, 1987.
- 52 Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.





Fig. 6.1
“How can you form breakout groups at a @dagstuhl seminar? Collect research preferences and apply a notice-board-driven clustering algorithm. #sefiat” Twitter post by 18343 G1 Dagstuhl Seminar participant Simos Gerasimou.
<https://twitter.com/simosgerasimou/status/1032531956112781314>. Photo courtesy of Simos Gerasimou.

6.5 Genetic Improvement of Software

Organizers: Justyna Petke, Stephanie Forrest, William B. Langdon, and Claire Le Goues
Seminar No. 18052

Date: January 28–February 2, 2018 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.8.1.158

© Creative Commons BY 3.0 Unported license

© Justyna Petke, Stephanie Forrest, William B. Langdon, and Claire Le Goues



Participants: Brad Alexander, Wolfgang Banzhaf, Benoit Baudry, Bobby R. Bruce, Celso G. Camilo-Junior, Myra B. Cohen, Benjamin Danglot, Stephanie Forrest, Nicolas Harrand, Colin G. Johnson, Krzysztof Krawiec, William B. Langdon, Claire Le Goues, Alexandru Marginean, Michael Pradel, Joseph Renzullo, Eric Schulte, Lukas Sekanina, Marija Selakovic, Shin Hwei Tan, Christopher Timperley, Leonardo Trujillo, Emamurho Ugherughe, Sergi Valverde, Zdenek Vasicek, Markus Wagner, John R. Woodward, Shin Yoo

Genetic improvement (GI) uses automated search to find improved versions of existing software. It can be used for improvement of both functional and non-functional properties of software. Much of the early success came from the field of automated program repair. However, GI has also been successfully used to optimise for efficiency, energy and memory consumption as well as automated transplantation of a piece of functionality from one program to another. These results are impressive especially given that genetic improvement only arose as a separate research area in the last few years. Thus the time was ripe to organise a seminar that would gather researchers from GI and related areas together to summarise the current achievements and identify avenues for further research.

The seminar attracted researchers from various GI-related software engineering areas, ranging from automated software repair through genetic programming and software testing to biological and evolutionary computation. The talks covered the latest research and speculations on future research both in the practical applications of genetic improvement, such as energy consumption optimisation and automated parallelisation, to initial results on much lacking GI theory. In particular, GI theory and indeed software in general were discussed in terms of search landscape analysis. Other talks covered software testing and bug repair. The participants also identified a set of benchmarks and tools for GI. These have been published at the geneticimprovementofsoftware.com website to allow other researchers to compare their new technologies against the state-of-the-art.

The seven breakout groups' topics ranged from re-evaluating the basic components of the GI framework, such as fitness functions and traversing the GI search space, to identifying issues related to adoption of GI in industry. One of the issues has been explanation of the automatically generated changes, which might

be a roadblock in applying them in the real-world, especially safety-critical, software.

The seminar has already led to a few publications. For example, four papers accepted to the 4th International Genetic Improvement Workshop (GI-2018)³⁵, co-located with the International Conference on Software Engineering (ICSE), were written by one or more workshop participants. Indeed most were started in Dagstuhl. Several other collaborations have been established, with plans for visits and further research on topics identified at the seminar. We look forward to results of this work initiated at Dagstuhl.

Introduction Genetic improvement (GI) uses automated search to find improved versions of existing software [6,8]. It uses optimisation, machine learning techniques, particularly search based software engineering techniques such as genetic programming [1, 2, 9]. to improve existing software. The improved program need not behave identically to the original. For example, automatic bug fixing improves program code by reducing or eliminating buggy behaviour, whilst automatic transplantation adds new functionality derived from elsewhere. In other cases the improved software should behave identically to the old version but is better because, for example: it runs faster, it uses less memory, it uses less energy or it runs on a different type of computer.

GI differs from, for example, formal program translation, in that it primarily verifies the behaviour of the new mutant version by running both the new and the old software on test inputs and comparing their output and performance in order to see if the new software can still do what is wanted of the original program and is now better. Using less constrained search allows not only functional improvements but also each search step is typically far cheaper, allowing GI to scale to substantial

³⁵ <http://geneticimprovementofsoftware.com/>

programs. Genetic improvement can be used to create large numbers of versions of programs, each tailored to be better for a particular use or for a particular computer, or indeed (e.g. to defeat the authors of computer viruses) simply to be different. Other cases where software need to be changed include porting to new environments (e.g. parallel computing [3] mobile devices) or for code obfuscation to prevent reverse engineering [7].

Genetic improvement can be used with multi-objective optimisation to consider improving software along multiple dimensions or to consider trade-offs between several objectives, such as asking GI to evolve programs which trade speed against the quality of answers they give. Of course, it may be possible to find programs which are both faster and give better answers. Mostly Genetic Improvement makes typically small changes or edits (also known as mutations) to the program's source code, but sometimes the mutations are made to assembly code, byte code or binary machine code.

GI arose as a separate field of research only in the last few years. Even though its origins could be traced back to the work by Ryan & Walsh [18] in 1995, it is the work by Arcuri [10] and White [20] that led to the development and wider uptake of the GI techniques. The novelty lay in applying heuristics to search for code mutations that improved existing software. Both Arcuri and White applied genetic programming (GP), with Arcuri using also hill-climbing and random search on a small set of problems. Rather than trying to evolve a program from scratch, as in traditional GP, Arcuri and White took the approach of seeding [5] the initial population with copies of the original program. Next, instead of focusing on evolving a program fulfilling a particular task, as has been done before, Arcuri and White used GP to improve their programs either to fix existing bugs or to improve the non-functional properties of software, in particular, its efficiency and energy consumption. Both Arcuri and White, however, applied their, now known as, GI techniques, to relatively small benchmarks having little resemblance to large scale real-world problems.

The bug fixing approach was taken up by Forrest, Le Goues and Weimer et al. [12, 15, 19] and adapted for large software systems. One of the insights that allowed for this adoption was an observation that full program variants need not be evolved, yet only a sequence of edits, which are then applied to the original program. Validity of the resultant modified software was then evaluated on a set of test cases, assumed to capture desired program behaviour, as in previous work. This strand of research led to the development of first GP-based automated software repair tool called GenProg [15]. Success of this automated bug fixing work led to several best paper awards and two 'Humie' awards (international prizes for human-competitive results produced by genetic and evolutionary computation <http://www.human-competitive.org/>) and inspired work on other automated software repair tools, including Angelix [16], which uses a form of constraint solving to synthesise bug fixes.

Research on improvement of non-functional software properties has yet to garner the attention and software development effort as the work on automated bug fixing. Langdon et al. [3, 13, 14] published several articles on efficiency improvement and parallelisation using GI. They were able to improve efficiency of large pieces of state-of-the-art software. Moreover, the genetically improved version of a bioinformatics software called BarraCUDA is the first instance of a genetically improved piece of software adapted into development [4, 14].

Petke et al. [17] set themselves a challenge of improving efficiency of a highly-optimised piece of software that has been improved by expert human developers over a period of several years. In particular, a famous Boolean satisfiability (SAT) solver

was chosen, called MiniSAT. It implements the core technologies of SAT solving and inspired a MiniSAT-hack track at the annual international SAT solver competitions, where anyone can submit their own version of MiniSAT. Petke et al. showed that further efficiency improvements can be made by using this source of genetic material for the GP process and specializing the solver for a particular downstream application. This work showed the initial potential of what is now called automated software transplantation and was awarded a Silver 'Humie'. Further work on automated software transplantation won an ACM SIGSOFT distinguished paper award and a Gold 'Humie' at this year's Genetic and Evolutionary Computation Conference (GECCO-2017) [11].

■ Aims of the Seminar

The seminar brought together researchers in this new field of software engineering to investigate what is achievable with current technology and the current impediments to progress (if indeed there are any) of what can be achieved within the field in the future and how GI can affect the software development process.

With the growing popularity of the field, multiple awards and fast progress GI research in the field, it is the right time to gather top the academics in GI and related fields to push the boundaries of what genetic improvement can achieve even further.

This seminar brought researchers working in genetic improvement and related areas, such as automated program repair, software testing and genetic programming, together. It summarized achievements in automated software optimisation. We will use this summary as a basis to investigate how optimisation approaches from the different fields represented at the seminar can be combined to produce a robust industry-ready set of techniques for software improvement.

■ References

- 1 Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, and Frank D. Francone. *Genetic Programming – An Introduction; On the Automatic Evolution of Computer Programs and its Applications*. Morgan Kaufmann, San Francisco, CA, USA, January 1998.
- 2 John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- 3 W. B. Langdon and M. Harman. Evolving a CUDA kernel from an nVidia template. In Pilar Sobrevilla, editor, *2010 IEEE World Congress on Computational Intelligence*, pages 2376–2383, Barcelona, 18–23 July 2010. IEEE.
- 4 W. B. Langdon and Brian Yee Hong Lam. Genetically improved BarraCUDA. *BioData Mining*, 20(28), 2 August 2017.
- 5 W. B. Langdon and J. P. Nordin. Seeding GP populations. In Riccardo Poli, Wolfgang Banzhaf, William B. Langdon, Julian F. Miller, Peter Nordin, and Terence C. Fogarty, editors, *Genetic Programming, Proceedings of EuroGP'2000*, volume 1802 of *LNCS*, pages 304–315, Edinburgh, 15–16 April 2000. Springer-Verlag.
- 6 William B. Langdon. Genetically improved software. In Amir H. Gandomi, Amir H. Alavi, and Conor Ryan, editors, *Handbook of Genetic Programming Applications*, chapter 8, pages 181–220. Springer, 2015.
- 7 Justyna Petke. Genetic improvement for code obfuscation. In Justyna Petke, David R. White, and Westley Weimer, editors, *Genetic Improvement 2016 Workshop*, pages 1135–1136, Denver, July 20–24 2016. ACM.
- 8 Justyna Petke, Saemundur O. Haraldsson, Mark Harman, William B. Langdon, David R. White, and John R. Woodward. Genetic improvement of software: a comprehensive survey. *IEEE Transactions on Evolutionary Computation*. In press.
- 9 Riccardo Poli, William B. Langdon, and Nicholas Freitag McPhee. *A field guide to genetic programming*. Published via <http://lulu.com> and freely available at <http://www.gp-field-guide.org.uk>, 2008. (With contributions by J. R. Koza).
- 10 Andrea Arcuri. Automatic software generation and improvement through search based techniques. PhD. Univ. of Birmingham, 2009.
- 11 Earl T. Barr, Mark Harman, Yue Jia, Alexandru Marginean, and Justyna Petke. Automated software transplantation. In *ISSTA*, pages 257–269, 2015.
- 12 Stephanie Forrest, ThanhVu Nguyen, Westley Weimer, and Claire Le Goues. A genetic programming approach to automated software repair. In *GECCO*, pages 947–954, 2009.
- 13 William B. Langdon and Mark Harman. Optimizing existing software with genetic programming. *IEEE Transactions on Evolutionary Computation*, 19(1):118–135, 2015.
- 14 William B. Langdon, Brian Yee Hong Lam, Justyna Petke, and Mark Harman. Improving CUDA DNA analysis software with genetic programming. In *GECCO*, pages 1063–1070, 2015.
- 15 Claire Le Goues, Michael Dewey-Vogt, Stephanie Forrest, and Westley Weimer. A systematic study of automated program repair: Fixing 55 out of 105 bugs for \$8 each. In *ICSE*, pages 3–13, 2012.
- 16 Sergey Mechtaev, Jooyong Yi, and Abhik Roychoudhury. Angelix: scalable multiline program patch synthesis via symbolic analysis. In *ICSE*, pages 691–701, 2016.
- 17 Justyna Petke, Mark Harman, William B. Langdon, and Westley Weimer. Using genetic improvement and code transplants to specialise a C++ program to a problem class. In *EuroGP*, pages 137–149, 2014.
- 18 Paul Walsh and Conor Ryan. Automatic conversion of programs from serial to parallel using genetic programming - the Paragen system. In *ParCo*, pages 415–422, 1995.
- 19 Westley Weimer, ThanhVu Nguyen, Claire Le Goues, and Stephanie Forrest. Automatically finding patches using genetic programming. In *ICSE*, pages 364–374, 2009.
- 20 David R. White. Genetic programming for low-resource systems. PhD. Univ. of York, 2009.

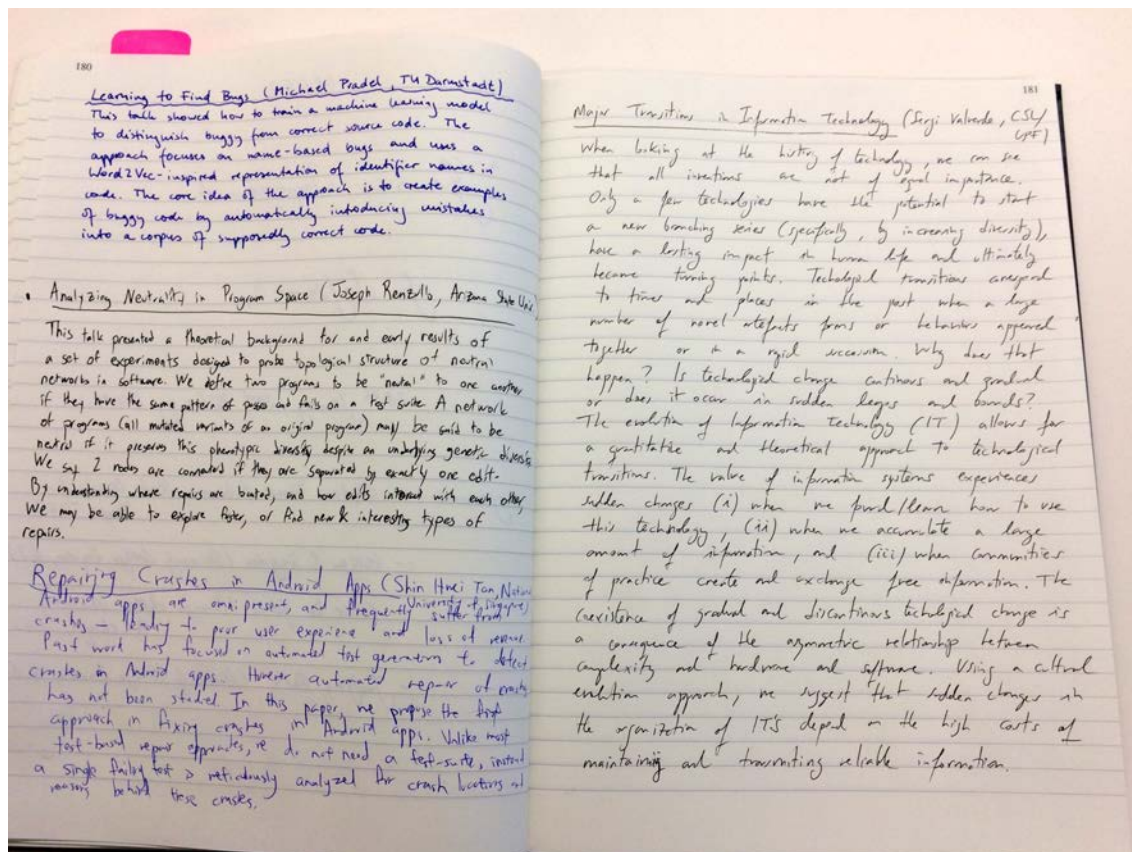


Fig. 6.2

"I love the tradition @dagstuhl of collecting hand-written abstracts. It's just great."

Twitter post by 18052 Dagstuhl Seminar participant Sergi Valverde.

<https://twitter.com/svalver/status/958361552377335808>. Photo courtesy of Sergi Valverde.

6.6 Evidence About Programmers for Programming Language Design

Organizers: Andreas Stefik, Bonita Sharif, Brad A. Myers, and Stefan Hanenberg
Seminar No. 18061

Date: February 4–9, 2018 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.8.2.1
 © Creative Commons BY 3.0 Unported license
 © Michael Coblenz



Participants: Jonathan Aldrich, Craig Anslow, Ameer Armaly, Johannes Bechberger, Brett A. Becker, Andrew Begel, Tanja Blascheck, Neil C. C. Brown, Michael Coblenz, Igor Crk, John M. Daughtry, Fabian Deitelhoff, Rob DeLine, Brian Dorn, Andrew Duchowski, Scott Fleming, Baker Franke, Reiner Hähnle, Matthias Hauswirth, Felienne Hermans, Johannes Hofmeister, Ciera Jaspan, Antti-Juhani Kaijanaho, Andrew J. Ko, Thomas LaToza, Andrew Macvean, Jonathan I. Maletic, Amelia A. McNamara, Briana B. Morrison, Brad A. Myers, Lutz Prechelt, Sibylle Schupp, Bonita Sharif, Andreas Stefik, Walter F. Tichy, Phillip Merlin Uesbeck, Lea Verou

Programming languages underlie and have significant impact on software development, especially in terms of the ability of programmers to achieve their goals. Although designers of programming languages can already reason about the *formal* properties of their languages, few tools are available to assess the impact of design decisions on programmers and software engineers.

At Dagstuhl Seminar 18061, a diverse set of participants gathered to review the existing body of evidence about programmers that has implications on programming language design. Participants also reviewed existing research methods, such as eye tracking, that may help better understand the impact of language design decisions on programmers. Participants brainstormed a long list of possible research questions for investigation (§4 of the full report), and then divided into working groups (§5 of the full report) to focus on several areas of research interest, including novices, context switching and cognitive load, language features, emotional attachment to languages, and representativeness of subjects in studies. In each area, participants proposed research methods and questions that they felt would be valuable to address in the future. Then, the group discussed and prioritized these research questions.

The seminar included a discussion of the need for an evidence standard in empirical studies of programming languages, focusing on content of the evidence standard, adoption mechanisms, and criteria for what it might include in our field. Finally, the seminar concluded with a discussion of future directions for research, including a list of research questions that the participants were planning on collaborating on in the near future.

6.7 Planning and Operations Research

Organizers: J. Christopher Beck, Daniele Magazzeni, Gabriele Röger, and Willem-Jan Van Hoeve

Seminar No. 18071

Date: February 11–16, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.2.26

© Creative Commons BY 3.0 Unported license

© Florian Pommerening

Participants: Roman Bartak, J. Christopher Beck, Adi Botea, Christina N. Burt, Hadrien Cambazard, Michael Cashmore, Alessandro Cimatti, Mathijs de Weerd, Jeremy D. Frank, Patrik Haslum, Emmanuel Hebrard, Malte Helmert, John N. Hooker, Serdar Kadioglu, Michael Katz, Thorsten Koch, Sven Koenig, Michele Lombardi, Daniele Magazzeni, Andrea Micheli, Christian Muise, Eva Onaindia, Gilles Pesant, Chiara Piacentini, Nicola Policella, Florian Pommerening, Mark Roberts, Gabriele Röger, Louis-Martin Rousseau, Hana Rudová, Domenico Salvagnin, Scott Sanner, Pierre Schaus, Mathijs Spaan, Charlotte Truchet, Willem-Jan Van Hoeve, Parisa Zehtabi



This seminar brought together leading experts in the fields of AI planning, constraint programming and operations research. These areas historically come from different roots but are all concerned with supporting decision making in complex systems which a huge space of interacting options. While the approach and focus is different, some concepts have been developed in multiple areas and some solution techniques are or could be transferred. There is also a growing intersection of the areas that considers hybrid problems or uses solvers developed in one area to solve problems from a different area, for example by compiling planning problems into MIP or by using CP for subproblems in a MIP solver. Solvers of a different community are often used as black boxes and the deeper understanding of each other's area of expertise that was developed in this seminar will help to foster collaboration and transfer knowledge between the areas.

The seminar started with eleven short but intense tutorials on Monday and Tuesday morning. The tutorials on Monday conveyed the basics of AI Planning, MIP, and CP. They also already introduced the main connections between the fields by talking about compilations from planning to CP and MIP and using LPs as heuristics in planning. The tutorials on Tuesday delved deeper into areas that became the focus of discussion later in the seminar, such as non-deterministic planning, Markov decision processes, and decision diagrams. Front-loading these tutorials worked well to bring everyone up to speed and created a good basis for the rest of the seminar.

The rest of the seminar was organized into working groups that included one to three short presentations followed by a longer discussion all focused on a central topic. Three of these sessions were organized as break-out sessions where the participants split into two groups, each discussion one topic and then reconvening to present the main points discussed in each group to each other. The schedule for each day was created on the evening before which kept the topics flexible and allowed the organizers to include

topics that came up during the discussion. Notes on each of the working groups and abstracts of the tutorials are included in the rest of this report.

6.8 Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization

Organizers: Pierre Bonami, Ambros M. Gleixner, Jeff Linderoth, and Ruth Misener
Seminar No. 18081

Date: February 18–23, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.2.64

© Creative Commons BY 3.0 Unported license

© Pierre Bonami, Ambros M. Gleixner, Jeff Linderoth and Ruth Misener



Participants: Tobias Achterberg, Claire Adjiman, Shabbir Ahmed, Kurt M. Anstreicher, Radu Baltean-Lugojan, Pietro Belotti, David Bernal Neira, Timo Berthold, Christian Bliet, Pierre Bonami, Fani Boukouvala, Andrea Callia D’Iddio, Sanjeeb Dash, Alberto Del Pia, Santanu Dey, Matteo Fischetti, Ambros M. Gleixner, Ignacio Grossmann, Andreas Grothey, Oktay Gunluk, Akshay Gupte, Hassan Hijazi, Aida Khajavirad, Carl Damon Laird, Amélie Lambert, Jon Lee, Sven Leyffer, Leo Liberti, Frauke Liers, Jeff Linderoth, Andrea Lodi, James Luedtke, Ashutosh Mahajan, Alexander Martin, Ruth Misener, Miten Mistry, Benjamin Müller, Sebastian Sager, Nikolaos V. Sahinidis, Felipe Serrano, Mohit Tawarmalani, Juan Pablo Vielma, Stefan Vigerske, Robert Weismantel, Angelika Wiegele, Sven Wiese

This workshop aimed to address this mismatch between natural optimization models for important scientific problems and practical optimization solvers for their solution. By bringing together experts in both theory and implementation, this workshop energized efforts making MINLP as ubiquitous a paradigm for both modeling and solving important decision problems as mixed-integer linear programming (MIP) and nonlinear programming (NLP) have become in recent years. In particular, we highlighted:

- **MINLP Solver Software** Early in the workshop, the main developers of MINLP software packages outlined the current state of their software. This served as a needs analysis for the community to identify crucial areas for future development. We also dedicated a break-out session discussing best practices for conducting scientifically-meaningful computational experiments in MINLP.
- **Intersecting Mixed-Integer & Nonlinear Programming** MINLP is a superset of both mixed integer linear optimization and nonlinear optimization, so we leveraged the best methods from both by incorporating both sets of experts.
- **Driving Applications** Applications experts, e.g. in petrochemicals, manufacturing, and gas networks, offered their perspectives on what practitioners need from MINLP solvers. We dedicated an entire break-out session to energy applications and explored what are the needs for MINLP within the energy domain. During the open problem session, several other applications experts outlined other open problems in engineering.
- **Connections between MINLP and machine learning** Many machine learning challenges can be formulated as MINLP. Also, machine learning can significantly improve MINLP solver software. We explored these connections at length in a break-out session.

This seminar brought together an assortment of computer scientists with expertise in mathematical optimization. Many of the presentations were more theoretical and suggested new technologies that the solver software could incorporate. Other presentations were more practical and discussed building solver software or applying that software to specific domain applications.

As a result of this seminar, we are planning a special issue in the journal “Optimization & Engineering”. We are also working to turn the notes from our open problem session into a larger document that will start a conversation with the entire mathematical optimisation community. Participants broadly expressed that this week at Dagstuhl helped them workshop their papers, so several academic papers will explicitly mention the Dagstuhl seminar. Finally, a new set of metrics for comparing MINLP solvers were developed at this meeting and will greatly aid future solver testing.

6.9 Formal Methods for the Synthesis of Biomolecular Circuits

Organizers: Yaakov Benenson, Neil Dalchau, Heinz Koepl, and Oded Maler

Seminar No. 18082

Date: February 18–23, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.2.88

© Creative Commons BY 3.0 Unported license

© Heinz Koepl

Participants: Aaron Adler, Ben Barak, Chris Barnes, Jacob Beal, Yaakov Benenson, Milan Ceska, Neil Dalchau, Sara-Jane Dunn, François Fages, Eric Fanchon, Thomas Gorochowski, Maleen Hanst, Nathan Hillson, Johannes Kabisch, Heinz Koepl, Jan Madsen, Oded Maler, Gareth Molyneux, Radu Muschevici, Chris J. Myers, Irene Otero-Muras, James Scott-Brown, Boyan Yordanov, Paolo Zuliani



The seminar brought together experts in formal methods for the verification and synthesis of hardware and software with wet-lab and dry-lab synthetic biologists to (1) achieve a common understanding of the current state of design methodology in synthetic biology; (2) to identify the limitations of current approaches and (3) to investigate dedicated solutions to the synthesis problem in synthetic biology. Some of these methods are based on leveraging experience and methods from electronic design automation (EDA) and from program synthesis and verification. In addition, ideas for entirely new methodologies specifically tailored for synthetic biology are likely to emerge. For example, features that are not apparent in electronic circuits such as heterogeneity and variability between the cells and between the circuits embedded in different cells, were addressed.

Apart from talk by participants, the seminar also featured break out session that were well received by the participants. In particular, we had sessions on “Modeling context-dependency of synthetic circuits” on “Metrology in Synthetic Biology” and on “Formal Specification for Biological Circuit Synthesis”.

6.10 Data Consistency in Distributed Systems: Algorithms, Programs, and Databases

Organizers: Annette Bieniusa, Alexey Gotsman, Bettina Kemme, and Marc Shapiro
Seminar No. 18091

Date: February 25–March 2, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.2.101

© Creative Commons BY 3.0 Unported license

© Annette Bieniusa, Alexey Gotsman, Bettina Kemme, and Marc Shapiro



Participants: Peter Alvaro, Mahesh Balakrishnan, Carlos Baquero, Annette Bieniusa, Ahmed Bouajjani, Manuel Bravo, Sebastian Burckhardt, Andrea Cerone, Gregory Chockler, Khuzaima Daudjee, Diego Didona, Amr El Abbadi, Carla Ferreira, Alexey Gotsman, Suresh Jagannathan, Bettina Kemme, Brad King, Kyle Kingsbury, Martin Kleppmann, Christopher Meiklejohn, Roland Meyer, Maged M. Michael, Pascal Molli, Roberto Palmieri, Matthieu Perrin, Gustavo Petri, Nuno Preguica, Luis Rodrigues, Rodrigo Rodrigues, Masoud Saeida Ardekani, Sebastian Schweizer, Marc Shapiro, Pierre Sutra, Viktor Vafeiadis, Peter Van Roy

Large-scale distributed systems have become ubiquitous, and there are a variety of options to develop, deploy, and operate such applications. Typically, this type of application is data-centric: it retrieves, stores, modifies, forwards, and processes data from different sources. However, guaranteeing availability, preventing data loss, and providing efficient storage solutions are still major challenges that a growing number of programmers are facing when developing large-scale distributed systems. In our seminar, we brought together academic and industrial researchers and practitioners to discuss the status quo of data consistency in distributed systems. As result of talks and discussions, we identified several topics of interest that can be grouped into the following four areas.

Theoretical foundations: The seminar included a tutorial on specification of consistency guarantees provided by distributed systems and talks on comparing different styles of specification and expressing replicated data type semantics in Datalog. Different specification styles are suitable for different purposes and more work is needed to identify the most appropriate ones. The seminar also included talks on formally reasoning about which consistency levels are enough to satisfy correctness properties of applications. The talks demonstrated that formal verification is a promising approach to cope with the challenge of selecting appropriate consistency levels.

Distributed systems and database technologies: With the growing number of replicated data stores, the two fields of distributed systems and databases are moving closer together. The communities should be made more aware of each others results. A common concern in agreement, i.e., ensuring that database copies are updated correctly. Traditionally, the distributed systems community has based many of their approaches on classical consensus algorithms or looked at weaker consistency models. In contrast, database systems focused most

work on 2-phase commit protocols and eager update protocols. At the same time, the database community also considered other ACID aspects that required to combine commit protocols with concurrency control protocols and recovery schemes. In the last decade however, and in particular with practical implementations of the Paxos consensus algorithms, and the use of file replication in storage systems for availability, work of the two communities has come closer together. A challenge in this context is that work that emerges from the different communities still makes slightly different assumptions about failure and correctness models. They can often be quite subtle so that the differences are not obvious, even to the experts. And they can lead to very different approaches to find solutions. Bridging this gap in terms of understanding each other, and the implications of correctness and failure models remains a challenging task. As an example, the separation of the concepts of atomicity, isolation and durability in the database world offers many opportunities for optimization, but includes extra complexity when analyzing which algorithms are appropriate in which situations.

Conflict-handling in highly-scalable systems: In the last years, conflict-free replicated data types (CRDTs) have been adopted by an ever-growing number of products and companies to deal with high-availability requirements under concurrent modifications of data. Recent advances in related techniques for collaborative editing might make it possible that hundreds of people work together on a shared document or data item with limited performance impact. Several talks presented programming guidelines, static analyses, and related tools for safe usage of CRDTs in situations where eventual consistency is not enough to maintain application invariants.

Programming models for distributed systems: Micro-services have become a standard approach for constructing large-scale distributed systems, though microservice composition

and scalability raises a lot of questions. Some presentations discussed current work on actor-based and data-flow programming. Design for testability and test frameworks are crucial for providing reliable services, but they currently require a lot of experience as

of today. We believe that future progress on programming models and new results in theoretical foundations will help to simplify this challenging task and support programmers in building safe systems.




Fig. 6.3

“@dagstuhl is a perfect place for #MINLPDagstuhl. Nonconvexities everywhere. @pietroBelotti”

Twitter post by 18081 Dagstuhl Seminar participant Jeff Linderth.

<https://twitter.com/JeffLinderth/status/966690589122187264>. Photo courtesy of Jeff Linderth.

6.11 The Logical Execution Time Paradigm: New Perspectives for Multicore Systems

Organizers: Rolf Ernst, Stefan Kuntz, Sophie Quinton and Martin Simons
Seminar No. 18092

Date: February 25–28, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.2.122

© Creative Commons BY 3.0 Unported license

© Rolf Ernst, Stefan Kuntz, Martin Simons, Borislav Nikolić, Sophie Quinton, and Hermann von Hasseln



Participants: Leonie Ahrendts, James H. Anderson, Matthias Beckert, Alessandro Biondi, Bert Boeddeker, Björn B. Brandenburg, Sylvain Cotard, Marco Di Natale, Benoit Dupont de Dinechin, Rolf Ernst, Glenn Farrall, Gerhard Fohler, Alain Girault, Mathieu Jan, Karl Henrik Johansson, Sebastian Kehr, Christoph M. Kirsch, Stefan Kuntz, Ralph Mader, Martina Maggio, Florence Maraninchi, Jorge Luis Martinez Garcia, Andreas Naderlinger, Moritz Neukirchner, Borislav Nikolic, Nathan Otterness, Claire Pagetti, Paolo Pazzaglia, Christophe Prévot, Sophie Quinton, Stefan Resmerita, Hermann von Hasseln, Eugene Yip, Dirk Ziegenbein

The Logical Execution Time (LET) abstraction, which was originally introduced as a real-time programming paradigm, has gained traction recently in the automotive industry with the shift to multicore architectures. The objective of this Dagstuhl Seminar was to investigate new opportunities and challenges raised by the use of LET as a basis for implementing parallel execution of control software.

LET abstracts from the actual timing behavior of real-time tasks on the physical platform: Independent of when a task executes, the time interval between its reading input and writing output is fixed by the LET. This introduces a separation between functionality on the one hand, and mapping and scheduling on the other hand. It also provides a clean interface between the timing model used by the control engineer and that of the software engineer.

The LET paradigm was considered until recently by the automotive industry as not efficient enough in terms of buffer space and timing performance. The shift to embedded multicore processors has represented a game changer: The design and verification of multicore systems is a challenging area of research that is still very much in progress. Predictability clearly is a crucial issue which cannot be tackled without changes in the design process. Several OEMs and suppliers have come to the conclusion that LET might be a key enabler and a standardization effort is already under way in the automotive community to integrate LET into AUTOSAR.

The seminar brought together researchers and practitioners from different backgrounds to discuss and sketch solutions to the problems raised by the use of LET in multicore systems, with a focus on the automotive domain. The program was structured around the following topics: (i) Implementations of LET; (ii) LET and related paradigms; (iii) LET and control; (iv) Future directions of LET. The fruitful discussions covered the following issues:

- LET was designed as a programming paradigm but is now being used as a mechanism for predictable communication. How can the principles of LET be adapted accordingly? How should LET values be chosen?
- LETs act as deadlines for tasks, which means that they must be dimensioned for the worst-case response time of tasks. This may be too inefficient in practice. Alternatives exist where a bounded number of deadline misses may be tolerated. How should LET exceptions (violations of the specified LET) be handled then? How can deadline miss patterns which still guarantee functional correctness (e.g., system stability) be established?
- How should the LET constructs be integrated into AUTOSAR? More generally, how should the design and verification process in the automotive industry be modified to integrate the LET paradigm?
- How does the use of the LET paradigm for multicore systems fit into the more general context of achieving predictability of multicore systems?

This seminar provided a unique opportunity for participants from the automotive industry to get feedback from academia on their effort to adopt the LET paradigm. At the same time, it allowed other participants to confront their own models and/or solutions with industrial reality and identify new research challenges. This seminar furthermore brought together research communities which do not so often interact with each other, e.g. the synchronous, control and real-time communities.

■ Organization of the seminar

The seminar took place from 25th to 28th February 2018. The first day started with an introduction by the organizers, followed by a talk from one of the co-founders of the LET paradigm – Christoph Kirsch. The following two sessions included talks

providing an industrial view on the challenges of implementing LET in the multi-core automotive setting. The first day continued with a session comprised of talks presenting the academic view on LET-related challenges, and concluded with breakout sessions (detailed below). The second day of the seminar started with two sessions in which LET was compared to related paradigms, such as the synchronous model. The afternoon talks focused on the connection between LET and control as well as on a possible application of the LET approach to the domain of graphical processing units. The second day concluded with another set of breakout sessions. The third day included talks exploring future directions of LET, and a final set of breakout sessions.

Breakout sessions led to very interesting and fruitful discussions, and covered, among others, the following aspects:

- **Dimensioning of LET intervals:** The main focus was on how to efficiently dimension LET intervals to fit specific applications, which is currently a very pragmatic and experience based activity. Moreover, the two uses of LET in the automotive setting were identified: (i) Functional LET and (ii) Implementation LET.
- **Buffer optimization within LET:** The main focus was on the management of buffers in a LET-based implementation. The following topics were identified as relevant and thus discussed: minimizing the number of used buffers, strategies to handle memory contentions when accessing buffers, location of buffers in the memory hierarchy of hardware platforms and locality affinities between buffers, impact of spatial partitioning or periodicity of LET frames (harmonic or not) the buffers.
- **The synchronous approach vs LET:** The focus was on the comparison between the synchronous and LET models, with a discussion of their advantages and limitations, and their positioning in the context of the needs of the automotive industry, with a special emphasis on a transition from a singlecore to a multicore setting.
- **Control and LET:** The main focus was on the use of the LET paradigm to implement controllers. The following topics were identified as relevant and thus discussed: Is LET the correct paradigm for controller implementation? What is a viable period choice? How are potential deadline misses handled? Can a proper fault model be conveniently incorporated into the LET methodology? Can LET lead to new contributions in the control research domain?

More details on breakout sessions are available in a dedicated section of this document, after the overview of the talks given during the seminar.

■ Outcome of the seminar

The seminar has already enabled several collaborations: (i) a white paper on the topic is under preparation; (ii) a special session at EMSOFT'18 will be proposed. In addition, since participants expressed very positive opinions about the seminar and were in favor of reproducing the experience, a follow-up seminar will be considered.

Finally, as organizers, we would like to thank all of the participants for their strong interaction, interesting talks, fruitful group discussions, and work on open problems.

6.12 Scheduling

Organizers: Magnús M. Halldórsson, Nicole Megow, and Clifford Stein
Seminar No. 18101

Date: March 4–9, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.1

© Creative Commons BY 3.0 Unported license

© Magnús M. Halldórsson, Nicole Megow, and Clifford Stein



Participants: Susanne Albers, Yossi Azar, Nikhil Bansal, Sanjoy K. Baruah, Syamantak Das, Jelena Diakonikolas, Michael Dinitz, Fritz Eisenbrand, Thomas Erlebach, Klaus-Tycho Foerster, Seth Gilbert, Anupam Gupta, Magnús M. Halldórsson, Sungjin Im, Thomas Kesselheim, Samir Khuller, Valerie King, Fabian Daniel Kuhn, Amit Kumar, Alberto Marchetti-Spaccamela, Nicole Megow, Danupon Nanongkai, Yvonne-Anne Pignolet, Kirk Pruhs, Thomas Rothvoss, Barna Saha, Jared Saia, Gabriel Scalosub, Christian Scheideler, Kevin Schewior, Jiri Sgall, David Shmoys, René Sitters, Martin Skutella, Frits C. R. Spieksma, Clifford Stein, Leen Stougie, Jukka Suomela, Ola Svensson, Marc Uetz, Rob van Stee, Anke van Zuylen, Jose Verschae, Tjark Vredeveld

This fifth meeting in a series of Dagstuhl “Scheduling” seminars brought together part of the community of algorithmic researchers who focus on scheduling, and part of the community of algorithmic researchers who focus on networking in general, and resource management within networks in particular. These communities are far from unknown to each other as they attend the same general academic conferences. But as each community has its own specialized conferences, there is less interaction between these communities than there should be. Further there are differences in the types of algorithmic problems these communities are naturally drawn towards.

The primary objective of the seminar was to expose each community to the important models, problems and techniques from the other community, and to facilitate dialog and collaboration between researchers. The program included 22 invited main talks including an inspiring talk on practical applications at ABB Corporate Research, 8 short spot-light talks, two open problem sessions in the beginning of the week, and ample unstructured time for research and interaction. The overall atmosphere among the 44 participants was very interactive.

A highlight of the seminar was a joint Wednesday-session with the Dagstuhl Seminar 18102 “Dynamic Traffic Models in Transportation Science”. It was a fortunate coincidence that both seminars were scheduled in parallel. Indeed, questions related to networks, scheduling and resource sharing arise naturally in traffic control and transportation science. It was an inspiring secondary outcome of the workshop to realize this strong overlap in interests which led to interesting discussions between researchers of the different communities.

6.13 Dynamic Traffic Models in Transportation Science

Organizers: Roberto Cominetti, Tobias Harks, Carolina Osorio, and Britta Peis
Seminar No. 18102

Date: March 4–9, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.21

© Creative Commons BY 3.0 Unported license

© Tobias Harks, Roberto Cominetti, Carolina Osorio, and Britta Peis

Participants: Umang Bhaskar, Roberto Cominetti, Gunnar Flötteröd, Martin Gairing, Cristóbal Guzmán, Tobias Harks, Martin Hofer, Anja Huber, Max Klimm, Ekkehard Köhler, Kai Nagel, Neil Olver, Carolina Osorio, Britta Peis, Rahul Savani, Marco Scarsini, Guido Schäfer, Heiko Schilling, Miriam Schlöter, Daniel Schmand, Marc Schröder, Alexander Skopalik, Nicolás Stier-Moses, Sebastian Stiller, Martin Strehler, Chris Tampère, Theresa Thunig, Veerle Timmermans, Laura Vargas-Koch, Bernhard von Stengel, Dave Watling



Traffic assignment models play an important role for traffic planners to predict traffic distributions, especially, in light of possible changes of the infrastructure, e.g., road constructions, traffic light controls, speed limits, tolls, etc. The prevailing *mathematical* approaches used in the transportation science literature to predict such distributions can be roughly classified into static traffic assignment models based on aggregated static multi-commodity flow formulations and dynamic traffic assignment (DTA) models based on the methodology of flows over time. While static models have seen several decades of development and practical use, they abstract away too many important details and, thus, become less attractive. On the other hand, dynamic models are known to be notoriously hard to analyze in terms of existence, uniqueness and computability of dynamic equilibria.

In light of the prevailing computational difficulties for realistic-sized networks, the systematic optimization of such networks (e.g., by designing the network infrastructure, link tolls, or traffic light controls) becomes even more challenging as the resulting mathematical programs with equilibrium constraints contain already in the lower level presumably “hard” optimization-, complementarity- or variational inequality problems; not to speak of the resulting optimization problem for the first level.

On the other hand, there is a trend in the transportation science community to use *large-scale computer-based microsimulations* for predicting traffic distributions. The striking advantage of microscopic simulations over DTA models is that the latter usually ignores the feedback of changing network conditions on user behavior dimensions such as flexible departure time choice, mode choice, activity schedule choice, and such. Current simulation tools integrate all these dimensions and many more. The increasing model complexity, however, is by far not matched by the existing theory of dynamic traffic assignments.

The seminar brought together leading researchers from three different communities – Simulations (SIM), Dynamic Traffic

Assignment (DTA) and Algorithmic Game Theory (AGT). This years seminar was centered around three topics:

- *Horizontal queueing models.* Most of the static traffic assignment models assume that queues can occur, but do not take up any physical space. In order to make the current models more realistic one should assume that queues might effect traffic on other nearby road segments, thus, include possible spill-back effects.
- *Oligopolistic competition.* With the rise of autonomous vehicles new routing decisions need to be made. As a novel aspect, individual vehicles might be interested in selfishly optimizing their routes, but cooperate with other vehicles using the same software in order to decrease the average journey time.
- *Risk-averse travelers.* Current static traffic models often assume that each player is rational, and has the sole purpose of minimizing travel time or distance. However, the exact travel time of many routes might be uncertain at the moment of departure. Hence, travelers might stick to a more predictable route and might be unwilling to explore possibly better alternatives.

Again, the seminar was a big success both in terms of stimulating new and very fruitful collaborations. We got enthusiastic feedback from many participants which is also reflected in the survey conducted by Dagstuhl.

6.14 Loop Optimization

Organizers: Sebastian Hack, Paul H. J. Kelly and Christian Lengauer
Seminar No. 18111

Date: March 11–16, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.39

© Creative Commons BY 3.0 Unported license

© Sebastian Hack, Paul H. J. Kelly and Christian Lengauer



Participants: Cédric Bastoul, Barbara M. Chapman, Shigeru Chiba, Charisee Chiw, Philippe Clauss, Albert Cohen, James W. Demmel, Johannes Doerfert, Andi Drebes, Paul Feautrier, Stefan Ganser, Armin Größlinger, Tobias Grosse, Sebastian Hack, Julian Hammer, Frank Hannig, Alexandra Jimborean, Paul H. J. Kelly, Sriram Krishnamoorthy, Michael Kruse, Roland Leißa, Christian Lengauer, Fabio Luporini, Benoit Meister, Lawrence Mitchell, Madan Musuvathi, Victor Nicolet, Philip Pfaffe, Antoniu Pop, Louis-Noël Pouchet, Jonathan Ragan-Kelley, P. (Saday) Sadayappan, Jun Shirako, Andreas Simbürger, Daniele G. Spampinato, Michel Steuwer, Tianjiao Sun, Nicolas Vasilache, Richard M. Veras, Sven Verdoolaege, Ayal Zaks

■ Motivation

Loop optimization is at the heart of effective program optimization – even if the source language is too abstract to contain loop constructs explicitly as, e.g., in a functional style or a domain-specific language. Loops provide a major opportunity to improve the performance of a program because they represent compactly a large volume of accessed data and executed instructions. Because the clock frequency of processors fails to continue to grow (end of Dennard scaling), the only way in which the execution of programs can be accelerated is by increasing their throughput with a compiler: by increasing parallelism and improving data locality. This puts loop optimization in the center of performance optimization.

■ Context

The quick and easy way to optimize a loop nest, still frequently used in practice, is by restructuring the source program, e.g., by permuting, tiling or skewing the loop nest. Beside being laborious and error-prone, this approach favors modifications that can be easily recognized and carried out, but which need not be the most suitable choice. A much better approach is to search automatically for optimization options in a mathematical model of the iteration space, in which all options are equally detectable and the quality of each option can be assessed precisely.

Recently, the polyhedral compilation community has produced a set of robust and powerful libraries that contain a variety of algorithms for the manipulation of Presburger sets, including all standard polyhedral compilation techniques. They can be incorporated in a program analysis to make other compiler optimizations more precise and powerful, like optimizers and code generators for domain-specific languages, or aggressive optimizers for high-performance computing.

Polyhedral loop optimization relies on strict constraints on

the structure of the loop nest and may incur a computationally complex program analysis, based on integer linear programming. The optimization problems become much simpler when information at load or run time is available, i.e., the optimization is done just-in-time. Also, the search for the best optimization can be supported by other techniques, e.g., auto-tuning, machine learning or genetic algorithms. While these techniques are all fully automatic, engineering of software with robust performance characteristics requires programmers to have some level of explicit control over the data distribution and communication costs. However, manually optimized code is far too complicated to maintain. Thus, a major research area concerns the design of tools that allow developers to guide or direct analysis (e.g., via dependence summaries or domain-specific code generation) and optimization (e.g., via directives, sketches and abstractions for schedules and data partitioning).

■ Goal

The goal of this seminar was to generate a new synergy in loop optimization research by bringing together representatives of the major different schools of thought in this field. The key unifying idea is to formulate loop optimization as a mathematical problem, by characterizing the optimization space and objectives with respect to a suitable model.

One school is focused on reasoning about scheduling and parallelization using a geometric, “polyhedral”, model of iteration spaces which supports powerful tools for measuring parallelism, locality and communication – but which is quite limited in its applicability.

Another major school treats program optimization as program synthesis, for example by equational rewriting, generating a potentially large space of variants which can be pruned with respect to properties like load balance and locality. This approach

has flourished in certain application domains, but also suffers from problems with generalization.

A third family of loop optimization approaches tackles program optimization through program generation and symbolic evaluation. Generative approaches, such as explicit staging, support programmers in taking explicit control over implementation details at a high level of abstraction.

The seminar explored the interplay of these various loop optimization techniques and fostered the communication in the wide-ranging research community of model-based loop optimization. Participants represented the various loop optimization approaches but also application domains in high-performance computing.

■ Conclusions

The seminar succeeded in making the participants aware of common goals and relations between different approaches. Consensus emerged on the potential and importance of tensor contractions and tensor comprehensions as an intermediate representation. There was also some excitement in connecting the classical dependence-based optimization with newly emerging ideas in deriving parallel algorithms from sequentially-dependent code automatically. Guided automatic search and inference turned out to be a dominant theme. Another important insight was that the optimization criteria currently in use are often too coarse-grained and do not deliver satisfactory performance. More precise hardware models are needed to guide optimization. This will require a closer collaboration with the performance modeling and engineering community.

It was agreed that publications and collaborations fueled by the seminar will acknowledge Schloss Dagstuhl.

6.15 Coding Theory for Inference, Learning and Optimization

Organizers: Po-Ling Loh, Arya Mazumdar, Dimitris Papailiopoulos, and Rüdiger Urbanke
Seminar No. 18112

Date: March 11–16, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.60

© Creative Commons BY 3.0 Unported license

© Po-Ling Loh, Arya Mazumdar, Dimitris Papailiopoulos, and Rüdiger Urbanke



Participants: Dimitris Achlioptas, Alexander Barg, Martin Bossert, Elette Boyle, Amin Coja-Oghlan, Anna Gál, Venkatesan Guruswami, Hamed S. Hassani, Sihuang Hu, Sidharth Jaggi, Marc Lelarge, Po-Ling Loh, Nicolas Macris, Arya Mazumdar, Olgica Milenkovic, Dimitris Papailiopoulos, Ankit Singh Rawat, Changho Suh, Itzhak Tamo, Rüdiger Urbanke, Pascal Vontobel, Mary Wootters

Codes are widely used in engineering applications to offer reliability and fault tolerance. The high-level idea of coding is to exploit redundancy in order to create robustness against system noise. The theoretical properties of codes have been studied for decades both from a purely mathematical point of view, as well as in various engineering contexts. The latter have resulted in constructions that have been incorporated into our daily lives: No storage device, cell phone transmission, or Wi-Fi connection would be possible without well-constructed codes.

Recent research has connected concepts in coding theory to non-traditional applications in learning, computation and inference, where codes have been used to design more efficient inference algorithms and build robust, large-scale, distributed computational pipelines. Moreover, ideas derived from Shannon theory and the algebraic properties of random codes have resulted in novel research that sheds light on fundamental phase transition phenomena in several long-standing combinatorial and graph-theoretic problems.

The main goal of our seminar was to accelerate research in the growing field of coding theory for computation and learning, and maximize the transformative role of codes in non-traditional application areas. The seminar brought together 22 researchers from across the world specializing in information theory, machine learning, theoretical computer science, optimization, and statistics. The schedule for each day included a tutorial talk by a senior researcher, followed by shorter talks by participants on recent or ongoing work. The afternoons were devoted to informal breakout sessions for groups to discuss open questions. Two of the larger breakout sessions focused on distributed optimization and group testing.

Seminar participants reported that they enjoyed hearing about new ideas, as well as delving into deeper technical discussions about open problems in coding theory. Some topics deserving special mention include the use of techniques in statistical

mechanics; locally decodable and recoverable codes; submodular function optimization; hypergraph clustering; private information retrieval; and contagion on graphs. All participants valued the ample time for discussions between and after talks, as it provided a fruitful atmosphere for collaborating on new topics.

6.16 Machine Learning and Model Checking Join Forces

Organizers: Nils Jansen, Joost-Pieter Katoen, Pushmeet Kohli, and Jan Kretinsky
Seminar No. 18121

Date: March 18–23, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.74

© Creative Commons BY 3.0 Unported license

© Nils Jansen, Joost-Pieter Katoen, Pushmeet Kohli, and Jan Kretinsky

Participants: Alessandro Abate, Erika Abraham, Ezio Bartocci, Roderick Bloem, Luca Bortolussi, Tomáš Brázdil, Marc Brockschmidt, Rudy Bunel, Michael Carbin, Rayna Dimitrova, Krishnamurthy Dvijotham, Rüdiger Ehlers, Andreas Berre Eriksen, Radu Grosu, Arnd Hartmanns, Laura Humphrey, Manfred Jaeger, Nils Jansen, Sebastian Junges, Joost-Pieter Katoen, Pushmeet Kohli, Jan Kretinsky, Kim Guldstrand Larsen, Alexis Linard, Tobias Meggendorfer, Daniel Neider, Guillermo A. Pérez, Ruzica Piskac, Hasan Poonawala, Pavithra Prabhakar, Jean-Francois Raskin, Guido Sanguinetti, Daniel Selsam, Sanjit A. Seshia, Armando Solar-Lezama, Ufuk Topcu, Jana Tumova, Jonathan Uesato, Frits Vaandrager, Min Wen, Leonore Winterer



This Dagstuhl Seminar aimed at bringing together researchers working in the fields of machine learning and model checking. Growing application areas for machine learning, such as autonomous driving, require the exclusion or likely avoidance of unsafe behaviors. An important question is then, how confidence in system behaviors obtained from machine learning can be transferred to formal verification. Vice versa, industrial usage of model checking still suffers from scalability issues for large applications. Leveraging the capabilities of machine learning to assess large data sets will help to enable the verification for more realistic systems.

Based on the concrete discussions and inputs from all the participants, we identified the following topics as great challenges to the combination of the fields of machine learning and model checking.

- Safety Verification of Deep Neural Networks
- Formal Program Synthesis and Analysis using Machine Learning
- Representation of Strategies and Controllers
- Explainable Artificial Intelligence
- Challenges for Machine Learning in Motion Planning
- Guarantees on Reinforcement Learning in Verification
- Social and Legal Issues in Artificial Intelligence
- Exploiting Weaknesses in Reinforcement Learning

6.17 Automatic Quality Assurance and Release

Organizers: Bram Adams, Benoit Baudry, Sigrid Eldh, and Andy Zaidman
Seminar No. 18122

Date: March 18–21, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.3.94

© Creative Commons BY 3.0 Unported license

© Bram Adams, Benoit Baudry, Sigrid Eldh, and Andy Zaidman



Participants: Bram Adams, Benoit Baudry, Moritz Beller, Benjamin Danglot, Tamara Domic, Sigrid Eldh, Daniele Gagliardi, Georgios Gousios, Zhen Ming (Jack) Jiang, Foutse Khomh, Philipp Leitner, Lucy Ellen Lwakatare, Vincent Massol, Shane McIntosh, Martin Monperrus, Sarah Nadi, Andrew Neitsch, Christopher J. Parnin, Gerald Schermann, Weiyi (Ian) Shang, Hui Song, Oscar Luis Vera Perez, Hyrum K. Wright, Andy Zaidman, Fiorella Zampetti

The seminar explored the relationship between DevOps and quality assurance from a software engineering perspective. DevOps has been gaining traction since around 2012, with initiatives formed both in industry and academia. While the importance of DevOps as an enabler in higher quality software is intuitively clear to both industry and academia, we have discussed commonalities in views, but also the challenges that lie ahead for this discipline.

In essence, human factors are very important, because DevOps is not only a technology, it is a way of working and organizing teams. In this light, we have also discussed the resistance that some team members or even entire organisations seem to have towards automating quality assurance through DevOps. Section 4.2 of the full report summarizes a group discussion that eventually triggered a set of reflections on this topic of human aspects of DevOps. Yet, we have also discussed how DevOps can be an enabler for onboarding new team members through the availability of a standardized DevOps infrastructure (Section 4.4 of the full report). The whole group observed the general lack of empirical evidence on the importance and benefits of DevOps in modern software engineering. This final point is tightly connected to another important theme in our discussion: educating software engineers in the ways and associated technologies of DevOps.

The main goal of this seminar was to bridge the knowledge divide on how researchers and industry professionals reason about and implement DevOps for automatic quality assurance. Through the seminar, we have built up a common understanding of DevOps tools and practices, but we have also identified major challenges for this field of research as well as for the teaching of DevOps principles and practices.

This Dagstuhl was a 2.5 day seminar, which we structured around 4 invited talks that served as keynotes to introduce key topics for discussions. These talks, summarized in Sections 3.1 through 3.3 of the full report, were given at the beginning of each

morning and afternoon to inspire topics for further discussions on a given topic. The group split into smaller sub-groups after each keynote, in order to focus discussions and reflections on a specific topic. All these discussions have been summarized in the form of a blog post, while in Dagstuhl, and are provided in this report.

In addition to keynotes and subgroup discussions, we had a plenary session to start the seminar, where each participant had 2 slides for a short introduction; we had a “speed-dating” session on Tuesday evening; and we organized a panel discussion about the future of the field on the last morning (Section 6.3 of the full report).

6.18 Program Equivalence

Organizers: Shuvendu K. Lahiri, Andrzej Murawski, Ofer Strichman, and Mattias Ulbrich
Seminar No. 18151

Date: April 8–13, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.1

© Creative Commons BY 3.0 Unported license

© Shuvendu K. Lahiri, Andrzej Murawski, Ofer Strichman, and Mattias Ulbrich

Participants: Anindya Banerjee, Gilles Barthe, Nick Benton, Dirk Beyer, Soham Chakraborty, Stefan Ciobaca, Constantin Enea, Grigory Fedyukovich, Dan R. Ghica, Arie Gurfinkel, Guilhem Jaber, Vasileios Koutavas, Steve Kremer, Shuvendu K. Lahiri, James Laird, Xavier Leroy, Yi Li, Sergey Mechtaev, Andrzej Murawski, Kedar Namjoshi, David A. Naumann, Julia Rubin, Philipp Rümmer, Neha Rungta, Chaked Saydoff, Rahul Sharma, Stephen Siegel, Marcelo Sousa, Ofer Strichman, Aaron Tomb, Nikos Tzevelekos, Mattias Ulbrich, Niels Voorneveld



Program equivalence is arguably one of the most interesting and at the same time important problems in formal verification. It has attracted the interest of several communities, ranging from the field of denotational semantics and the problem of Full Abstraction, to software verification and Regression Testing. The aim of this meeting was to bring together the different approaches and techniques of the current state of the art and to facilitate the cross-pollination of research between these areas.

This interdisciplinary community met once before in the workshop on program equivalence in London (April 2016). There was a general agreement among the participants that a research community around this topic should be established in the form of a workshop and eventually a conference, and that the interest in this topic continuously grows around the world, including a growing interest in the industry. Furthermore, currently there is little overlap in the conferences that some of the key players attend, to the point that many participants were little aware of other participants' work.

We were happy to witness that indeed participants learned greatly from this week, collaborations were established, and cross fertilization between the communities occurred. We hope to meet again in Dagstuhl in the future!

6.19 Blockchains, Smart Contracts and Future Applications

Organizers: Foteini Baldimtsi, Stefan Katzenbeisser, Volkmar Lotz, and Edgar Weippl
Seminar No. 18152

Date: April 8–13, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.20

© Creative Commons BY 3.0 Unported license

© Edgar Weippl, Foteini Baldimtsi, Stefan Katzenbeisser, and Volkmar Lotz



Participants: Zohar Aviv, Foteini Baldimtsi, Alex Biryukov, Rainer Böhme, Jan Camenisch, Samuel Christie, Ittay Eyal, Sebastian Faust, Peter Gazi, Dieter Gollmann, Raimund Gross, Bernhard Haslhofer, Aljosha Judmayer, Stefan Katzenbeisser, Kwok-Yan Lam, Juho Lindman, Volkmar Lotz, Sarah Meiklejohn, Bart Preneel, Alessandra Scafuro, Philipp Schindler, Sofie Schock, Nicholas Stifter, Thorsten Strufe, Edgar Weippl, Alexei Zamyatin

In its beginnings, the technical and socio-economical feasibility of Bitcoin was met with much skepticism; however, this has since changed as both research and practice have outlined the merits of distributed ledger technologies, commonly referred to as “blockchains”. Possible applications of blockchains reach from decentralized settlement layers over complex smart contract systems to tailored authenticated data structures that implement systems for identity or supply chain management. Nevertheless, beyond the immediate opportunities and applications lie many open questions regarding the long-term perspective of both permissionless and permissioned blockchain technologies. For example, while scalability and sustainability are currently topics of active research and development, other aspects such as usability, interoperability and cryptoeconomics have received considerably less attention. In order to anticipate and address future key topics and questions related to blockchain technologies, this seminar strove to provide an interdisciplinary breeding ground.

The participants focused on future applications and developments of this technology and discussed how such complex systems can thrive over a long period of time. Thereby, we started our seminar by outlining and collecting current and potentially future issues from the diverse viewpoints of the participants. These issues include not only current limitations of the underlying technologies, but also problems encountered in real-world applications.

As an example, we considered the various economic, legal and technological uncertainties and problems that have arisen as a consequence of the recent contentious forks in both the Bitcoin (August 2017) and Ethereum (July 2016) networks. While the possibility of such forks was previously well known, it can be argued that provisional measures and research on effectively dealing with them was immature and could have been addressed much sooner. In any case, the ramifications of these events have

and will continue to influence the discussion and development of blockchain technologies.

Beside establishing the relevant issues through numerous talks, subgroups of participants were formed to discuss a specific set of topics. Over the course of the seminar, participants were encouraged to move between groups and provide input to various topics. We hope to have thus enriched the discussion with different viewpoints and to have facilitated a rewarding range of outcomes; at the point of writing, two papers directly resulting from this Dagstuhl seminar are submitted for review. The goal of the seminar was to develop a shared and open agenda that shapes and directs research and development in the area of distributed ledger technologies to face current and future challenges as well as contribute to the positive development of this field.

The talks and working groups of this first Dagstuhl seminar on Blockchains, Smart Contracts and their future applications focused inter alia on the following topics:

- current and future protocols, including alternative consensus protocols
- governance
- interdisciplinary aspects of Blockchain technology (economy, law)
- cross-chain communication
- scalability and costs
- Goldfinger and other attack vectors



Fig. 6.4

“View from my room #MINLPDagstuhl @dagstuhl” Twitter post by 18081 Dagstuhl Seminar participant Matteo Fischetti.

<https://twitter.com/MFischetti/status/965569053287964672>. Photo courtesy of Matteo Fischetti.

6.20 Visualization of Biological Data – Crossroads

Organizers: Jan Aerts, Nils Gehlenborg, Georgeta Elisabeta Marai, and Kay Katja Nieselt
Seminar No. 18161

Date: April 15–20, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.32

© Creative Commons BY 3.0 Unported license

© Jan Aerts, Nils Gehlenborg, Georgeta Elisabeta Marai, and Kay Katja Nieselt



Participants: Jan Aerts, Katja Bühler, Sheelagh Carpendale, James L. Chen, Arlene Chung, Anamaria Crisan, Mirjam Figaschewski, Angus Forbes, Nils Gehlenborg, Carsten Görg, David H. Gotz, Helena Jambor, Jessie Kennedy, Karsten Klein, Anne Knudsen, Barbora Kozlíková, Michael Krone, Martin Krzywinski, Alexander Lex, Raghu Machiraju, Georgeta Elisabeta Marai, Lennart Martens, Ewy Mathé, Torsten Möller, Scooter Morris, Cydney Nielsen, Kay Katja Nieselt, Bruno Pinaud, James Procter, William Ray, Jens Rittscher, Jos B.T.M. Roerdink, Timo Ropinski, Ryo Sakai, Falk Schreiber, Christian Stolte, Marc Streit, Granger Sutton, Danielle Szafir, Cagatay Turkey, Michel A. Westenberg, Blaz Zupan

The rapidly expanding application of experimental high-throughput and high-resolution methods in biology is creating enormous challenges for the visualization of biological data. To meet these challenges, a large variety of expertise from the visualization, bioinformatics and biology domains is required. These encompass visualization and design knowledge, algorithm design, strong implementation skills for analyzing and visualizing big data, statistical knowledge, and specific domain knowledge for different application problems. In particular, it is of increasing importance to develop powerful and integrative visualization methods combined with computational analytical methods. Furthermore, because of the growing relevance of visualization for bioinformatics, teaching visualization should also become part of the bioinformatics curriculum.

With this Dagstuhl Seminar we wanted to continue the process of community building across the disciplines of biology, bioinformatics, and visualization. We aim to bring together researchers from the different domains to discuss how to continue the BioVis interdisciplinary dialogue, to foster the development of an international community, to discuss the state-of-the-art and identify areas of research that might benefit from joint efforts of all groups involved.

Based on the topics identified in the seminar proposal, as well as the interest and expertise of the confirmed participants, the following four topics were chosen as focus areas for the seminar, in addition to the overarching topic of collaboration between the data visualization, bioinformatics, and biology communities:

Visualization challenges related to high-dimensional medical data. Patient data is increasingly available in many forms including genomic, transcriptomic, epigenetic, proteomic, histologic, radiologic, and clinical, resulting in large (100s of TBs, 1000s of patients), heterogeneous (dozens of data types per patient) data repositories. Repositories such as The Cancer Genome Atlas (TCGA) contain a multitude of patient

records which can be used for patient stratification, for high-risk group and response to treatment discoveries, or for disease subtype/biomarker discoveries. Still, patient records from the clinic are used singularly to diagnose patients in the clinic without including likely insights from other sources. Similarly, molecular expression signatures from the omic sources barely impinge on the clinical observations. There is an urgent need to bridge the divide the precision medicine gap between the laboratory and the clinic, as well as a need to bridge the quantitative sciences with biology. Additionally, many precision medicine studies plan to include sensor data (e.g. physical activity, sleep, and other patient-worn sensors) that will add another dimension of complexity that analysis and visualization tools need to take into account.

This highly relevant topic focused on visual analytic tools and collaborations that will promote and leverage notions of patient similarity across the phenotypical scales. Scalable and robust machine learning methods will need to work synergistically to integrate evidence of similarity while meaningful visual encodings should simultaneously summarize and illuminate patient similitude at the individual and group level. This topic is closely related to some of the topics below.

Visualization of biological networks. Modeling the stochasticity of genetic circuits is an important field of research in systems biology, and can help elucidate the mechanisms of cell behavior, which in turn can be the basis of diseases. These models can further enable predictions of important phenotypic cellular states. However, the analysis of stochastic probability distributions is difficult due to their spatiotemporal and multidimensional nature, and due to the typically large number of simulations run under varying settings. Moreover, stochastic network researchers often emphasize that what is of biological significance is often not of statistical significance – numerical analyses often miss small or rare events of particular biological relevance. A visual approach

can help, in contrast, in mining the network dynamics through the landscape defined by these probability distributions.

Another major challenge relates to finding “stable behavior” of networks, including those recruited in signal transduction. Multistability and bistability have been often studied in metabolic chemically reactive networks. Necessary conditions have been formulated to imply the emergence of stable phenotypes. However, these methods have been deployed on small networks. Recently many groups have recognized that scalable methods can be explored using steady state or quasi steady state models that are derived from stoichiometry and rate-action kinetics. These unfortunately suffer from the lack of methods that will examine the large parametric space. Consider this: N interacting molecules imply N^2 interactions and in turn the same order of the governing “parameters” (activation rates and abundances). For even mid-size portions of salient pathways (EGFR, B-cell Receptor activation, etc.) finding stable states is challenging. It is certainly the case that a complete graph is never realized and sparsity and network mining can be used to glean the necessary structure. Design of experiments followed by visualization of parametric spaces will be required to search for these stable points. Furthermore, the huge size of this space needs possibly new scalable approaches for the visualization.

Visualization for pan-genomics. With the advent of next-generation sequencing we can observe the increase of genome data both in the field of metagenomics (simultaneous assessment of many species) as well as within the field of pan-genomics. In metagenomics, the aim is to understand the composition and operation of complex microbial consortia in environmental samples. On the other hand in pangenomics genomes within a species are studied. While originally a pan-genome has been referred to as the full complement of genes in a clade (mainly a species in bacteria or archaea), this has recently been generalized to considering a pan-genome as any collection of genomic sequences to be analyzed jointly or to be used as a reference rather than a single genome.

In bioinformatics, both topics impose a number of computational challenges. For example, a recent review paper by Marschall et al. on “Computational Pan-Genomics: Status, Promises and Challenges” (DOI: 10.1093/bib/bbw089) addresses current efforts in this sub-area of bioinformatics. This area needs novel, qualitatively different computational methods and paradigms. While the development of new promising computational methods and new data structures both in metagenomics and pangenomics can be observed, a number of open challenges exist. One of them in the area of pangenomics is for example the transition from the representation of reference genomes as strings to representations as graphs. However, the important topic of pangenome visualization has not been addressed in the aforementioned review. Interestingly this has been taken up in a break-out session in a recent Dagstuhl seminar on “Next Generation Sequencing - Algorithms, and Software For Biomedical Applications”, and identified as a topic of urgent interest and demand. One observation for example is that in pan-genomes there are segments of conserved regions interspersed by highly variable regions. Open question here is how to visualize the highly variable regions, or how to interpret its content in the context of its neighborhood. Other open visualization topics involve the visual representation of the graph structure underlying pangenomes.

In the field of metagenomics some common visualization approaches, such as heatmaps or scatter plots in combination with principal component analyses, are used, however, many open challenges exist. In particular those visualization tools that are developed for genomics studies fall short in representing

large-scale, high dimensional metagenomics studies. Especially the magnitude of the data presents a challenge to meaningfully represent biologically valuable information from complex analysis results. Thus also in this topic the question of large-scale and heterogeneous data visualization is of central importance.

Curriculum development of biological data visualization. Parallel to the recognized need to teach bioinformatics students about big data in biology, there is a growing need to familiarise students with modern visual analytics methodologies applied to biological data, and to provide hands-on training. While several community members are teaching summer camps, tutorials, and workshops on biological data visualization, many of these educational sessions take the form of an introduction to specific tools. We find ourselves handling similar questions: what is exploratory data visualization, what is visual analytics, which frameworks to think about visualization exist, how can we explore design space, and how can we visualise biological data to gain insight into them, so that hypotheses can be generated or explored and further targeted analyses can be defined?

Despite the increasing importance of visualization for bioinformatics, there is currently a general lack of integration into the bioinformatics education, and a useful and appropriate curriculum has not yet been developed. In this topic the following questions will be addressed: What should a modern and seminal curriculum for visualization in bioinformatics look like? How far along the introductory visualization courses should this curriculum go, while allowing biological data topics as well? What are the essential topics, and how can comprehensive training be achieved?

The schedule for the seminar was developed by the organizers based on previous successful Dagstuhl seminars. Emphasis was given to a balance between prepared talks and panels and break groups for less structured discussions focused on a selection of highly relevant topics. Three types of plenary presentations were available to participants who had indicated interest in presenting during the seminar: overview talks (20 minutes plus 10 minutes for questions), regular talks (10 minutes plus 5 minutes for questions), and panel presentations (5 minutes per speaker followed by a 20 – 25 minute discussion). The break out groups met multiple times for several hours during the week and reported back to the overall group on several occasions. This format successfully brought bioinformatics and visualization researchers onto the same platform, and enabled researchers to reach a common, deep understanding through their questions and answers. It also stimulated very long, intense, and fruitful discussions that were deeply appreciated by all participants.

This report describes in detail the outcomes of this meeting. Our outcomes include a set of white papers summarizing the breakout sessions, overviews of the talks, and a detailed curriculum for biological data visualization courses.

6.21 Normative Multi-Agent Systems

Organizers: Mehdi Dastani, Jürgen Dix, Harko Verhagen, and Serena Villata
Seminar No. 18171

Date: April 22–27, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.72

© Creative Commons BY 3.0 Unported license

© Mehdi Dastani, Jürgen Dix, and Harko Verhagen



Participants: Tobias Ahlbrecht, Natasha Alechina, Kevin D. Ashley, Matteo Baldoni, Christoph Benzmüller, Célia da Costa Pereira, Mehdi Dastani, Tiago de Lima, Davide Dell'Anna, Jürgen Dix, Ali Farjami, Dov M. Gabbay, Aditya K. Ghose, Matthias Grabmair, Joris Hulstijn, Wojtek Jamroga, Özgür Kafali, Sabrina Kirrane, Brian Logan, Emiliano Lorini, Martin Neumann, Pablo Noriega, Julian Padget, Adrian Paschke, Nicolas Payette, Ken Satoh, Matthias Scheutz, Viviane Torres da Silva, Leon van der Torre, Harko Verhagen, Douglas Walton, Michael Winikoff, Vahid Yazdanpanah

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

The seminar was a blend of talks, discussions and group work. It began on the first day with short “teaser talks” (10 + 5 minutes) related to the main topic of *norms and responsibility*, one given by almost each participant. The talks were meant to be inspiring and thought-provoking, channeling ideas for the following days. While some missed the established procedure with longer talks, the new format was overall very well received and allowed for many different thoughts and concepts to be presented and discussed in relatively short time.

Four working groups formed at the end of the first day for the norm-related topics responsibility, new logics, ethics/values and (machine) learning.

The aim of the group sessions, on the second and fourth day, was to get a shared understanding of the specific topics and to identify future research possibilities. Each group reported back in a plenary session at the end of each group work day, where the groups also tried to establish interconnections between them.

Responsibility. This group discussed how to grasp the very abstract concept of responsibility. A big chunk was dedicated to the formalization of responsibility. Many (vastly different) assumptions were laid out. The problem of “delegating responsibility” was discussed with special intensity. The group (being by far the largest one) split later to discuss different notions of responsibility on the basis of selected examples. A working paper was produced, included in Section 4.1 of the full report.

New logics. The aim of this group was to find out how to tackle norms and responsibility in terms of logics, especially how new logics for this task could be devised.

Ethics/values. This group discussed the more ethics-oriented

aspects of normative systems. Values provide an additional layer for normative reasoning: e.g. “how acceptable is it to violate a given norm?” The group produced a draft of a paper on “The Value(s) of Water” connecting NorMAS to the AI for Good initiative. Work is planned to continue during 2018 resulting in a paper for publication, e.g. in ACM communications or a similar outlet.

(Machine) Learning. The learning group discussed the opportunity of integrating norms and responsibility into machine learning procedures. As those are usually opaque, this presents as a notable challenge. For example, the learning’s input data has to be pre-processed to get a normatively acting system. Also, the learned sub-symbolic system should be enhanced with “regular” symbolic reasoning, which can be better regulated by norms and analysed for responsibility.

The fourth day was further enriched by a brainstorming session to identify possible applications. The subsequent clustering revealed the topics

- **transport**, e.g. smart grid/home, intelligent cars,
- **tools**, e.g. for autonomous service composition, legal reasoning, or supporting software/requirements engineering,
- **climate & agriculture**, e.g. agents negotiating fertilizer and water use, or an app that helps monitoring personal climate-affecting activities,
- **societies**, e.g. norms improving sustainability, monitoring of online forums for bad behavior or hate speech detection,
- **security**, e.g. protecting personal freedom by dynamically analysing normative consequences of law proposals, monitoring a company’s compliance with EU regulations, improving access to restricted access datasets, or making societies resilient for data surveillance by means of contract negotiations,
- **health**, e.g. ethical decision-making, norms for improving

personal health and fitness, defining wellbeing by norms, handling of patient/health data, and a big interest in healthcare robots,

- **energy**, e.g. modelling energy security with norms, managing air quality, observing long-term consequences, agents monitoring (personal) energy use to identify bad behavior, or regulating industrial relations or the energy and material footprint.

The application areas were discussed in a plenary session and formed the input to the discussion on future plans for the NorMAS community. Several conferences were identified to target proposals for a NorMAS-related workshop as part of the event. The community sees many relevant application areas not in the least in autonomous internet services and physical agents such as robots, vehicles and drones, where social reasoning will be of the utmost importance. Bringing the work from NorMAS to these areas will be highly beneficial to the involved communities.

6.22 Algebraic Effect Handlers go Mainstream

Organizers: Sivaramakrishnan Krishnamoorthy Chandrasekaran, Daan Leijen, Matija Pretnar, and Tom Schrijvers

Seminar No. 18172

Date: April 22–27, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.104

© Creative Commons BY 3.0 Unported license

© Sivaramakrishnan Krishnamoorthy Chandrasekaran, Daan Leijen, Matija Pretnar, and Tom Schrijvers



Participants: Amal Ahmed, Robert Atkey, Lennart Augustsson, Andrej Bauer, Oliver Bracevac, Jonathan Immanuel Brachthäuser, Edwin Brady, Stephen Dolan, Jeremy Gibbons, Daniel Hillerström, Mauro Jaskelioff, Ohad Kammar, Andrew Kennedy, Oleg Kiselyov, Sivaramakrishnan Krishnamoorthy Chandrasekaran, Daan Leijen, Sam Lindley, Andres Löh, Žiga Lukšič, Anil Madhavapeddy, Conor McBride, Adriaan Moors, Matija Pretnar, Andreas Rossberg, Tom Schrijvers, Perdita Stevens, Wouter Swierstra, Leo White, Nicolas Wu, Jeremy Yallop

Algebraic effects and their handlers have been steadily gaining attention as a programming language feature for composablely expressing user-defined computational effects. Algebraic effect handlers generalise many control-flow abstractions such as exception handling, iterators, `async/await`, or backtracking, and in turn allow them to be expressed as libraries rather than implementing them as primitives as many language implementations do. While several prototype languages that incorporate effect handlers exist, they have not yet been adopted into mainstream languages. This Dagstuhl Seminar 18172 “Algebraic Effect Handlers Go Mainstream” touched upon various topics that hinder adoption into mainstream languages. To this end, the participants in this seminar included a healthy mix of academics who study algebraic effects and handlers, and developers of mainstream languages such as Haskell, OCaml, Scala, WebAssembly, and Hack.

This seminar follows the earlier, wildly successful Dagstuhl Seminar 16112 “From Theory to Practice of Algebraic Effects and Handlers” which was dedicated to addressing fundamental issues in the theory and practice of algebraic effect handlers. We adopted a similar structure for this seminar. We had talks each day in the morning, scheduled a few days ahead. The folks from the industry were invited to present their perspectives on some of the challenges that could potentially be addressed with the help of effect handlers. The afternoons were left free for working in self-organised groups and show-and-tell sessions with results from the previous days. We also had impromptu lectures on the origins of algebraic effects and handlers, which were quite well received and one of the highlights of the seminar.

Between the lectures and working-in-groups, the afternoons were rather full. Hence, a few participants offered after-dinner “cheesy talks” just after the cheese was served in the evening. The participants were treated to entertaining talks over delightful cheese and fine wine. We encourage the organisers to leave part of the day unplanned and go with what the participants feel like

doing on that day. The serendipitous success are what makes Dagstuhl Seminars special.

We are delighted with the outcome of the seminar. There were interesting discussions around the problem of *encapsulation* and leaking of effects in certain higher order use cases, with several promising solutions discussed. It was identified that the problem of encapsulation and leaking effect names is analogous to the name binding in lambda calculus. Another group made significant progress in extending WebAssembly with support for effect handlers. The proposal builds on top of support for exceptions in WebAssembly. During the seminar week, the syntax extensions and operational semantics were worked out, with work begun on the reference implementation. During the seminar, Andrej Bauer pointed out that several prototype implementations that incorporate effect handlers exist, each with their own syntax and semantics. This makes it difficult to translate ideas across different research groups. Hence, Andrej proposed and initiated effects and handlers rosetta stone – a repository of examples demonstrating programming with effects and handlers in various programming languages. This repository is hosted on GitHub and has had several contributions during and after the seminar.

In conclusion, the seminar inspired discussions and brought to light the challenges in incorporating effect handlers in mainstream languages. During the previous seminar (16112), the discussions were centered around whether it was even possible to incorporate effect handlers into mainstream languages. During this seminar, the discussions were mainly on the ergonomics of effect handlers in mainstream languages. This is a testament to the success of the Dagstuhl Seminars in fostering cutting edge research.



Fig. 6.5
Aerial photography of Schloss Dagstuhl.

6.23 Towards Accountable Systems

Organizers: David Eyers, Christopher Millard, Margo Seltzer, and Jatinder Singh
Seminar No. 18181

Date: April 29–May 4, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.126

© Creative Commons BY 3.0 Unported license

© David Eyers, Christopher Millard, Margo Seltzer, and Jatinder Singh



Participants: Virgilio Almeida, Jean Bacon, Jennifer Cobbe, Jon Crowcroft, Lilian Edwards, David Eyers, Krishna P. Gummadi, Tristan Henderson, Martin Henze, Melanie Herschel, Heleen Louise Janssen, Joshua A. Kroll, Bertram Ludäscher, Derek McAuley, Christopher Millard, Ken Moody, Maximilian Ott, Frank Pallas, Thomas Pasquier, Silvia Puglisi, Margo Seltzer, Jatinder Singh, Barbara Staudt Lerner, Michael Veale, Ben Wagner, Michael Winikoff, Martina Zitterbart

Background and Motivation Technology is becoming increasingly pervasive, impacting all aspects of everyday life. Our use of apps and online services is tracked and extensively processed (data analytics), and the results are used for various purposes, predominately advertising. Monitoring and surveillance by sensors in smart cities creates vast amounts of data, much of which can be identifiably linked with people. Smart home, health and lifestyle monitoring, and other sensor technologies yield sensitive personal data; mobile phones reveal people’s positions, and their calls are tracked leading to data that can be used to determine social linkages and sometimes mental wellbeing. Such collection and analysis of personal data raises serious privacy concerns. A key aspiration is to provide end-users with a means to understand their digital footprints, and control the propagation, aggregation and retention of their data.

Concerns over data movement, location, processing and access have led to increasing regulation, both national and international. An example is the recently adopted EU General Data Protection Regulation (GDPR) that reinforces and expands individual rights, as well as restrictions and obligations regarding personal data. However, data moves easily beyond geographical boundaries, and use of cloud computing resources may mean that stored data may be replicated in multiple locations worldwide, with potential for conflicts between applicable laws and jurisdictions. Governments may demand access to data (whether stored locally or remotely) and this may result in complex legal disputes. Regulations, codes of conduct, and best practices can incentivise the use of particular technical mechanisms for data management. Examples include encryption and anonymisation, for example when using medical data for research. However, there are often misalignments between legal/regulatory aims and the capabilities of the technologies.

Key issues concern how to demonstrate compliance with regulations, such as those regarding how data is handled and

used, and, in cases of failure, how to hold the appropriate entities accountable. This is a particular challenge for wide-scale, federated, or cross-border systems. In large or complex systems, data may be handled by many different parties, falling under various management regimes and jurisdictions. Such concerns are not only horizontal (e.g., data being exchanged between parties, across geographic regions) but also vertical, where different levels of the services stack are managed by different parties (e.g., a company application running over a Heroku PaaS that runs over Amazon IaaS). Most end-users (people!) are oblivious to the potential complexity of such systems, let alone the complexity of the legal requirements that underpin such architectures. In general, the lack of transparency and uncertainty about the means for compliance with legal obligations, along with a lack of technical means for managing such concerns, may inhibit innovative technology development (a “chilling factor”), may escalate compliance costs, may trigger inappropriate policy responses, and may work to undermine public trust in technology.

These concerns will only grow in prominence, given the increasing deployment of sensors, generating ever-more data; actuators, giving systems physical effects; and the use of machine learning, facilitating automation. In response, this seminar brought together experts from the computer science and legal communities, spanning academia and industry, to explore issues of accountability as it relates to data and systems. The seminar aimed to: (i) raise awareness of and establish new research directions concerning issues of accountability as they relate to systems, given directions in systems technologies; (ii) explore developing legal and regulatory requirements; and (iii) investigate issues of user empowerment. A key goal was to increase awareness that law, regulation and requirements for data usage, management, security, confidentiality, quality and provenance should align with the technology, and *vice versa*: technologists should be legally-aware and lawyers should be technology-aware.

Seminar Structure Due to the diverse backgrounds of the participants, the first day was focused on introductions and ensuring that everyone had a common grounding in key topics. This included a series of guided discussion sessions: Lilian Edwards provided an introduction to legal and regulatory considerations, particularly the European Union General Data Protection Regulation (GDPR); Jon Crowcroft introduced emerging technical architectures such as edge computing; Bertram Ludäscher led a session exploring data provenance; and Ben Wagner introduced broader ethical and social concerns. A motivating case study was also presented highlighting how an apparently enthusiastic view of emerging Internet of Things technologies might obscure a plethora of questionable social and policy implications.

The structure of the week included multiple breakout sessions in which working groups examined particular topics (below) and reported back summaries of their discussions at plenary sessions. The working group sessions were interspersed with an interactive case study session, that focused on the technological compliance concerns of a hypothetical global hotel chain seeking to introduce a series of IoT and cloud technologies in the current regulatory environment, and a session in which participants were able to present their recent research, abstracts for (most of) which are included in this report.

Moving forward The topics explored by the working groups at the seminar spanned policy, legal and technical considerations. The topics were seeded by the organisers but were ultimately gathered from the participants through a preference allocation process. The chosen topics included:

- *Trust in systems.*
- *Who is, could or should be accountable in complex systems?*
- *Engineering accountable systems.*
- *Is there a place for data provenance in accountable systems?*
- *Anonymity, identity and accountability.*
- *Thinking beyond consent.*
- *Automating the exercising of rights for collective oversight.*

Each group was asked to produce an abstract summarising the key issues, challenges and ways forward from the discussion. These abstracts are included in this report, and indicate many potential opportunities for research.

Generally, it was felt that *this seminar represented only the start of this important discussion*. It is clear that there is a substantial and urgent need for closer interactions between the technical and legal domains, such that (i) the computer science communities better understand the legal requirements and constraints that impact the design, implementation and deployment of technology; and (ii) the legal communities gain more of a grounding in the nature, capabilities, and potential of the technology itself. It was also recognised that there is potential for better collaboration amongst different computer science communities; for example, to have greater interactions between those working in systems, provenance and machine learning.

In light of this, key to moving forward is to work to form collaborative research proposals, and to organise relevant meetings, in order to drive progress on the topics, challenges and research opportunities identified during this seminar. As issues of accountability increase in importance and urgency, it is vital that researchers across academia, industry and civil society work together to proactively confront these challenges.

6.24 Software Business, Platforms, and Ecosystems: Fundamentals of Software Production Research

Organizers: Pekka Abrahamsson, Jan Bosch, Sjaak Brinkkemper, and Alexander Mädche
Seminar No. 18182

Date: April 29–May 2, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.4.164

© Creative Commons BY 3.0 Unported license

© Slinger Jansen and Karl Werder



Participants: Pekka Abrahamsson, Rahul C. Basole, Jan Bosch, Sjaak Brinkkemper, Christoph Bussler, Michael Cusumano, Jens Förderer, Samuel A. Fricker, Paul Grünbacher, Robert Heinrich, Armin Heinzl, Mika Helenius, Georg Herzwurm, Helena Holmström Olsson, Sami Hyrynsalmi, Slinger Jansen, Zhi Jin, Hans-Bernd Kittlaus, Thomas Kude, Alexander Mädche, Tiziana Margaria, Efi Papatheocharous, Balasubramaniam Ramesh, Guenther Ruhe, Kari Smolander, Diomidis Spinellis, Pasi Tyrväinen, Xiaofeng Wang, Karl Werder, Krzysztof Wnuk

Software producing organizations (SPOs) face challenges every day. Whether they are open source consortia or commercial software product companies, they all face the challenges of changing demands, rapidly evolving technology, and a dynamic ecosystem in which their products and services need to operate. SPOs need to rethink their operating models and benefit from current and future trends. E.g. agile software development and DevOps allow them to respond swiftly to changes in their environment, embracing uncertainty. Particularly in conjunction with machine learning and artificial intelligence, SPOs can generate strategic competitive advantages. Particularly companies with a long history in a given domain, such as SAP and Volkswagen, seem to be too comfortable with their status quo. Meanwhile, smaller companies drive innovation on many fronts. Examples are Provenance that benefits from blockchain technology to revolutionize trust in goods, or Tesla and Local Motors that push autonomous cars into consumer markets.

The challenge to make these organizations successful is multi-disciplinary. First, there exist technology challenges, such as eliciting and prioritizing requirements, dealing with platforms and technology standards, and operating in complex technology landscapes that constrain and enable their technology. Secondly, there exist adoption challenges: organizations need to find ways to convince their target users to adopt their technologies and to coordinate evolving technologies to provide the most valuable end-user experience. Thirdly, there exist business model challenges, where these organizations must find ways to maximize profit from their innovations and technologies. Because of the pervasiveness of software, the challenges are observed everywhere in the economy, whether it is logistics, online marketing, or e-health. Furthermore, they are applicable to organizations in every stage of development, whether it is a software startup or a software giant that has influenced the market consistently for decades.

Hence, this Dagstuhl Seminar invited thought leaders from

academia and industry to share their knowledge and experiences. Participants were asked to share a short position statement of max 300 words and participate in the development of a groundbreaking research agenda. These efforts aimed to increase visibility and impact of software production research and to set a course for the next decades. In addition, the seminar helped bringing together scholars and industry practitioners from different communities, such as product management, technology management, information systems, software engineering, and human-computer interaction in order to sharpen and define the joint community of Software-intensive Business (see Section 4.3.1 of the full report).

A central outcome of the seminar was the agreement to use the term Software-intensive Business in order to describe the joint community with members of great diversity. Furthermore, the seminar focused on

- defining core concepts and identifying a roadmap
- Software-intensive Business and technology artifacts
- research needs in continuous experimentation & innovation
- lifecycle and research of software ecosystems
- research data for Software-intensive Businesses

As a major result from the seminar, the following achievements have been identified:

1. research a clear agenda for the field of Software-intensive Business research
2. carving out trends and research challenges in further depth
3. forming groups for continuous collaborations on different elements of the research agenda
4. organize bi-weekly meetings on-line for community building and research sharing.



Fig. 6.6

“A rare opportunity to admire Don Knuth’s books together at the @dagstuhl magnificent Computer Science library.”

Twitter post by 18182 Dagstuhl Seminar participant Diomidis Spinellis.

<https://twitter.com/CoolSWEng/status/991606321827938305>. Photo courtesy of Diomidis Spinellis.

6.25 Secure Compilation

Organizers: Amal Ahmed, Deepak Garg, Catalin Hritcu, and Frank Piessens

Seminar No. 18201

Date: May 13–18, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.5.1

© Creative Commons BY 3.0 Unported license

© Amal Ahmed, Deepak Garg, Catalin Hritcu, and Frank Piessens



Participants: Amal Ahmed, Gilles Barthe, Nick Benton, Frédéric Besson, Pramod Bhatotia, Lars Birkedal, Roberto Blanco, William Bowman, Stefan Brunthaler, David Chisnall, John T. Criswell, Dominique Devriese, Derek Dreyer, Akram El-Korashy, Cédric Fournet, Deepak Garg, Chris Hawblitzel, Catalin Hritcu, Chung-Kil Hur, Limin Jia, Gabriele Keller, Vincent Laporte, Xavier Leroy, Toby Murray, Magnus Myreen, Santosh Nagarakatte, Kedar Namjoshi, David A. Naumann, Max S. New, Scott Owens, Zoe Paraskevopoulou, Marco Patrignani, Daniel Patterson, Frank Piessens, Tamara Rezk, Christine Rizkallah, Peter Sewell, Deian Stefan, Andrew Tolmach, Stelios Tsampas, Neline van Ginkel, Stephanie Weirich, Steve Zdancewic

Today’s computer systems are distressingly insecure. The semantics of mainstream low-level languages like C and C++ is inherently insecure, and even for safer languages, establishing security with respect to a high-level semantics does not prevent devastating low-level attacks. In particular, all the abstraction and security guarantees of the source language are currently lost when interacting with lower-level code, for instance when using low-level libraries. For a concrete example, all modern languages provide a notion of structured control flow and an invoked procedure is expected to return to the right place. However, today’s compilation chains (compilers, linkers, loaders, runtime systems, hardware) cannot efficiently enforce this abstraction: linked low-level code can call and return to arbitrary instructions or smash the stack, blatantly violating the high-level abstraction.

Secure compilation is an emerging field that puts together advances in security, programming languages, compilers, verification, systems, and hardware architectures in order to devise secure compiler chains that eliminate many of today’s low-level vulnerabilities. Secure compilation aims to protect high-level language abstractions in compiled code, even against low-level attacks, and to allow sound reasoning about security in the source language. The emerging secure compilation community aims to achieve this by:

1. **Identifying and formalizing secure compilation criteria and attacker models.** What are the properties we want secure compilers to have, and under what attacker models? Should a secure compilation chain preserve observational equivalence of programs? Should it preserve some class of security properties of the source programs? Should it guarantee invariants on the run-time state of the compiled program (like for instance well-formedness of the call-stack)? And what are realistic attacker models? Can attackers only interact with compiled programs by providing input and reading output? Or can they link arbitrary low-level code to the program?

Well-studied notions like fully abstract compilation provide partial answers: a fully abstract compiler chain preserves observational equivalence under an attacker model where attackers are target-level contexts. Even where this is the desired end-to-end security goal, it can still be too hard to enforce, for instance in cases where target level contexts can measure time.

2. **Efficient enforcement mechanisms.** The main reason today’s compiler chains are not secure is that enforcing abstractions in low-level compiled code can be very inefficient. In order to overcome this problem, the secure compilation community is investigating various efficient security enforcement mechanisms: from the use of static checking of low-level code to rule out linking with ill-behaved contexts, to software rewriting (e.g., software fault isolation), dynamic monitoring, and randomization. One key enabler is that hardware support for security is steadily increasing.
3. **Developing effective formal verification techniques.** Secure compilation properties like full abstraction are generally much harder to prove than compiler correctness. Intuitively, in order to show full abstraction one has to be able to back-translate any low-level context attacking the compiled code to an equivalent high-level context that can attack the original source code. This back-translation is, however, nontrivial, and while several proof techniques have been proposed (e.g., based on logical relations, bisimulations, game semantics, multi-language semantics, embedded interpreters, etc.), scaling these techniques to realistic secure compilers is a challenging research problem. This challenge becomes even more pronounced if one expects a strong level of assurance, as provided by formal verification using a proof assistant.

The Secure Compilation Dagstuhl Seminar 18201 attracted a large number of excellent researchers with diverse backgrounds.

The 45 participants represented the programming languages, formal verification, security, and systems communities, which led to many interesting points of view and enriching discussions. Some of these discussions were ignited by the “guided discussions” on the 3 aspects above and by the 35 talks contributed by the participants. The contributed talks spanned a very large number of topics: investigating various secure compilation criteria and attacker models, building prototype secure compilation chains, proposing different enforcement techniques, studying the relation to verified compilation and compositional compiler correctness, specifying and restricting undefined behavior, protecting against side-channels, studying intermediate representations, performing translation validation, securing multi-language interoperability, controlling information-flow, compartmentalizing software, enforcing memory safety, compiling constant-time cryptography, securing compiler optimizations, designing more secure (domain-specific) languages, enforcing security policies, formally specifying the semantics of realistic languages and ISAs, compartmentalization, capability machines, tagged architectures, integrating with existing compilation chains like LLVM, making exploits more difficult by diversification, multi-language interoperability, etc. Talks were interspersed with lively discussions, since by default each speaker could only use half of the time for presenting and had to use the other half for answering questions and engaging with the audience.

Given the high interest spurred by this first edition and the positive feedback received afterwards, we believe that this Dagstuhl Seminar should be repeated in the future. Particular aspects that could still be improved in future editions is focusing more on secure compilation and spurring more participation from the practical security and systems communities.

6.26 Inter-Vehicular Communication Towards Cooperative Driving

Organizers: Onur Altintas, Suman Banerjee, Falko Dressler, and Geert Heijenk

Seminar No. 18202

Date: May 13–16, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.5.31

© Creative Commons BY 3.0 Unported license

© Falko Dressler, Onur Altintas, Suman Banerjee, Geert Heijenk, and Katrin Sjöberg



Participants: Onur Altintas, Ali Balador, Aruna Balasubramanian, Suman Banerjee, Claudia Campolo, Wai Chen, Sinem Coleri Ergen, Falko Dressler, Eylem Ekici, Sonia Heemstra de Groot, Thorsten Hehn, Geert Heijenk, Albert Held, Frank Kargl, Florian Klingler, Renato Lo Cigno, Jörg Ott, Elmar Schoch, Michele Segata, Christoph Sommer, Jonathan Sprinkle, Thomas Strang, Erik Ström, Isabel Wagner, Lars Wischhof, Lars Wolf, Andrea Zanella, Hongwei Zhang

Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of communication protocols to support safety applications, intelligent navigation, multi-player gaming and others. Very large projects have been initiated to validate the theoretic work in field tests and protocols are being standardized. With the increasing interest from industry, security and privacy have also become crucial aspects in the stage of protocol design in order to support a smooth and carefully planned roll-out. We are now entering an era that might change the game in road traffic management. This is supported by the U.S. federal government announcement in December 2016 that National Highway Traffic Safety Administration (NHTSA) plans to make V2V devices in new vehicles mandatory. This coincides with the final standardization of higher layer networking protocols in Europe by the ETSI.

The vehicular networking research also complements the ongoing activities towards automated driving. Very successful activities started with the Google and lead to first projects on the road such as the Singapore driverless taxi service or the platooning experiments in Scandinavia and now Germany.

The management and control of network connections among vehicles and between vehicles and an existing network infrastructure is currently one of the most challenging research fields in the networking domain. Using the terms Vehicular Ad-hoc Networks (VANETs), Inter-Vehicle Communication (IVC), Car-2-X (C2X), or Vehicle-2-X (V2X), many applications – as interesting as challenging – have been envisioned and (at least) partially realized. In this context, a very active research field has developed. There is a long list of desirable applications that can be grouped into four IVC categories:

1. eSafety applications that try to make driving safer, e.g. road hazard warning;

2. traffic efficiency applications aiming at more efficient and thus greener traffic, e.g., detection of traffic jams;
3. manufacturer oriented applications, e.g., automatic software updates; and
4. comfort applications, e.g. automatic map updates.

In 2010, a first Dagstuhl Seminar (10402) was organized on the topic of inter-vehicular communication. The motivation was to bring together experts in this field to investigate the state of the art and to highlight where sufficient solutions already existed. The main outcome of this very inspiring seminar was that there are indeed areas within this research where scientific findings are being consolidated and adopted by industry. This was the consensus of quite intriguing discussions among participants from both industry and academia. Yet, even more aspects have been identified where substantial research is still needed. These challenges have been summarized in the following IEEE Communications Magazine article [1].

A follow-up seminar (13392) was organized in 2013. The goal was to again bring together leading researchers both from academia and industry to discuss if and where the previously identified challenges have been adequately addressed, and to highlight where sufficient solutions exist today, where better alternatives need to be found, and also to give directions where to look for such alternatives. Furthermore, it was the goal of this workshop to go one step beyond and identify where IVC can contribute to the basic foundations of computer science or where previously unconsidered foundations can contribute to IVC. It turned out that quite a number of research questions were still open or insufficiently addressed. This particularly included scalability and real-time capabilities. These challenges have been summarized in the following IEEE Communications Magazine article [2].

We now shifted the focus of this seminar from basic net-

working principles to networked control applications. We were particularly interested in the first two IVC categories that are thought to yield substantial benefits for the emerging “cooperative automated driving” domain. It is of utmost importance to bring together expertise from classical computer science (computer networking, simulation and modeling, operating system design), from electrical engineering (digital signal processing, communication networks), as well as from automated driving (mechanical engineering, image processing, control theory). Building upon the great success of the first two seminars, with this follow-up seminar, we aimed to again bring together experts from all these fields from both academia and industry.

The seminar focused intensively on discussions in several working groups. To kick-off these discussions, we invited two keynote talks “Cooperative Driving A Control of a Networking Problem?” by Renato Lo Cigno and “Cooperative driving – maneuvers, perception, and IVC” by Lars Wolf. These keynotes were complemented by four additional talks:

Human-in-the-Loop: Towards Deeply Integrated Hybridized Systems (Falko Dressler), Machine Learning for Cooperative Driving (Geert Heijenk), Measuring Privacy in Vehicular Networks (Isabel Wagner), and Predictable V2X Networking for Application-Networking Co-Design (Hongwei Zhang). We finally organized the following working groups on some of the most challenging issues related to inter-vehicular communication and cooperative driving:

- Ultra-Reliable Low-Latency and Heterogeneous V2X Networking,
- Human-in-the-Loop,
- Safety-critical Vehicular Network Applications,
- Security and Privacy,
- Network and Cloud based Control, and
- Sensing and Data Management.

For most of these working groups, we provide in-depth feedback from the experts in this report.

■ References

- 1 Falko Dressler, Hannes Hartenstein, Onur Altintas, and Ozan K. Tonguz. Inter-Vehicle Communication – Quo Vadis. *IEEE Communications Magazine*, 52(6):170–177, June 2014.
- 2 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz, and Lars Wischhof. Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar. *IEEE Communications Magazine*, 49(5):158–164, May 2011.

6.27 Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance

Organizers: Javier Esparza, Pierre Fraigniaud, Anca Muscholl, and Sergio Rajsbaum
Seminar No. 18211

Date: May 21–25, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.5.60

© Creative Commons BY 3.0 Unported license

© Anca Muscholl, Javier Esparza, Pierre Fraigniaud, and Sergio Rajsbaum



Participants: Paul C. Attie, Hagit Attiya, A. R. Balasubramanian, Michael Blondin, Benedikt Bollig, Borzoo Bonakdarpour, Ahmed Bouajjani, Dan Brownstein, Keren Censor-Hillel, Aiswarya Cyriac, Giorgio Delzanno, Cezara Dragoi, Jo Ebergen, Yuval Emek, Constantin Enea, Javier Esparza, Bernd Finkbeiner, Marie Fortin, Pierre Fraigniaud, Matthias Függer, Paul Gastin, Swen Jacobs, Igor Konnov, Sandeep Kulkarni, Marijana Lazic, Jérémy Ledent, Martin Leucker, Stephan Merz, Roland Meyer, Yoram Moses, Anca Muscholl, Rotem Oshman, Mor Perry, Sergio Rajsbaum, David A. Rosenblueth, Arnaud Sangnier, Ana Sokolova, Paola Spoletini, Corentin Travers, Igor Walukiewicz, Philipp Woelfel

The original motivation of this workshop has to do with the evolution of research in Computer Science. The first ACM conference on Principles of Distributed Computing (PODC) was held in 1982. The proceedings of its first editions included papers on distributed algorithms³⁶, formal methods for distributed systems³⁷, or a combination of the two. However, in 1990 the area of formal methods for distributed computing branched out, and started its own conference, the International Conference on Concurrency Theory (CONCUR), now in its 27th edition. PODC and CONCUR have become the premier conferences in their respective fields, and, after over 20 years of almost independent evolution, feel the need to close a gap that slows down progress, limits the applicability of the results, and causes repetitions and inconsistencies.

Our seminar aimed at achieving synergy by bringing together the two research areas, both with deep understanding of distributed computation, but different perspectives. We had two longer tutorials, one about concurrent data structures by Ph. Woelfel and one about verification of concurrent programs by A. Bouajjani. In addition, we had several survey talks, on correctness in concurrent programming (H. Attiya), distributed runtime verification (B. Bonakdarpour), distributed property testing (K. Censor-Hillel), distributed synthesis (B. Finkbeiner), and parametrized verification (I. Konnov).

The scientific programme was quite dense, given that we had only 4 days and almost all participants proposed to give a talk. Exchanges were very lively, and the discussion that we had with all participants showed that this kind of workshop is a great opportunity to compare our approaches and find new research directions, inspired by the perspectives of the other community. We warmly thank Marie Fortin for the editorial work on this report

and the Dagstuhl staff for the excellent conditions provided for our seminar.

³⁶ Algorithms designed to run on computer hardware constructed from interconnected processors.

³⁷ Mathematically based techniques for the specification, development and verification of software and hardware systems.



Fig. 6.7
Aerial photography of Schloss Dagstuhl.

6.28 On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces

Organizers: Kasper Hornbaek, David Kirsh, Joseph A. Paradiso, and Jürgen Steimle
Seminar No. 18212

Date: May 21–24, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.5.80

© Creative Commons BY 3.0 Unported license

© Kasper Hornbaek, David Kirsh, Joseph A. Paradiso, and Jürgen Steimle



Participants: Eduard Arzt, Nadia Bianchi-Berthouze, Liwei Chan, Clément Duhart, Michael Haller, Nur Hamdan, Kayla J. Heffernan, Christian Holz, Kasper Hornbaek, David Kirsh, Antonio Krüger, Pedro Lopes, Paul Lukowicz, Karon MacLean, Aditya Shekhar Nittala, Joseph A. Paradiso, Eric Paulos, Anastasia Pavlidou, Henning Pohl, Ivan Poupyrev, Joan Sol Roo, Chris Schmandt, Albrecht Schmidt, Jürgen Steimle, Paul Strohmeier, Katia Vega

■ Motivation

For the past 40 years, input to computers has been given with mouse and keyboard. Over the last decade, multi-touch has become popular for small devices (e.g., phones and tablets) as well as for large displays (e.g., interactive tabletops and wall-sized screens). All these forms of input require the user to hold or touch a device. Conversely, output has happened on large screens external to the body (e.g., a desktop) or small ones on the body (e.g., smartwatches). The field of human-computer interaction (HCI) has worked to understand these user interfaces (UIs) and how people use them, in addition to establishing principles of design and models of performance to help design them so that they are useful and usable.

Recently, however, HCI researchers have been interested in allowing new forms of on-body technologies. One vision is to integrate technology with the body so as to use and supplement its capabilities. In particular, researchers have focused on sensing users' movement and gestures, aiming to allow users to interact using their body rather than by using a device. Early work included Bolt's put-that-there system developed in the late 1970s, and recent advances in computer vision have allowed the tracking of users' hands, arms, and bodies, leading to a flurry of motion-based gaming controls and inventive, body-based games. The number and variety of research prototypes of non-device UIs have also exploded over the past few years, showing how movements in front of a large display can control navigation, how users can gesture in mid-air, how scratching or poking the skin of one's forearm can be a means of input, and how electric muscle stimulation can be used to move users' limbs as output. Further, HCI researchers have been exploring the theoretical opportunities in using the body for interaction, describing principles for whole-body interaction, embodied interaction, and body-centric interaction, as well as highlighting some of the philosophical and psychological challenges associated with using the body as

an interface. First promising applications are being investigated or have been demonstrated in mobile computing, healthcare, or sports. A new UI paradigm seems to be emerging.

The main objective of the seminar was to explore on-body interaction through two research areas: embodied cognition and sensor/actuator engineering. The former has driven a lot of thinking and models around on-body technologies and the potential of body-based interaction. The latter has been behind many of the sensors and actuators that have enabled prototypes to be built and to demonstrate the potential of on-body technology. We did this bringing together a group of researchers from embodied cognition (including psychology, robotics, human-computer interaction, art/design, and sociology) as well as sensor/actuator engineering (including computer science, materials science, electrical engineering). Second, we had this diverse group of researchers outline a research agenda for on-body technologies, in part using a bottom-up process at the seminar, in part using structured answers to questions in advance of the seminar.

■ Topics

In line with the objectives above, the seminar focused on three areas of investigation:

- **Embodied Cognition:** Embodied cognition is a term covering research in linguistics, robotics, artificial intelligence, philosophy, and psychology (e.g., Anderson 2003, Wilson 2002). The core idea in embodied cognition is that our bodies shape thinking broadly understood (including reasoning, memory, and emotion). In contrast to most psychological foundations of HCI, embodied cognition argues that one cannot study the human as a system comprising input (senses), processing (thinking), and output (motor activity), because sensor-motor activity affects thinking fundamentally and, conversely but less radically, because our body reflects more about our

thinking than is commonly expected. Thus, bodies and thinking are intertwined, as reflected in embodied cognition book titles like “How the Body Shapes the Way We Think” [2] and “How the Body Shapes the Mind” [1]. Embodied cognition has become a prominent candidate for outlining what we can and cannot do in on-body interaction.

- **Sensor/Actuator Engineering:** The engineering of technologies that transform the human body into an interface is a very active research area. A widely used approach uses techniques from visual computing for capturing body gestures and touch input on the body using RGB or depth cameras, while projecting visual output with a body-worn projector. Other approaches build on the transdermal propagation of ultrasound or electromagnetic waves to identify the location of touch contact on human skin. EMG can be used to capture human muscle movement, while Electrical Muscle Stimulation can generate muscle output. Radar is another technology that has been successfully demonstrated very recently for capturing gestural input. A further recent strand in research uses slim skin electronics for sensing and output on the body. These technologies are opening up new avenues for human-computer interaction, by contributing body-based sensing input and output modalities with an increasing resolution and more body compatible form factors.
- **New On-Body Technologies:** This area concerns how we can combine embodied cognition and sensor/actuator engineering to design on-body technologies. The design of on-body technologies was a key discussion topic, in particular, how to drive the technical development from work on embodied cognition and the body, how to evaluate on-body technology, and how to take the peculiarities and possibilities of the body into consideration. The application areas of on-body technologies were another consideration.

■ Activities

The first day of the seminar was reserved for presentations, to establish common ground for discussions. All participants introduced themselves, their background, and their vision in short position talks.

Four long talks reviewed the state-of-the-art and presented recent work in key areas. In his talk “Embodied Cognition: What does having a body gives us?”, David Kirsh emphasized on four topics: Effectivity, Enactive perception, Interactive Cognition, and Experience. They all explore what having a body gives us that goes beyond just having a sensor in space. Katia Vega’s talk, entitled “Beauty Technologies”, focused on the possibilities to embed technology on and inside the skin. Nadia Bianchi-Berthouze gave a talk entitled “The Affective Body in Interaction”, discussing the high-level principles of affective computing and creating body-affective-aware-computing technology, which involves sensing the affect and emotion of the users and using them for interaction. In his talk “Cosmetic Computing: Actions and Urgencies towards an Inclusive, Equitable Landscape of On-Body Technologies”, Eric Paulos urged the need for transdisciplinary and interdisciplinary approaches and proposed a framing around “Cosmetic Computing”.

The evening featured a demo session. An impressive total number of 8 interactive demos and exhibits were demonstrated in

the historical ambiance of the Rokoko-style music hall. Those demos comprised, amongst other, e-textiles, interactive tattoos and make-up, new bio-inspired materials and tactile actuation technologies.



Fig. 6.8

Demo session featuring latest body-based technologies, held in the historical Music Hall of Dagstuhl castle.

The second day consisted of work in **breakout groups**. First, groups identified **challenges** for future work in the field of on-body interaction, grouped into four main areas: Integration of the body and the device; Cognition and Affect; Interaction; and Applications. Next, the participants worked together to identify **positive visions** of a future with body-based interfaces. Promising aspects that were identified include sensory augmentation of human body for graceful ageing, personalized medication and the idea of legal/democratic framework for controlling wearable technology.. To identify potential risks associated with body-based technologies and interaction, the group also developed **negative visions**. Key problems and risks that were identified include a loss of physical embodiment and substantial security risks of our bodies (and potentially even emotions) being externally controlled.

In an session, entitled academic speed-dating, we randomly paired two participants with each other. Their goal was to develop within 7 minutes an idea and a title for a paper they would write together. The format turned out to be very well-received and to stimulate research ideas at unforeseen intersections between the participants’ interest and expertise.

■ Conclusion

The seminar set out to bring together diverse researchers to discuss the overlap between embodied cognition and sensor/actuator engineering. The group managed to cover advances in on-body sensors and actuator, some of the cognitive consequences of on-body technologies, and open issues in applications of on-body technologies. Further, a range of open questions and exciting research questions were discussed, which will likely foster future collaboration and serve as a generator of future research on on-body technologies.

■ References

- 1 Shaun Gallagher. *How the body shapes the mind*. Clarendon Press, 2006.
- 2 Rolf Pfeifer and Josh Bongard. *How the body shapes the way we think: a new view of intelligence*. MIT press, 2006.

6.29 The Constraint Satisfaction Problem: Complexity and Approximability

Organizers: Martin Grohe, Venkatesan Guruswami, and Stanislav Živný
Seminar No. 18231

Date: June 3–8, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.6.1

© Creative Commons BY 3.0 Unported license

© Martin Grohe, Venkatesan Guruswami, Dániel Marx, and Stanislav Živný



Participants: Isolde Adler, Per Austrin, Libor Barto, Christoph Berkholz, Manuel Bodirsky, Joshua Brakensiek, Andrei A. Bulatov, Clement Carbonnel, Hubie Chen, Victor Dalmau, Peter Fulla, Leslie Ann Goldberg, Alexander Golovnev, Martin Grohe, Venkatesan Guruswami, Johan Hastad, Mark R. Jerrum, Peter Jonsson, Alexandr Kazda, Eun Jung Kim, Vladimir Kolmogorov, Michael Kompatscher, Marcin Kozik, Andrei Krokhn, Victor Lagerqvist, Euiwoong Lee, Yury Makarychev, Barnaby Martin, Dana Moshkovitz, Antoine Mottet, Ryan O'Donnell, Miroslav Olsak, Jakub Opršal, Marcin Pilipczuk, Michael Pinsker, Akbar Rafiey, Miguel Romero, Biman Roy, Tselil Schramm, Stefan Szeider, Johan Thapper, Madhur Tulsiani, Caterina Viola, Ross Willard, Yuichi Yoshida, Dmitry Zhuk, Stanislav Živný

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 colourability, and systems of equations. The CSP framework originated 30–35 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 the 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 counting assignments or involve 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, non-trivially approximable, fixed-parameter tractable, or definable in a weak logic). One of the most striking features of this research direction is the variety of different branches of mathematics (including universal algebra and logic, combinatorics and graph theory, probability theory and mathematical programming) that are used to achieve deep insights in the study of the CSP. In the last decade, research activity in this area has significantly intensified and hugely impressive progress was made.

The recent flurry of activity on the topic of the seminar is witnessed by four previous Dagstuhl seminars, titled “Complexity of constraints” (06401) and “The CSP: complexity and approximability” (09441, 12541, 15301), that were held in 2006, 2009, 2012, and 2015 respectively. This seminar was a follow-up to the 2009, 2012, and 2015 seminars. Indeed, the exchange of ideas at the 2009, 2012, and 2015 seminars has led to ambitious new research projects and to establishing regular communication channels. There is clearly the potential for further systematic interaction that will keep on cross-fertilising the areas and opening new research directions. The 2018 seminar brought together 47 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 participants presented, in 24 talks, their recent results on a number of important questions concerning the topic of the seminar. One particular feature of this seminar is a significant increase in the number of talks involving multiple subareas and approaches within its research direction – a definite sign of the growing synergy, which is one of the main goals of this series of seminars.

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 a multitude of impressive results reported during

the seminar and 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 another seminar focusing 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 problems (say, between P and NP-complete, assuming $P \neq NP$), 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 the 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 was 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 the 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. The above conjectures have algebraic versions which also predict in algebraic terms where the boundary between harder problems and easier problems lies. 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), culminating in two independent proofs of the general conjecture announced in 2017 by Bulatov and Zhuk.

- Bulatov and Zhuk gave detailed talks on the main insights into their proofs.
- Kolmogorov described an algorithm for Boolean CSPs under the restriction that every variable appears in exactly two constraints and all constraints are even Δ -matroids.

The valued CSP (VCSP) is a significant generalisation of the CSP that involves both feasibility and optimisation aspects. While the computational complexity of finite-domain VCSPs is by now well understood, the infinite-domain VCSPs are fairly unexplored.

- Viola gave a talk on submodular VCSPs on infinite domains.
- Kazda presented his results on the structure of weighted clones, which are intimately related to the computational complexity of VCSPs.

Approximability of CSPs. The use of approximation algorithms is one of the most fruitful approaches to coping with NP-hardness. Hard optimisation problems, however, exhibit diverse behavior with respect to approximability, making it an exciting research area that is by now well-developed but far from fully understood.

An emerging topic bridging the complexity of the CSP with approximation aspects is *promise constraint satisfaction* (PCSP). The PCSP is a generalization of the CSP in which the

constraints come in pairs of “stricter” and “weaker” versions. In a PCSP instance, the task is to find an assignment satisfying the weaker constraints under the promise that there is an assignment satisfying the strict constraints.

- Brakensiek gave an introductory talk to this exciting research direction and also presented a dichotomy classification for symmetric Boolean PCSPs.
- Opršal explained the very recently introduced algebraic approach to the computational complexity of PCSPs.
- Barto presented his results on PCSPs and cyclic operations.

Many approximation algorithms for CSPs are based on convex relaxations.

- Berkholz gave an overview on relaxations for Boolean CSPs based on algebraic methods.
- Schramm explained the power of semidefinite programming relaxations for random CSPs.
- Tulsiani presented results on the limits of linear programming relaxations for CSPs.
- Makarychev showed how to obtain an integrality gap for the Călinescu-Karloff-Rabani linear programming relaxation of the Multiway-Cut problem.
- Austrin established the currently best known inapproximability result for Min UnCut, which is a special Boolean CSP.

Some of the most exciting developments in approximability in the last decade revolve around the *unique games conjecture*, or UGC, of Khot (2002). This bold conjecture asserts that, for CSPs with a certain constraint language over a large enough domain, it is NP-hard to distinguish almost satisfiable instances from those where only a small fraction of constraints can be satisfied. This conjecture is known to imply tight inapproximability results for many classical optimisation 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).

- Moshkovitz presented recent developments on the so-called 2-to-2 PCP theorem, which covers important special cases of the UGC.

Logic and the complexity of CSPs. Logic has been used in two distinct ways in the study of the CSP. One of them, starting from earlier work of Kolaitis and Vardi, is *descriptive complexity*, where one tries to classify CSPs as classes of instances with respect to definability in a given logic. The other way is to use logic to specify CSP instances, which can be done very naturally. The latter direction leads to generalisations such as the quantified CSP (QCSP), as well as to the study of CSPs over infinite domains, where important links with the algebraic approach were found.

- Roy presented a dichotomy theorem for the inverse satisfiability problem.
- Bodirsky gave a talk on two methods of reducing infinite-domain CSPs to finite-domain CSPs.
- Pinsker explained recent results on the algebraic approach to infinite-domain CSPs. These results are related to the so-called loop conditions, which were in more detail discussed by Kozik.
- Kompatscher presented a proof of the equivalence of two dichotomy conjectures for infinite-domain CSPs.
- Mottet gave a new proof of the dichotomy for MMSNP and discussed consequences for infinite-domain CSPs.

- Martin described recent results for temporal and spatial problems, which are special cases of infinite-domain CSPs.

Exact exponential complexity of CSPs. The area of parameterised complexity is closely related to the area of exact exponential complexity, in which the goal is to design the most efficient exponential-time algorithms. There has been significant progress on the exact exponential complexity of CSPs.

- Golovnev presented results that give optimal lower bounds on the running time of algorithms for deciding if there is a homomorphism from one graph to another.
- The complexity of counting solutions for CSPs and related problems from statistical physics were presented by Goldberg and Jerrum.



Fig. 6.9

Impression from the visit of the Saarland Minister President Tobias Hans during his summer tour 2018.

Photo courtesy of Saarland/mn.

6.30 High-Performance Graph Algorithms

Organizers: Henning Meyerhenke, Richard Peng, and Ilya Safro
Seminar No. 18241

Date: June 10–15, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.6.19

© Creative Commons BY 3.0 Unported license

© Henning Meyerhenke, Richard Peng, and Ilya Safro



Participants: Nesreen K. Ahmed, Eugenio Angriman, David A. Bader, Maciej Besta, Rob Bisseling, Timothy Chu, Pierluigi Crescenzi, Timothy Alden Davis, Irene Finocchi, John Gilbert, David F. Gleich, Riko Jacob, George Karypis, Michel A. Kinsy, Marsha Kleinbauer, Christine Klymko, Yiannis Koutis, Danai Koutra, Rasmus Kyng, Nelly Litvak, Fredrik Manne, Henning Meyerhenke, Marco Minutoli, Lalla Mouatadid, Danupon Nanongkai, Lorenzo Orecchia, Richard Peng, Manuel Penschuck, Cynthia A. Phillips, Alex Pothen, Maria Predari, Vijaya Ramachandran, Sushant Sachdeva, Ilya Safro, Peter Sanders, Christian Schulz, Julian Shun, Blair D. Sullivan, Charalampos E. Tsourakakis

Many presentations in this Dagstuhl seminar emphasized recent trends regarding typical inputs and their effect on graph algorithm development. From a high-level perspective, one can divide the presentations into two categories: either more focused on algorithm theory or more focused on practical algorithmic results. Many talks considered both theoretical and practical aspects. Furthermore, attention was given to intermix talks with theoretical and practically motivated starting points in order to encourage discussions among attendees. We were happy to see such discussions, as well as synergy of both aspects, carrying over to working groups on open problems.

Theory-focused talks were given by Sachdeva, Nanongkai, Jacob, Mouatadid, Kyng, Tsourakakis, and Litvak. They considered numerous topics such as Laplacian solvers and related optimization techniques, dynamic graph algorithms, external-memory graph algorithms, graph decompositions, and generative models.

The talks with emphasis on practical performance can be further subdivided into three subclasses: (i) graph mining, network analysis and optimization, (ii) parallel, distributed and streaming graph algorithms and (iii) graph generation. The talks given by Koutra, Ahmed, Klymko, Angriman, Gleich and Schulz fall into the first subclass, with a wide variation of algorithmic problems under consideration. Likewise, it was interesting to see the variety in computing platforms and tools (for example shared memory, message passing, distributed systems, streaming from databases, GraphBLAS) used in the eight talks of subclass two, presented by Besta, Shun, Predari, Ramachandran, Pothen, Bader, Finocchi and Davis. Finally, the talks by Phillips, Sanders and Penschuck as well as Crescenzi dealt with generating very large graphs with properties also found in real-world graphs – which is important, among others, for convincing scaling studies in algorithm engineering.

6.31 Secure Routing for the Internet

Organizers: Phillipa Gill, Adrian Perrig, and Matthias Wählisch
Seminar No. 18242

Date: June 10–13, 2018 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.8.6.40
 © Creative Commons BY 3.0 Unported license
 © Vasileios Kotronis

Participants: Mai Ben-Adar Bessos, Nikita Borisov, Georg Carle, Shinyoung Cho, Ítalo Cunha, Marc C. Dacier, Phillipa Gill, Joel M. Halpern, Raphael Hiesgen, Carlee Joe-Wong, Mattijs Jonker, Vasileios Kotronis, Taeho Lee, Hemi Leibowitz, Victoria Manfredi, Marcin Nawrocki, Christos Pappas, Adrian Perrig, Alvaro Retana, Andreas Reuter, Thomas C. Schmidt, Laurent Vanbever, Pierre-Antoine Vervier, Stefano Vissicchio, Rüdiger Volk, Matthias Wählisch, Bing Wang



The seminar was focused on the following aspects of routing security, mostly in the context of traditional inter-domain routing security: (i) Protocol design vs tooling, (ii) sources of relevant routing data and their accuracy/collection challenges, including policy databases, (iii) the need for metadata and dataset “labelling”, (iv) monitoring and detection of routing attacks and anomalous incidents, such as BGP hijacks and route leaks, incentives for network operators to adopt routing security protocols, (v) testbeds for routing experiments, (vi) hijacks as enabling attacks against ToR and Bitcoin, on the application level, (vii) prevention of routing attacks, (viii) anonymity, privacy and (anti-)censorship. Moreover, we discussed in depth about (ix) PKI and cryptographic verification and protection mechanisms, and their use in securing routing infrastructures, such as the RPKI and BGPsec protocols. Finally, we (x) approached BGP flowspecs, DDoS attacks and QoS in the Internet as separate topics of interest in the field. Another goal of the seminar was to touch upon (xi) future network routing architectures which offer routing security “by design”, especially in light of demanding upcoming applications such as IoT, car-to-car communications, sensor swarms, and wireless routing at scale, and identify related security and privacy concerns and objectives.

Besides the specific goals of the seminar, it is also worth noting some interesting aspects of Dagstuhl seminars in general, that played a critical role in fueling the related talks, discussions and reports. In summary, the 3-day seminar in which we participated, focused not solely on the presentation of established results but also on ideas, sketches, and open (research and operations) problems. The pace and program was guided by topics and presentations that evolved through discussions. This report contains an executive summary of the material that was transcribed during the entire seminar.

Overall, some participants of the seminar seem to be more “pessimistic” about routing security. Both the research and operator communities need to consolidate more data sources to facilitate progress. Any deployment progress is only possible if

operator incentives are improved, however, it remains an open problem on how to provide strong incentives. In practice, a good technical solution is insufficient without first tackling the “politics”. We discussed about routing/network testbeds and the role they can play in emulating and verifying many of the discussed concepts. However, in the wild (or the “real world”), it is surprisingly hard to implement something like RPKI; even more so for BGPsec. We all need a better understanding of the problem space; formal taxonomies of routing attacks, such as hijacks, would be of great help on this front. Regarding improving BGP itself, we have seen many prevention mechanisms, whose deployment is the end-goal for the Internet. However, as we have to live with BGP at least in the intermediate term, we can also explore research on overlay solutions to achieve the properties that we need, at least for the time being. These solutions need to support incremental deployment for obvious reasons.

In general, deployment progress has been slow which is feared not change in the near future. It is reassuring to see that a lot of work is being done in the measurement area; we were also reminded how hard is it to get the ground truth, labelled with useful metadata. Some fundamentally new and secure approaches were discussed, for instance the SCION secure Internet architecture, however, the deployment of new inter-domain routing protocols is very challenging. To improve the deployment incentives of secure routing protocols for operators, the creation of a catalog of routing incidents could be beneficial.

Moreover, it seems that the community may have underestimated the importance of monitoring tools and their utility in the wild. We have learned about new data sets, as well as interesting insights on the Impact of prefix hijacks on the application layer. In general though, we were hoping to see more enthusiasm for new solutions.

Finally, it is worth noting that having a mixed group of researchers and operators is very important to exchange information and discuss potential approaches, which made the seminar an interesting and worthwhile experience.

6.32 Database Architectures for Modern Hardware

Organizers: Peter A. Boncz, Goetz Graefe, Bingsheng He, and Kai-Uwe Sattler
Seminar No. 18251

Date: June 17–22, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.6.63

© Creative Commons BY 3.0 Unported license

© Peter A. Boncz, Goetz Graefe, Bingsheng He, and Kai-Uwe Sattler



Participants: Anastasia Ailamaki, Gustavo Alonso, Witold Andrzejewski, Carsten Binnig, Peter A. Boncz, Philippe Bonnet, Sebastian Breß, Holger Fröning, Goetz Graefe, Bingsheng He, Alfons Kemper, Thomas Leich, Viktor Leis, Daniel Lemire, Justin Levandoski, Stefan Manegold, Klaus Meyer-Wegener, Onur Mutlu, Thomas Neumann, Anisoara Nica, Ippokratis Pandis, Andrew Pavlo, Thilo Pionteck, Holger Pirk, Danica Porobic, Gunter Saake, Ken Salem, Kai-Uwe Sattler, Caetano Sauer, Bernhard Seeger, Evangelia Sitaridi, Jan Skrzypczak, Olaf Spinczyk, Ryan Stutsman, Jürgen Teich, Tianzheng Wang, Zeke Wang, Marcin Zukowski

Over the last years, the social and commercial relevance of efficient data management has led to the development of database systems as foundation of almost all complex software systems. Hence there is a wide acceptance of architectural patterns for database systems which are based on assumptions on classic hardware setups. However, the currently used database concepts and systems are not well prepared to support emerging application domains such as eSciences, Internet of Things or Digital Humanities. From a user's perspective, flexible domain-specific query languages or at least access interfaces are required, novel data models for these application domains have to be integrated, and consistency guarantees which reduce flexibility and performance should be adaptable according to the requirements. Finally, volume, variety, veracity as well as velocity of data caused by ubiquitous sensors have to be mastered by massive scalability and online processing by providing traditional qualities of database systems like consistency, isolation and descriptive query languages. At the same time, current and future hardware trends provide new opportunities such as:

- many-core CPUs: Next-generation CPUs will provide hundreds of compute cores already in the commodity range. In order to allow high degrees of parallelism some architectures already provide hardware support for the necessary synchronization, e.g. transactional memory. However, it is not clear yet how to fully utilize these degrees of parallelism and synchronization mechanism for database processing.
- co-processors like GPU and FPGA: Special-purpose computing units such as GPUs and FPGAs allow for parallelism at much higher degrees accelerating compute-intensive tasks significantly. Moreover, heterogeneous hardware designs such as coupled CPU-FPGA and CPU-GPU architectures represent a trend of close integration between classic hardware and emerging hardware. However, such designs require new architectural concepts for data management.

- novel storage technologies like NVRAM and SSD: Even modern in-memory database system solutions rely mostly on block-based media (e.g. SSD and HDD) for ensuring persistence of data. Emerging memory technologies such as non-volatile memory (NVRAM) promise byte-addressable persistence with latencies close to DRAM. Currently, the usage of this technology is discussed for instant failure recovery of databases, but the role of NVRAM in future data management system architectures is still open.
- high-speed networks: Both in scale-up and scale-out scenarios efficient interconnects play a crucial role. Today, high-speed networks based on 10 Gbit/s Ethernet or InfiniBand support already Remote DMA, i.e. direct access to memory of a remote node. However, this requires to deal with distributed systems properties (unreliability, locality) and it is still unclear how database systems can utilize this mechanism.

In order to open up the exemplarily mentioned application domains together with exploiting the potential of future hardware generations it becomes necessary now to fundamentally rethink current database architectures.

One of the main challenges of this rethinking is that it requires expertise from different research disciplines: hardware design, computer architectures, networking, operating systems, distributed systems, software engineering, and database systems.

Thus, the goal of this Dagstuhl Seminar was to bring together researchers and practitioners from these areas representing both the software and hardware sides and therefore different disciplines to foster cross-cutting architectural discussions. In this way, the seminar extended the series of previous Dagstuhl seminars on database systems aspects, such as “Robust Query Processing” (10381, 12321, 17222) as well as “Databases on Future Hardware” (17101).

The seminar was organized into six working groups where

the participants discussed opportunities and challenges in order to exploit different features of modern hardware and operating system primitives for data processing:

- Database accelerators: Based on an analysis of use cases for database accelerators from the level of individual operators and algorithms up to the level of complex database tasks, the group discussed ways of exploiting and evaluating accelerator technologies as well as future research directions with respect to hardware acceleration in databases.
- Memory hierarchies: The group discussed design recipes for database nodes with non-trivial memory hierarchies containing not only disk and RAM but also non-volatile memory. Within such a hierarchy different caching strategies are employed: exclusive caching for functionally equivalent levels and inclusive caching for levels with different functionality.
- Remote direct memory access: The group discussed ways of exploiting RDMA in data-intensive applications. Particularly, an interface providing a set of useful abstractions for network-aware data-intensive processing called DPI was proposed. Similar to MPI, DPI is designed as an interface that can have multiple implementations for different networking technologies to enable the exploitation of RDMA and in-network processing.

- Heterogeneous database architectures: This topic was addressed by two working groups. Both groups discussed a database software architecture that is capable of making use of multiple hardware devices (GPU, TPU, FPGA, ASICs), in addition to the CPU for handling database workloads. The principle goal was an architecture that would never be worse than a state-of-the-art CPU-centered database architecture, but would get significant benefit on those workloads were the heterogeneous devices can exploit their strengths. The first group developed a morsel-driven architecture, where pipelines are broken up into sub-pipelines and adaptive execution strategies are exploited. The second group discussed operating system support and primitives for heterogeneous architectures.
- Machine learning in database systems: The goal of this working group was to investigate the application of machine learning methods for estimating operator selectivities as part of query optimization. Such an approach could overcome the inaccuracies of traditional cost estimation techniques especially for queries comprised of complex predicates and multiple joins.

The progress and outcome of the individual working groups was presented in a daily plenary session, details of the results are given below.

References

- 1 Gustavo Alonso, Michaela Blott, Jens Teubner: *Databases on Future Hardware* (Dagstuhl Seminar 17101). Dagstuhl Reports 7(3):1–18 (2017)
- 2 Renata Borovica-Gajic, Goetz Graefe, Allison Lee: *Robust Performance in Database Query Processing* (Dagstuhl Seminar 17222). Dagstuhl Reports 7(5):169–180 (2017)
- 3 Goetz Graefe, Wey Guy, Harumi A. Kuno, Glenn N. Paulley: *Robust Query Processing* (Dagstuhl Seminar 12321). Dagstuhl Reports 2(8):1–15 (2012)
- 4 Goetz Graefe, Arnd Christian König, Harumi Anne Kuno, Volker Markl, Kai-Uwe Sattler: *Robust Query Processing*. Dagstuhl Seminar Proceedings 10381, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany 2010

6.33 Ubiquitous Gaze Sensing and Interaction

Organizers: Lewis Chuang, Andrew Duchowski, Pernilla Qvarfordt, and Daniel Weiskopf
Seminar No. 18252

Date: June 18–21, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.6.77

© Creative Commons BY 3.0 Unported license

© Tanja Blascheck, Lewis Chuang, Andrew Duchowski, Pernilla Qvarfordt, and Daniel Weiskopf



Participants: Amy Alberts, M. Stella Atkins, Roman Bednarik, Hans-Joachim Bieg, Maria Bielikova, Leslie Blaha, Tanja Blascheck, Andreas Bulling, Lewis Chuang, Andrew Duchowski, Sara Irina Fabrikant, Nina Gehrer, Hans Gellersen, Kenneth Holmqvist, Eakta Jain, Radu Jianu, Enkelejda Kasneci, Peter Kiefer, Krzysztof Krejtz, Kuno Kurzhals, David P. Luebke, Radoslaw Mantiuk, Diako Mardanbegi, Thies Pfeiffer, Pernilla Qvarfordt, Michael Raschke, Martin Raubal, Laura Trutoiu, Daniel Weiskopf

The miniaturization of optical devices and advances in computer vision, as well as a lower cost point, have led to an increased integration of gaze sensing capabilities in computing systems, from desktop computing to mobile devices and wearables. With these advances in technology, new application areas for gaze sensing are emerging. Eye tracking is no longer restricted to a well-controlled laboratory setting, but moving into everyday settings. When technology makes forays into new environments, there are many questions to be resolved and challenges to be met, from computational to applications and interaction. Ubiquitous gaze sensing and interaction require a framework that can accommodate compatible solutions from data acquisition to signal processing to pattern classification and computer vision to visualization and analytics. Including gaze data into interactive applications requires knowledge of natural gaze behaviors as well as how gaze is coordinate with other modalities and actions.

Therefore, this Dagstuhl Seminar brought together computer scientists and gaze researchers to explore future ubiquitous applications and to identify requirements for reliable gaze sensing technology. Ubiquitous gaze sensing and interaction cannot be achieved by research discipline, but require knowledge and scientific advancement in multiple fields. And, of utmost importance is that researchers from different disciplines meet, interact, and address their common challenges. For this reason, experts in computer graphics, signal processing, visualization, human-computer interaction, data analytics, pattern analysis and classification along with researchers who employ gaze tracking across diverse disciplines attended: geo-information systems, medicine, aviation, psychology, neuroscience, etc. This fostered a dialogue and allowed: (1) computing scientists to understand the problems that are faced in recording and interpreting gaze data, (2) gaze researchers to consider how modern computing techniques could potentially advance their research. In addition,

we discussed the ethical and privacy concerns of deploying gaze monitoring devices in everyday scenarios.

The workshop was organized to identify identifying possible **scenarios** and pinpointing the associated **challenges** of developing and deploying ubiquitous gaze sensing during the first day. Challenges identified by multiple scenarios, or the ones that were considered to be significant were the focus of in-depth cross-disciplinary groups. These challenges were discussed on the second day. In three sessions taking place during the day, five challenges were debated. “Data Privacy” and “Gaze + X” were two of the most important topics and received multiple dedicated sessions of discussion due to the high interest of the participants.

On the third day the Dagstuhl Seminar finally discussed future work and how to get the research community engaged in researching the various interesting topics covered. Some of the suggestions were to organize workshops at conferences and organizing a special issue focused on ubiquitous gaze sensing. Several of the discussion groups started brainstorming on papers covering the important topics raised at the workshop.

6.34 Discipline Convergence in Networked Systems

Organizers: Yungang Bao, Lars Eggert, Simon Peter, and Noa Zilberman
Seminar No. 18261

Date: June 24–29, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.6.149

© Creative Commons BY 3.0 Unported license

© Yungang Bao, Lars Eggert, Simon Peter, and Noa Zilberman

Participants: Gustavo Alonso, Yungang Bao, Claude Bartheles, Angelos Bilas, Pietro Bressana, Trevor Carlson, Julian Chesterfield, Dilma Da Silva, Felix Eberhardt, Lars Eggert, Tim Harris, David Hay, Matthias Hille, Timo Hönig, Michio Honda, Stefan Klauck, Dirk Kutscher, Giuseppe Lettieri, Sue Moon, Jacob Nelson, Jörg Ott, Simon Peter, Max Plauth, Dan Ports, Timothy Roscoe, Henning Schulzrinne, Golan Schzukin, Leendert van Doorn, Eric Van Hensbergen, Irene Y. Zhang, Noa Zilberman



Networked computing systems have reached a watershed, as the amount of networked-data generated by user applications exceeds the processing capability of any single computer. This requires an integrated system design, unlike the traditional layered approaches. This seminar therefore brought together experts from the operating systems, distributed systems, computer architecture, networks, storage and databases communities, to advance the state of the art in discipline convergence in networked systems.

The networking community has advanced in giant leaps, making high bandwidth networking and software-defined networking (SDN) commodity. Furthermore, the advent of network function virtualization (NFV) has started the convergence of computing technologies and networking technologies. The computing community, on the other hand, struggled to overcome power density limitations, resource- efficiency and quality-of-service etc. for cloud computing as well as end host computing (or edge computing), and cannot keep up.

Revolutionary networked system design approaches are now emerging, seeking to increase performance, efficiency and security through the convergence of disciplines: compute, storage and networking. This seminar investigated both hardware and software challenges, and attempted to bridge the gaps between different communities in order to compensate the challenges in some areas with emerging breakthroughs from other areas. Over the course of the 5-day seminar, seventeen presentations were given on various aspects of data center networking. Taking the presentations as input, the workshop then broke into five working groups to discuss research aspects of operating systems, distributed systems, computer architecture, networks, storage, and databases. The talks as well as the outcome of the breakout session and the concluding statements are summarized in this report.

6.35 10 Years of Web Science: Closing The Loop

Organizers: Susan Halford, James A. Hendler, Eirini Ntoutsis, and Steffen Staab
Seminar No. 18262

Date: June 24–29, 2018 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.8.6.173

© Creative Commons BY 3.0 Unported license

© Eirini Ntoutsis



Participants: Robert Ackland, Ricardo A. Baeza-Yates, Bettina Berendt, Noshir S. Contractor, David De Roure, Kemal A. Delic, Nikolaus Forgó, Fabien Gandon, Susan Halford, Wendy Hall, Lynda Hardman, Andreas Hotho, Katharina E. Kinder-Kurlanda, Claudia Müller-Birn, Wolfgang Nejdl, Eirini Ntoutsis, Paolo Parigi, Evaggelia Pitoura, Oshani Seneviratne, Elena Simperl, Steffen Staab, Guglielmo Tamburrini, Pinelopi Troullinou

This Dagstuhl Seminar aimed at bringing together researchers from different disciplines related to Web Science, namely computer science, sociology, philosophy and law to discuss on future of Web Science and how it can stay faithful to its initial mission for societal good. Several recent incidents like the online psychological experiment by Facebook have provoked widespread public concern regarding the effect of such experiments and interventions and there is no agreement on expertise and ethics knowledge about how to do Web experimental research.

The Web is a complex sociotechnical system where humans and (intelligent) machines interact in unexpected ways; such hybrid societies of natural and artificial intelligence raise new challenges for Web Science which go beyond technical challenges into ethical, legal and societal implications. The role of Artificial Intelligence in these developments was discussed extensively in terms of both opportunities and risks.

Based on the discussions and inputs from all participants, we have split the discussion into three main working groups:

- Working group on innovative methods for Web Science
- Working group on values
- Working group on Web Science and Artificial Intelligence

The group will continue its work in the aforementioned topics and a manifesto is foreseen to be ready by the end of the year.

6.36 In Situ Visualization for Computational Science

Organizers: Janine C. Bennett and Hank Childs and Christoph Garth and Bernd Hentschel
Seminar No. 18271

Date: July 1–6, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.7.1

© Creative Commons BY 3.0 Unported license

© Janine C. Bennett, Hank Childs, Christoph Garth, Bernd Hentschel

Participants: Andrew Bauer, Janine C. Bennett, E. Wes Bethel, Peer-Timo Bremer, Thierry Carrard, Hank Childs, Matthieu Dorier, Steffen Frey, Christoph Garth, Nicolas R. Gauger, Markus Hadwiger, Charles D. Hansen, Katrin Heitmann, Bernd Hentschel, Ingrid Hotz, Katherine E. Isaacs, Jens Krüger, Matthew Larsen, Peter Messmer, Kenneth Moreland, Benson Muite, Kenji Ono, Manish Parashar, Valerio Pascucci, John Patchett, Tom Peterka, Dirk Pleiter, David Pugmire, Bruno Raffin, Alejandro Ribes Cortes, Niklas Röber, Ulrich Rüde, Filip Sadlo, Han-Wei Shen, Robert Sisneros, Madhusudhanan Srinivasan, Gunther H. Weber, Rüdiger Westermann, Hongfeng Yu



The workshop identified ten challenges for in situ processing that require significant research. These challenges were identified by spending the first day of the workshop with participants giving short presentations on their experiences with in situ processing, with a special focus on unsolved problems. The participant perspectives were then organized into the ten research challenges. Over the following days, sub-groups discussed each of the ten challenges and then presented the key points of their discussions to the group and received feedback. Shortly after the workshop, the leaders of each sub-group wrote summaries for its associated research challenge; these summaries are the basis of this report.

The ten challenges identified by our participants were:

- Data quality and reduction, i.e., reducing data in situ and then exploring it post hoc, which is likely the form that will enable exploration of large data sets on future supercomputers.
- Workflow specification, i.e., how to specify the composition of different tools and applications to facilitate the in situ discovery process.
- Workflow execution, i.e., how to efficiently execute specified workflows, including workflows that are very complex.
- Exascale systems, which will have billion-way concurrency and disks that are slow relative to their ability to generate data.
- Algorithmic challenges, i.e., algorithms will need to integrate into in situ ecosystems and still perform efficiently.
- Use cases beyond exploratory analysis, i.e., ensembles for uncertainty quantification and decision optimization, computational steering, incorporation of other data sources, etc.
- Exascale data, i.e., the data produced by simulations on exascale machines will, in many cases, be fundamentally different than that of previous machines.
- Cost models, which can be used to predict performance before executing an algorithm and thus be used to optimize performance overall.
- The convergence of HPC and Big Data for visualization

and analysis, i.e., how can developments in one field, such as machine learning for Big Data, be used to accelerate techniques in the other?

- Software complexity, heterogeneity, and user-facing issues, i.e., the challenges that prevent user adoption of in situ techniques because in situ software is complex, computational resources are complex, etc.

From group discussion, two other important topics emerged that do not directly lead to open research questions, but rather are concerned with effective organization of the often highly interdisciplinary research into in situ techniques. To address these, two panels were held to facilitate effective discussion. Finally, the workshop featured technical presentations by participants on recent results related to in situ visualization.

6.37 Synergies between Adaptive Analysis of Algorithms, Parameterized Complexity, Compressed Data Structures and Compressed Indices

Organizers: J r my Barbay, Johannes Fischer, Stefan Kratsch, and Srinivasa Rao Satti
Seminar No. 18281

Date: July 8–13, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.7.44

  Creative Commons BY 3.0 Unported license

  J r my Barbay, Johannes Fischer, Stefan Kratsch, and Srinivasa Rao Satti



Participants: J r my Barbay, Philip Bille, Stefan B ttcher, Luca Castelli Aleardi, Stephane Durocher, Johannes Fischer, Till Fluschnik, Allyx Fontaine, Travis Gagie, Simon Gog, Meng He, Falko Hegerfeld, Shunsuke Inenaga, Bart Jansen, Artur Jez, Seungbum Jo, Ahmet Kara, David G. Kirkpatrick, Christian Knauer, Dominik K ppl, Stefan Kratsch, Florian Kurpicz, Zsuzsanna Lipt k, Sebastian Maneth, Ian Munro, Yakov Nekrich, Patrick K. Nicholson, Yoshio Okamoto, Ramamohan Paturi, Nicola Prezza, Rajeev Raman, Venkatesh Raman, Srinivasa Rao Satti, Mireille Regnier, Giovanna Rosone, Raimund Seidel, Jouni Sir n, Tatiana Starikovskaya, Rossano Venturini, Sandra Zilles

Seminar 18281, about the ‘‘Synergies between Adaptive Analysis of Algorithms, Parameterized Complexity, Compressed Data Structures and Compressed Indices’’, gathered researchers from four distinct research areas (with some researchers having results in up to three such areas, but none in all four):

1. the area of adaptive analysis of algorithms;
2. the study of parameterized complexity of NP-hard problems;
3. the area focused on compressed data structures; and
4. the area concerned with the study of compressed indices.

Goals The intuition behind gathering people from such diverse communities was that while all of these subareas of algorithms and data structures focus on ‘‘going beyond the worst-case’’ for classes of structurally restricted inputs, there has been a limited amount of interactions between them, and some results have been ‘‘discovered’’ twice. Therefore, the main goal of the seminar was to share knowledge and make joint progress through dedicated survey talks and plenty of time for discussions and work on open problems.

Structure The seminar consisted of

1. a first session of personal introductions, each participant presenting his expertise and themes of interests in two slides;
2. a small series of technical talks, some organized a long time in advance, and some improvised ‘‘on demand’’; and
3. a larger series of presentation of open problems, with ample time left for the participants to gather and work on such open problems.

Conclusion Most participants concurred that they learned a lot from the seminar, and acquired new contacts to foster further collaborations. In particular, interactions between the adaptive analysis of algorithms and the study of the parameterized complexity of NP-hard problems seemed relevant to the recent

development of conditional lower bounds for problems classically solved in polynomial time, an approach referred to as ‘‘Fine Grained Analysis’’ or ‘‘FPT in P’’.

Generally, it appears that the seminar struck a good balance between scheduled sessions for survey talks and presentation of open problems as well as free time for discussion and interaction. During the free time, many smaller groups got together for work on open problems or for informal presentations of more specialist topics with a smaller audience. We think that this setup, along with the longer than usual round of introductions on the first day, was very successful at bringing together the different research areas.



Fig. 6.10

“@dagstuhl manor in InfraRed. #photooftheday” Twitter post by 18271 Dagstuhl Seminar participant Madhu Srinivasan.

<https://twitter.com/vyslexic/status/1013522346236416001>.

The photo was taken with a specially modified “full-spectrum” Nikon DSLR. Photo courtesy of Madhu Srinivasan.

6.38 Extreme Classification

Organizers: Samy Bengio, Krzysztof Dembczyński, Thorsten Joachims, Marius Kloft, and Manik Varma

Seminar No. 18291

Date: July 15–20, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.7.62

© Creative Commons BY 3.0 Unported license

© Samy Bengio, Krzysztof Dembczyński, Thorsten Joachims, Marius Kloft, and Manik Varma



Participants: Maximilian Alber, Rohit Babbar, Samy Bengio, Alexander Binder, Evgenii Chzhen, Kunal Dahiya, Krzysztof Dembczyński, Urun Dogan, Matthias Enders, Asja Fischer, Johannes Fürnkranz, Thomas Gärtner, Edouard Grave, Yann Guermeur, Eyke Hüllermeier, Christian Igel, Himanshu Jain, Kalina Jasinska, Armand Joulin, Nikos Karampatziakis, Matthias Kirchler, Marius Kloft, Christoph H. Lampert, John Langford, Antoine Ledent, Christoph Lippert, Nicolas Mayoraz, Jinseok Nam, Alexandru Niculescu-Mizil, Yashoteja Prabhu, Pradeep Ravikumar, Adith Swaminathan, Manik Varma, Willem Waegeman, Marek Wydmuch

The topic of this seminar is in the general context of machine learning [10] which concerns the study and development of algorithms that learn from empirical data how to make accurate predictions about yet unseen data without being explicitly programmed. Multi-class and multi-label learning are classical problems in machine learning. The outputs here stem from a finite set of categories (classes), and the aim is to classify each input into one (multi-class) or multiple (multi-label) out of several possible target classes. Classical applications of multi-class and multi-label learning include handwritten optical character recognition [8], part-of-speech tagging [11], and text categorization [7]. However, with the advent of the big data era, learning problems can involve even millions of classes. As examples let us consider the following problems:

- Person recognition in Facebook images (there are billions of Facebook users; given an image, we might want to predict the subset of users present in the image for such applications like security, surveillance, social network analysis, etc.).
- Predicting Wikipedia tags for new Wikipedia articles or webpages (Wikipedia has almost 2 million tags now).
- Recommending Amazon items where each of the 100 million items on Amazon is a separate label.
- Search on Google/Bing where each of the 100 million queries is a separate label.
- Language modelling – predicting the next word in a sentence from the millions of words available.

The problems of this type are often referred to as *extreme classification*. They have posed new computational and statistical challenges and opened a new line of research within machine learning.

The main goal of extreme classification is to design learning and prediction algorithms, characterized by strong statistical guarantees, that exhibit sublinear time and space complexity in the number of classes. Unfortunately, the theoretical results obtained

so far are still not satisfactory and very limited. Moreover, the problems at this scale often suffer from unreliable learning information, e.g., there is no chance to identify all positive labels and assign them precisely to training examples. The majority of labels is used very rarely, which leads to the problem of the long-tail distribution. In practical applications, learning algorithms run in rapidly changing environments. Hence, during testing/prediction phase new labels might appear that have not been present in the training set [2, 4]. This is the so-called zero-shot learning problem. Furthermore, typical performance measures used to assess the prediction quality of learning algorithms, such as 0/1 or Hamming loss, do not fit well to the nature of extreme classification problems. Therefore, other measures are often used such as precision@k [9] or the F-measure [6]. However, none of the above is appropriate to measure predictive performance in the long-tail problems or in the zero-shot setting. Hence, the goal is to design measures, which promote a high coverage of sparse labels [5].

The seminar aimed at bringing together researchers interested in extreme classification to encourage discussion on the above mentioned problems, identify the most important ones and promising research directions, foster collaboration and improve upon the state-of-the-art algorithms. The meeting in this regard was very successful as participants from both academia and industry as well as researchers from both core machine learning and applied areas such as recommender systems, computer vision, computational advertising, information retrieval and natural language processing, were given the opportunity to see similar problems from different angles.

The seminar consisted of invited talks, working groups, presentation of their results, and many informal discussions. The talks concerned among others such topics as: common applications of extreme classification, potential applications in bioinformatics and biotechnology, neural networks for extreme

classification, learning theory for problems with a large number of labels, approaches for dealing with tail labels, learning and prediction algorithms, extreme classification challenges in natural language processing, multi-task learning with large number of tasks, pitfalls of multi-class classification, recommendation systems and their connection to extreme classification, counterfactual learning and zero-shot learning. The short abstracts of these talks can be found below in this report. The four working groups focused on the following problems: loss functions and types of predictions in multi-label classification, deep networks for extreme classification, zero-shot learning and long tail labels, and generalization bounds and log-time-and-space algorithms. Short summaries of the results obtained by the working groups can also be found below.

During the seminar, we also discussed different definitions of extreme classification. The basic one determines extreme classification as a multi-class or multi-label problem with a very

large number of labels. The labels are rather typical identifiers without any explicit meaning. However, there usually exists some additional information about *similarities* between the labels (or this information can be extracted or learned from data). From this point of view, we can treat extreme classification as a learning problem with a weak structure over the labels. This is in difference to structured output prediction [1], where we assume much stronger knowledge about the structure. The most general definition, however, says that extreme classification concerns all problems with an extreme number of choices.

The talks, working groups, and discussions have helped to gain a better understanding of existing algorithms, theoretical challenges, and practical problems not yet solved. We believe that the seminar has initiated many new collaborations and strengthen the existing ones that will soon deliver new results for the extreme classification problems.

References

- 1 Bakir GH, Hofmann T, Schölkopf B, Smola AJ, Taskar B, Vishwanathan SVN (eds) Predicting Structured Data. Neural Information Processing, The MIT Press, 2007
- 2 Andrea Frome, Gregory S. Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc'Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems*, pages 2121–2129, 2013.
- 3 Yann Guermeur. l_p -norm Sauer-Shelah lemma for margin multi-category classifiers. *arXiv preprint arXiv:1609.07953*, 2016.
- 4 Bharath Hariharan, S.V.N. Vishwanathan, and Manik Varma. Efficient max-margin multi-label classification with applications to zero-shot learning. *Machine Learning*, 88(1-2):127–155, 2012.
- 5 Himanshu Jain, Yashoteja Prabhu, and Manik Varma. Extreme multi-label loss functions for recommendation, tagging, ranking & other missing label applications. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, pages 935–944, New York, NY, USA, 2016. ACM.
- 6 K. Jasinska, K. Dembczyński, R. Busa-Fekete, K. Pfannschmidt, T. Klerx, and E. Hüllermeier. Extreme F-measure maximization using sparse probability estimates. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning – Volume 48*, pages 1435–1444, 2016.
- 7 Thorsten Joachims. *Text categorization with support vector machines: Learning with many relevant features*. Springer, 1998.
- 8 Nei Kato, Masato Suzuki, Shin Ichiro Omachi, Hirotomo Aso, and Yoshiaki Nemoto. A handwritten character recognition system using directional element feature and asymmetric mahalanobis distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(3):258–262, 1999.
- 9 Yashoteja Prabhu and Manik Varma. Fastxml: A fast, accurate and stable tree-classifier for extreme multi-label learning. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 263–272. ACM, 2014.
- 10 Vladimir Vapnik. *Statistical learning theory*, volume 1. Wiley New York, 1998.
- 11 Atro Voutilainen. Part-of-speech tagging. *The Oxford handbook of computational linguistics*, pages 219–232, 2003.

6.39 Web Application Security

Organizers: Martin Johns, Nick Nikiforakis, Melanie Volkamer, and John Wilander
Seminar No. 18321

Date: August 5–8, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.1

© Creative Commons BY 3.0 Unported license

© Martin Johns, Nick Nikiforakis, Melanie Volkamer, John Wilander



Participants: Frederik Braun, Achim D. Brucker, Stefano Calzavara, Luca Compagna, Lieven Desmet, Steven Englehardt, Thomas Gross, Marian Harbach, Daniel Hausknecht, John Hazen, Mario Heiderich, Boris Hemkemeier, Martin Johns, Christoph Kerschbaumer, Pierre Laperdrix, Sebastian Lekies, Benjamin Livshits, Matteo Maffei, Marius Musch, Nick Nikiforakis, Lukasz Olejnik, Juan David Parra, Giancarlo Pellegrino, Karen Renaud, Tamara Rezk, Konrad Rieck, Andrei Sabelfeld, Sebastian Schinzel, Zubair Shafiq, Lynsay Shepherd, Dolière Francis Somé, Ben Stock, Daniel Veditz, Melanie Volkamer, Malte Wedel, Rigo Wenning, Mike West, John Wilander, Henrik Willert

■ Introduction

Motivation Since its birth in 1990, the Web has evolved from a simple, stateless delivery mechanism for static hyper-text documents to a fully-fledged run-time environment for distributed, multi-party applications. Even today, there is still a continuous demand for new features and capabilities which drives the Web’s evolution onwards. This unplanned and often chaotic development has led to several deeply ingrained security and privacy problems that plague the platform:

- The Web’s original hypertext, multi-origin nature which is manifested in the design of HTML and HTTP is in fundamental conflict with JavaScript’s Same-Origin Policy, the Web’s most important security mechanism.
- Important security properties, such as end-to-end communication security or endpoint identity are outside of the control of the actual applications. Instead, they depend on the security of external entities, such as domain name servers or certificate authorities.
- Data/code separation in web applications is practically infeasible, as the HTTP link between server-side application logic and client-side application interface requires an intermixing of protocol, data and code fragments within a single continuous character stream.
- HTTP is a stateless protocol without a native session or authentication tracking concept.
- Users are not aware of general or application specific threats. Protecting against these threats (incl. to know which security indicators to trust) is nowadays difficult and time consuming.

Using this fragile basis, critical applications are created, that long have left the strict client-server paradigm, on which the Web was initially built. Instead, scenarios are realized that involve several mutually distrusting entities in a single security and application context. In many cases the browser is the link that

connects the remote parties, either via direct JavaScript inclusion, web mashups, or through the usage of web protocols, such as OpenID and OAuth.

The accumulated ballast of the last two decades of web evolution, the ever growing functional demands of sophisticated web applications and the ambitious vision of the web platform’s drivers creates an exciting tension field which is in constant conflict with the required security assurances of high value business applications.

Since approximately ten years, academic security and privacy research has recognized the importance of the web platform and the unique characteristics and challenges of the web security and privacy topic. And while specific techniques, that originated from academic research, such as the Content Security Policy, have been adapted in practice, the fundamental security problems of the web remain and the overall vulnerability landscape is getting worse, as it can be seen in the constant flow of reported web security issues in bug trackers and vulnerability databases.

Academic web security research has started 2007 and usable security research started almost at the same time. In the context of this Dagstuhl Seminar, we will revisit the lessons learned from the last decade and revisit the success stories and mistakes that have been made. Questions, that have to be raised in include “What has worked?”, “What has been taken up by industry?”, “What failed and why?”, and – most importantly – “What did we learn?”

Seminar Objectives Today, several unconnected groups drive the topic, including Security, Privacy as well as Usable Security & Privacy Academics, standardization, and browser vendors. The seminar will facilitate essential exchange between them. This will allow academia to directly influence browser vendors and standardization representatives, and allow industry representatives to influence the research community.

■ Overview

Participants The seminar was well attended with 39 participants. A good balance of European and American researchers was present. Furthermore, the group represented a nice mix of participants of academia and industry. Compared to the previous editions, not only researchers from the web security area participated but also from the field of human factors in security.

Structure This was the third Dagstuhl seminar on Web application security. The seminar's organisation combined overview presentation of various subfields, highlight talks, and discussions in working groups. In particular the overview presentations were important to connect the two research fields web security from a more technical point of view and human factors in security. This way, also a good, comprehensive view on current activities and open problems in the realm of Web application security in particular from a user's point of view could be achieved and areas for potential future collaborations could be identified.

■ Summary

Talks The following people presented either an overview of their research field, very recent research results or overarching observations on the field of web application security. Please also refer to Section 3 for selected talk abstracts.

- Stefano Calzavara, University of Venezia, IT: REASON – A programmable architecture for secure browsing
- Luca Compagna, SAP Labs France – Mougins, FR: Analysis & Detection of Authentication Cross-Site Request Forgeries
- Lieven Desmet, KU Leuven, BE: Detecting and Preventing Malicious Domain Registrations in the .eu TLD
- Steven Englehardt, Mozilla – Mountain View, US: No Boundaries: Data exfiltration by directly embedded tracking scripts
- Thomas Gross, Newcastle University, GB: Investigating Cognitive and Affective Predictors Impacting Password Choice
- Mario Heiderich, Cure53 – Berlin, DE, DOMPurify: Client-Side Protection Against XSS and Markup Injection
- Boris Hemkemeier, Commerzbank AG – Frankfurt, DE: Web application security in vulnerable environments
- Martin Johns, TU Braunschweig, DE: WebAppSec @ Dagstuhl – The Third Iteration
- Christoph Kerschbaumer, Mozilla – San Francisco, US: Could we use Information Flow Tracking to generate more sophisticated blacklists?
- Pierre Laperdrix, Stony Brook University, US: Browser fingerprinting: current state and possible future
- Sebastian Lekies, Google Switzerland – Zürich, CH: Trusted Types: Prevent XSS with this one simple trick!
- Benjamin Livshits, Imperial College London, GB: Browser Extensions for the Web of Value
- Marius Musch, TU Braunschweig, DE: On measurement studies and reproducibility
- Lukasz Olejnik, Independent researcher, W3C TAG, FR: Private browsing modes guaranteed. On the example of Payment Request API
- Juan David Parra, Universität Passau, DE: Computational Resource Abuse through the Browser
- Giancarlo Pellegrino, Stanford University, US: Removing Browsers from the Equation: A New Direction for Web Application Security
- Tamara Rezk, INRIA Sophia Antipolis, FR: Content Security Policy Challenges

- Konrad Rieck, TU Braunschweig, DE: Beyond the Hype: Web Security and Machine Learning?
- Andrei Sabelfeld, Chalmers University of Technology – Göteborg, SE: A Challenge for Web of Things: Securing IoT Apps
- Sebastian Schinzel, FH Münster, DE: Handling HTML Emails after the Efail Attacks
- Zubair Shafiq, University of Iowa – Iowa City, US: The Arms Race between Ad Tech vs. Adblockers: Key Challenges and Opportunities
- Lysnay Shepherd, Abertay University – Dundee, GB: How to Design Browser Security and Privacy Alerts
- Dolière Francis Somé, INRIA Sophia Antipolis, FR: The Same Origin Policy and Browser Extensions
- Ben Stock, CISPA – Saarbrücken, DE: Persistent Client-Side Cross-Site Scripting in the Wild
- Melanie Volkamer, KIT – Karlsruher Institut für Technologie, DE: Web Security Meets Human Factors in Security
- Mike West, Google – München, DE: HTTP State Tokens

■ Conclusions

This seminar was the third Dagstuhl Seminar von Web Application Security, following Seminar 09141 (2009) and Seminar 12401 (2012). Thus, it was a great opportunity to reflect on a decade of web security research. In 2009 the field was largely undefined and that year's seminar offered a wild mix of various topics, some with lasting impact and many that went nowhere. Where the 2009 seminar was overly broad, the 2012 iteration had a comparatively narrow focus as the seminar was dominated by the notion that solving web security mainly revolves around solving the security properties of JavaScript.

This year's seminar reflected the ongoing maturing of the topic very well. Fundamental problems, such as Cross-site Scripting or the Web Browser security model, are well explored and their understanding served as a great foundation for the seminar's discussions. This allowed the extension of the topic toward important facets, such as privacy problems or human factors. While the addressed topics were too broad and the time for overarching discussions was limited due to the three-day format of the seminar, the sparked discussions were fruitful for several follow-up activities (see above). An underlying theme of the seminar can be summarized as “the last decade of web security has broad good progress and development but the overall problem is still neither fully understood nor solved”. Especially, the newly introduced dimension of integrating human factors in security, which was reflected through including several high-profile members of this community in the seminar, is still immature.

One of the seminar's prime objectives has been reached very nicely: The fostering of collaboration between the different web security communities. For one, several compelling interactions between practitioners from industry (such as SAP, Commerzbank and Cure53) and researcher from academia took place. Furthermore, thanks to the fact that all major web browser vendors (plus the new privacy-centric browser Brave) were represented at the seminar, both cross-browser vendor interaction as well as browser/academia collaborations were initiated, with the browser-based sanitizer initiative (see breakout session 4.3 of the full report) being a prominent example.

6.40 Human-Computer Integration

Organizers: Jonathan Grudin, Pattie Maes, and Florian Mueller
Seminar No. 18322

Date: August 5–10, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.18

© Creative Commons BY 3.0 Unported license

© Florian Mueller, Jonathan Grudin, Pattie Maes, Zhuying Li



Participants: Tom Erickson, Elizabeth Gerber, Steven Greenspan, Stefan Greuter, Jonathan Grudin, Ti Hoang, Masahiko Inami, Wendy Ju, Kai Kunze, Joseph La Delfa, Zhuying Li, Pedro Lopes, Pattie Maes, Joseph Marshall, Jochen Meyer, Florian Mueller, Suranga Nanayakkara, Jun Nishida, Marianna Obrist, Harald Reiterer, Thecla Schiphorst, Caitlyn E. Seim, Jürgen Steimle, Paul Strohmeier, Dag Svanaes, Dakuo Wang, Martin Weigel, Katrin Wolf

The rise of technology that supports a partnership between user and computer highlights an opportunity for a new era of “human-computer integration”, contrasting the previously dominant paradigm of computers functioning as tools. However, most work around these technologies only focused on the instrumental perspective to achieve extrinsic performance objectives. However, phenomenology emphasizes that it is also important to support the experiential perspective, which indicates that technology should also help people pay attention to their lived experiences and personal growth in order to deepen their understanding of their own bodies. This seminar focuses on embodied integration, where a computer tightly integrates with the person’s body. Although an increasing number of systems are emerging, a thorough understanding of how to design such systems is notably absent. The reason for this is the limited knowledge about how such embodied partnerships unfold, and what underlying theory could guide such developments. This seminar brought together leading experts from industry and academia, including those who are central to the development of products and ideas such as wearables, on-body robotics, and exertion systems. The goal was to address key questions around the design of embodied integration and to jump-start collaborations to pioneer new approaches for a human-computer integrated future.



Fig. 6.11

“**Somewhere over the rainbow @dagstuhl**” Twitter post by 18332 Dagstuhl Seminar participant Jan Mendling.

<https://twitter.com/janmendling/status/1029067772855242752>. Photo courtesy of Jan Mendling.

6.41 Algorithmic Foundations of Programmable Matter

Organizers: Spring Berman, Sándor P. Fekete, Matthew J. Patitz, and Christian Scheideler
Seminar No. 18331

Date: August 12–17, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.48

© Creative Commons BY 3.0 Unported license

© Spring Berman, Sándor P. Fekete, Matt Patitz, and Christian Scheideler



Participants: Aaron Becker, Spring Berman, Julien Bourgeois, Luca Cardelli, Kenneth C. Cheung, Joshua J. Daymude, Erik D. Demaine, David Doty, Sándor Fekete, Roderich Gross, Dan Halperin, Heiko Hamann, Kristian Hinnenthal, Lila Kari, MinJun Kim, Irina Kostitsyna, Dominik Krupke, Alcherio Martinoli, Friedhelm Meyer auf der Heide, Othon Michail, Joseph S. B. Mitchell, Nils Napp, Ram Prasad Narayanan, Pekka Orponen, Matthew J. Patitz, Andréa Richa, Marcel J. M. Roeloffzen, Kay Römer, Trent Rogers, Dorian Rudolph, Christian Scheideler, Stefan Schmid, Arne Schmidt, Chris Thachuk, Pierre Thalamy, André van Renssen, Jennifer L. Welch

The term “programmable matter” refers to any substance that can change its physical properties (shape, density, moduli, conductivity, optical properties, etc.) in a programmable fashion. The role of *algorithmic foundations* of programmable matter continues to grow in importance due to ongoing progress in a wide range of applications. Examples of cutting-edge application areas with a strong algorithmic flavor include *self-assembling systems*, in which chemical and biological substances such as DNA are designed to form predetermined shapes or carry out massively parallel computations; and *swarm robotics*, in which complex tasks are achieved through the local interactions of robots with highly limited individual capabilities, including micro- and nano-robots. Progress in these application areas has been achieved through close collaboration with algorithmic theoreticians, enabling the investigation of fundamental problems related to system geometry using methods from the field of *computational geometry*, and yielding techniques for decentralized computation from the field of *distributed computing*.

A previous Dagstuhl seminar (16271, Algorithmic Foundations of Programmable Matter) had laid the foundations for further progress by bringing together experts from different fields and focusing on expert surveys and breakout groups. We built on the success of that seminar by expanding its focus on particular challenges that arise from the application areas of programmable matter. For this purpose, we brought together a combination of established experts from DNA computing, swarm robotics, computational geometry, and distributed computing. On the senior level, participants included a number of leading authorities who are established in more than one of the mentioned topics; on the junior level, we had a good selection of highly talented scientists who are able to advance the field by specific contributions.

The seminar started with a plenary introduction of all participants, their research areas and their specific challenges and expectations for the seminar. This was followed by a number

of plenary sessions, in which experts gave overviews of broad developments and specific open problems.

- Erik Demaine gave an overview of challenges for **geometric algorithms** in the settings of reconfigurable robots (both modular and folding robots that can become any possible shape), robot swarms (which may be so small and simple that they have no identity), and self-assembly (building computers and replicators out of DNA tiles).
- Dave Doty and Chris Thachuk gave a survey of the basics of experimental and theoretical **DNA tile self-assembly**, concluding with suggestions for theoretical problems related to programmable control of the nucleation of assemblies. A second part consisted of a survey of DNA strand displacement, including the problem of orienting molecules on a surface with the use of DNA origami and some clever shapes that can “align” themselves into target placements.
- Andréa Richa presented an overview of **self-organizing particle systems**, describing programmable matter as an abstract collection of simple computational elements (particles) with limited memory that each execute fully distributed, local, asynchronous algorithms to self-organize and solve system-wide problems such as movement, (re)configuration, and coordination.
- Aaron Becker discussed the connection between **robot swarms and programmable matter**, in particular in a setting with a global input to a whole particle swarm, as well as open questions arising from the use of mobile robots to fold 2D planar stock into 3D bricks and to connect the bricks together.

Spread throughout the week, further presentations were given by Spring Berman (applications and open challenges in swarm robotics and a control-theoretic framework for robotic swarms and programmable matter), Julien Bourgeois (realizing programmable matter with modular robots), Luca Cardelli (sequence-

able DNA algorithms), Kenneth Cheung (programmable modular periodic metamaterials), Sándor Fekete (coordinated motion planning), Roderich Groß (capabilities of individual units in distributed robotic systems and making programmable matter self-propel efficiently), Dan Halperin (hard vs. easy tasks in multi-robot motion planning), Heiko Hamann (self-assembly and collective construction based on minimal surprise), Lila Kari (DNA smart-tile self-assembly and computational CRISPR), MinJun Kim (engineering particles for robot swarms and modular microrobotics), Alcherio Martinoli (fluid-mediated stochastic self-assembly), Friedhelm Meyer auf der Heide (continuous strategies for swarm robotics), Nils Napp (autonomous construction in unstructured environments), Pekka Orponen (algorithmic design of RNA nanostructures) and Christian Scheideler (a survey on hybrid programmable matter).

A key feature of the seminar was exceptionally intensive, interdisciplinary collaboration throughout the week, based on the use of the new interactive electronic tool `coauthor`. This tool³⁸, specifically developed for use in a workshop-like environment, is an excellent platform that provides a versatile medium for collaborative research discussions, and maintains easily accessible structured records for future reference. We have found that `coauthor` greatly facilitated the work done during the seminar, enabling not just identification of, but also dynamic research work on a number of new topics. These include (A) specific problems in the context of hybrid models for programmable matter, in which there is a set of active micro-robots that can move a large set of simple material tiles that cannot move themselves; (B) aspects of distributed boundary detection for self-organizing swarms; (C) fundamental issues related to the computational equivalence of completely different self-assembly systems and robotic models; and (D) questions of self-aligning geometric shapes that would allow more robust methods for DNA origami and self-assembly. For some aspects, we were able to resolve long-standing open problems; for others, we made significant progress that will undoubtedly lead to future publications. As a consequence, the seminar has triggered a number of new collaborations and a variety of followup projects that will undoubtedly contribute to further collaborative research activities.

³⁸ <https://github.com/edemaine/coauthor/>

6.42 Blockchain Technology for Collaborative Information Systems

Organizers: Marlon Dumas, Richard Hull, Jan Mendling, and Ingo Weber

Seminar No. 18332

Date: August 12–17, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.67

© Creative Commons BY 3.0 Unported license

© Marlon Dumas, Richard Hull, Jan Mendling, and Ingo Weber



Participants: Michael Coblenz, Søren Debois, Claudio Di Ciccio, Alevtina Dubovitskaya, Marlon Dumas, Fabiana Fournier, Avigdor Gal, Luciano García-Bañuelos, Stephan Haarmann, Richard Hull, Hans-Arno Jacobsen, Mieke Jans, Agnes Koschmider, Qinghua Lu, Raimundas Matulevičius, Jan Mendling, Petr Novotny, Sooyong Park, Stefanie Rinderle-Ma, Stefan Schulte, Jerome Simeon, Ludwig Stage, Mark Staples, Barbara Weber, Ingo Weber, Francesca Zerbato, Kaiwen Zhang

Blockchain technology enables an evolving set of parties to maintain a safe, permanent, and tamper-proof ledger of transactions without a central authority. This technology opens manifold opportunities to redesign Business-to-Business (B2B) collaborations in a wide range of fields, including supply chain, logistics, service agreements, healthcare, and Industry 4.0. Importantly, it can enable substantial efficiency gains in terms of cost and time it takes to set-up and perform collaborative processes, particularly in settings where there is a lack of trust between the parties involved in the collaboration. Traditionally, collaborative processes are executed by relying on trusted third-party providers such as Electronic Data Interchange (EDI) hubs or escrows. This centralized architecture creates entry barriers and hinders bottom-up innovation. Blockchains and smart contracts enable these processes to be executed in a distributed manner without delegating trust to central authorities nor requiring mutual trust between each pair of parties. Further, blockchain enables fine-grained access control, thus allowing multiple parties to selectively share their data with each other and to selectively grant permissions to perform transactions on these data.

While blockchain opens up new possibilities, it also raises a number of challenges because it requires us to re-think the way B2B collaborations are designed and implemented. In contrast to centralized collaborative processes, the transparent and decentralized nature of blockchains brings in new challenges related to compliance, control, and privacy, in addition to major scalability and performance challenges. This seminar brought together established and young researchers with forward-thinking industry representatives from both large and start-up companies, in order to establish a research roadmap for blockchain-based collaborative information systems, and to initiate concrete research collaborations between participants along this roadmap.

6.43 Formalization of Mathematics in Type Theory

Organizers: Andrej Bauer, Martín H. Escardó, Peter L. Lumsdaine, and Assia Mahboubi
Seminar No. 18341

Date: August 19–24, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.130

© Creative Commons BY 3.0 Unported license

© Andrej Bauer, Martín H. Escardó, Peter L. Lumsdaine, and Assia Mahboubi

Participants: Benedikt Ahrens, Carlo Angiuli, Andrej Bauer, Sophie Bernard, Yves Bertot, Auke Booij, Guillaume Brunerie, Jacques Carette, Mario Carneiro, Cyril Cohen, Manuel Eberl, Martín H. Escardó, Diane Gallois-Wong, Gaëtan Gilbert, Georges Gonthier, Daniel R. Grayson, Philipp Haselwarter, Florent Hivert, Johannes Hölzl, Kuen-Bang (Favonia) Hou, Simon Huber, Fabian Immler, Nicolai Kraus, Dan Licata, Peter L. Lumsdaine, Assia Mahboubi, Maria Emilia Maietti, Anders Mörtberg, Scott Morrison, Russell O’Connor, Ian Orton, Anja Petkovic, Claudio Sacerdoti Coen, Bas Spitters, Jonathan Sterling, Neil Strickland, Michael Trott, Hoang Le Truong, Josef Urban, Floris van Doorn, Makarius Wenzel, Bohua Zhan



We and all the participants were delighted to benefit from Dagstuhl’s inspiring environment.

Proof assistants are receiving increased attention from users with a background in mathematics, as opposed to their traditional users from theoretical computer science/logic/program verification, and this was the major focus of the meeting. This is true in particular of proof assistants based on dependent types, probably due in part to the advent of homotopy type theory, developed in the proof assistants Coq, Agda and Lean.

The audience of the seminar was thus rather unusual in composition, and featured several experienced researchers used to attending seminars at the Mathematisches Forschungsinstitut Oberwolfach, and visiting Schloss Dagstuhl for the first time. In order to foster discussion and fuse collaborations, we adopted a different format from the standard string of slide-based talks: talks in the morning, so that people get to know the work of each other, and working in groups in the afternoon. At the end of each day, before dinner, each group presented a summary of the outcomes of their meetings to all participants, which allowed inter-group discussion and collaboration. This had been tried before by some of the organizers, in the course of Dagstuhl seminar 16112, and worked just as well in our case.

Working group topics were proposed by the audience on the first day, by giving short presentations of a few minutes and writing topics in the board. Some were quite specialized and homogeneous (e.g. the cubical type theory group), and allowed people to have a focussed collaborative brainstorming on a specific open problem of the field. Some were more open-ended, and allowed people to confront various approaches to the same issue/concept in different systems (different proof assistants, computer algebra systems, etc.).

Some people did applied work, such as trying to compute the so-called Brunerie number from an existing proof in homotopy type theory, in order to identify and fix inefficiency problems in

proofs assistants based on cubical type theory. Some people used their spare time to solve the “Dagstuhl dinner” problem. Details of the topics discussed are in the reports produced by each group.

This was a rather productive meeting, and people from different scientific backgrounds not only met but talked together effectively, solving and identifying problems to work on collaboratively in future.

6.44 Modeling for Sustainability

Organizers: Gordon Blair, Betty H. C. Cheng, Lorenz Hilty, and Richard F. Paige
Seminar No. 18351

Date: August 26–31, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.8.146

© Creative Commons BY 3.0 Unported license

© Gordon Blair, Betty H. C. Cheng, Lorenz Hilty, and Richard F. Paige



Participants: Olivier Barais, Lucy Bastin, Christoph Becker, Didier Beloin-Saint-Pierre, Nelly Bencomo, Stefanie Betz, Keith Beven, Gordon Blair, Gael Blondelle, Betty H. C. Cheng, Ruzanna Chitchyan, Benoît Combemale, Nelly Condori-Fernandez, Letícia Duboc, François Fouquet, Joao Goncalves, Øystein Haugen, Lorenz Hilty, Jean-Marc Jézéquel, Eva Kern, Jörg Kienzle, Sedef Akinli Kocak, Peter D. Mosses, Gunter Mussbacher, Oscar M. Nierstrasz, Richard F. Paige, Birgit Penzenstadler, Bernhard Rumpe, Lionel Seinturier, Noelle Selin, Norbert Seyff, Eugene Syriani, Colin Venters, Jon Whittle, Paul Young

Many different kinds of models, from engineering models to scientific models, have to be integrated and coordinated to support sustainability systems such as smart grid or cities, i.e., dynamically adaptable resource management systems that aim to improve the techno-economic, social, and environmental dimensions of sustainability. Scientific models help understand sustainability concerns and evaluate alternatives, while engineering models support the development of sustainability systems. As the complexity of these systems increases, many challenges are posed to the computing disciplines to make data and model-based analysis results more accessible as well as integrate scientific and engineering models while balancing trade-offs among varied stakeholders. This seminar explored the intrinsic nature of both scientific and engineering models, the underlying differences in their respective foundations, and the challenges related to their integration, evolution, analysis, and simulation including the exploration of what-if scenarios.

Sustainability systems must provide facilities for the curation and monitoring of data sets and models and enable flexible (open) data and model integration, e.g., physical laws, scientific models, regulations and preferences, possibly coming from different technological foundations, abstractions, scale, technological spaces, and world views. This also includes the continuous, automated acquisition and analysis of new data sets, as well as automated export of data sets, scenarios, and decisions. The main function is to support the generation of what-if scenarios to project the effects on the different sustainability dimensions, and support the evaluation of externalities, especially for non rapidly renewable resources. Since the predictions are necessarily probabilistic, the system must be able to assess the uncertainty inherent in all its actions and provide suitable representations of uncertainty understandable by users. In addition to generating what-if scenarios to explore alternate model instantiations, the tool should be capable of generating suggestions for how to reach user-specified goals including quantifiable impacts and driving the dynamic adaptation of sustainability systems. These powerful

services must be made accessible to the population at large, regardless of their individual situation, social status, and level of education.

This seminar explored how Model-Driven Engineering (MDE) will help to develop such an approach, and in particular i) how modeling frameworks would support the integration of the various heterogeneous models, including both engineering and scientific models; ii) how domain specific languages (DSLs) would (a) support the required socio-technical coordination, i.e., engage engineers, scientists, decision makers, communities, and the general public; and (b) integrate analysis/probabilistic/user models into the control loop of smart CPS (cyber physical system). DSLs are also supposed to provide the right interface (in terms of abstractions/constructs) to be used as tools for discovering problems and evaluating ideas.

The seminar served to identify critical disciplines and stakeholders to address MDE for sustainability and the research roadmap of the MDE community with regards to the development of sustainability systems. In particular, the seminar identified and explored four key areas: 1) research challenges relevant to modeling for sustainability (M4S); 2) a multidisciplinary collection of relevant literature to provide the foundation for exploring the research challenges; 3) three case studies from different application domains that provide a vehicle for illustrating the M4S challenges and for validating relevant research techniques; and 4) the human and social aspects of M4S.

The cumulative results of the work performed at the seminar and subsequent collaborations will help to establish the required foundations for integrating engineering and scientific models, and to explore the required management facilities for evaluating what-if scenarios and driving adaptive systems. In addition, we envision to produce as an outcome of the seminar a representative case study that will be used by the community to assess and validate contributions in the field of modeling for sustainability.

6.45 Measuring the Complexity of Computational Content: From Combinatorial Problems to Analysis

Organizers: Vasco Brattka, Damir D. Dzhafarov, Alberto Marcone, and Arno Pauly
Seminar No. 18361

Date: September 2–7, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.9.1

© Creative Commons BY 3.0 Unported license

© Vasco Brattka, Damir D. Dzhafarov, Alberto Marcone, and Arno Pauly

Participants: Eric P. Astor, Vittorio Bard, Laurent Bienvenu, Vasco Brattka, Merlin Carl, Raphael Carroy, Matthew Colbrook, Matthew de Brecht, Hannes Diener, Francois G. Dorais, Damir D. Dzhafarov, Marta Fiori Carones, Guido Gherardi, Jun Le Goh, Peter Hertling, Denis R. Hirschfeldt, Jeffrey L. Hirst, Rupert Hölzl, Mathieu Hoyrup, Akitoshi Kawamura, Takayuki Kihara, Ulrich Kohlenbach, Wei Li, Alberto Marcone, Kenshi Miyabe, Carl Mummert, Takako Nemoto, Eike Neumann, Keng Meng Ng, Sabrina Ouazzani, Ludovic Patey, Arno Pauly, Cristóbal Rojas, Matthias Schröder, Victor Selivanov, Paul Shafer, Giovanni Soldà, Florian Steinberg, Patrick Uftring, Manlio Valenti, Klaus Weihrauch, Linda Brown Westrick, Keita Yokoyama



Reducibilities such as many-one, Turing or polynomial-time reducibility have been an extraordinarily important tool in theoretical computer science from its very beginning. In recent years these reducibilities have been transferred to the continuous setting, where they allow us to classify computational problems on real numbers and other continuous data types.

In the late 1980s Weihrauch has introduced a reducibility that can be seen as an analogue of many-one reducibility for (multi-valued) functions on infinite data types. This reducibility, now called *Weihrauch reducibility*, was studied since the 1990s by Weihrauch's school of computable analysis and flourished recently when Gherardi and Marcone proposed this reducibility as a tool for a uniform approach to reverse analysis.

Reverse mathematics aims to classify theorems according to the axioms that are needed to prove these theorems in second-order arithmetic. This proof theoretic approach yields non-uniform classifications of the computational content of certain theorems. However, many of these classifications also have uniform content and Weihrauch complexity allows us to study this uniform computational content directly using methods of computability theory.

This perspective has motivated Dorais, Dzhafarov, Hirst, Mileti and Shafer, on the one hand, Hirschfeldt and Jockusch, on the other hand, to study combinatorial problems using this approach. This research has led to a number of further reducibilities (computable reducibility, generalized Weihrauch reducibility and others) that can be seen as non-uniform or less resource sensitive versions of Weihrauch reducibility. Using this toolbox of reducibilities one can now adjust the instruments exactly according to the degree of uniformity and resource sensitivity that one wants to capture.

A precursor seminar³⁹ that was also held at Dagstuhl has

been instrumental in bringing together researchers from these different communities for the first time. This has created a common forum and fostered several research developments in this field. We believe that the current seminar was very successful in strengthening and deepening the collaborations between the involved communities. Ample time was left and successfully used for research in groups. A novelty of the current seminar was a special session at which solutions of open problems from the previous seminar were presented. To see that several of the major open problems of the previous meetings were solved in the meantime was inspiring and motivating! Some of the solutions involve new techniques with a wider applicability. Hopefully, we will see solutions to some of the open questions presented at the current seminar in the not too far future! Altogether, the seminar did proceed in a highly productive atmosphere, thanks to many excellent contributions from participants. Inspired by these contributions the organizers are planning to edit a special issue of the journal *Computability* dedicated to this seminar.

This report includes abstracts of many talks that were presented during the seminar, it includes a list of some of the open problems that were discussed, as well as a bibliography on Weihrauch complexity that was started during the previous meeting and that saw significant growth in the meantime. Altogether, this report reflects the extraordinary success of our seminar and we would like to use this opportunity to thank all participants for their valuable contributions and the Dagstuhl staff for their excellent support!

³⁹ 15392 Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis, see <https://doi.org/10.4230/DagRep.5.9.77>

6.46 Knowledge Graphs: New Directions for Knowledge

Representation on the Semantic Web

Organizers: Piero Andrea Bonatti, Stefan Decker, Axel Polleres, and Valentina Presutti

Seminar No. 18371

Date: September 9–14, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.9.29

© Creative Commons BY 3.0 Unported license

© Piero Andrea Bonatti, Michael Cochez, Stefan Decker, Axel Polleres, and Valentina Presutti



Participants: Wouter Beek, Christian Bizer, Eva Blomqvist, Piero Andrea Bonatti, Daniel Brickley, Sarven Capadisi, Michael Cochez, Claudia d'Amato, Gerard de Melo, Stefan Decker, Michel Dumontier, Paul Groth, Claudio Gutierrez, Andreas Harth, Aidan Hogan, Sabrina Kirrane, Markus Krötzsch, Barend Mons, Roberto Navigli, Sebastian Neumaier, Axel-Cyrille Ngonga-Ngomo, Andrea Giovanni Nuzzolese, Heiko Paulheim, Lydia Pintscher, Axel Polleres, Valentina Presutti, Sabbir Rashid, Sebastian Rudolph, Marta Sabou, Juan F. Sequeda, Cogan Matthew Shimizu, Elena Simperl, Dezhao Song, Steffen Staab, Volker Tresp, Marieke van Erp, Frank van Harmelen, Maria-Esther Vidal, Michael Witbrock, Sonja Zillner, Antoine Zimmermann

In 2001 Berners-Lee et al. stated that “The Semantic Web is an extension of the current web in which information is given well-defined meaning, better enabling computers and people to work in cooperation.”

The time since the publication of the paper and creation of the foundations for the Semantic Web can be roughly divided in three phases: The first phase focused on bringing Knowledge Representation to Web Standards, e.g., with the development of OWL. The second phase focused on data management, linked data and potential applications. In the third, more recent phase, with the emergence of real world applications and the Web emerging into devices and things, emphasis is put again on the notion of Knowledge, while maintaining the large graph aspect: Knowledge Graphs have numerous applications like semantic search based on entities and relations, disambiguation of natural language, deep reasoning (e.g. IBM Watson), machine reading (e.g. text summarisation), entity consolidation for Big Data, and text analytics. Others are exploring the application of Knowledge Graphs in industrial and scientific applications.

The shared characteristic by all these applications can be expressed as a challenge: the capability of combining diverse (e.g. symbolic and statistical) reasoning methods and knowledge representations while guaranteeing the required scalability, according to the reasoning task at hand. Methods include: Temporal knowledge and reasoning, Integrity constraints, Reasoning about contextual information and provenance, Probabilistic and fuzzy reasoning, Analogical reasoning, Reasoning with Prototypes and Defeasible Reasoning, Cognitive Frames, Ontology Design Patterns (ODP), and Neural Networks and other machine learning models.

With this Dagstuhl Seminar, we intend to bring together researchers that have faced and addressed the challenge of combining diverse reasoning methods and knowledge representations in

different domains and for different tasks with Knowledge Graphs and Linked Data experts with the purpose of drawing a sound research roadmap towards defining scalable Knowledge Representation and Reasoning principles within a unifying Knowledge Graph framework. Driving questions include:

- What are fundamental Knowledge Representation and Reasoning methods for Knowledge Graphs?
- How should the various Knowledge Representation, logical symbolic reasoning, as well as statistical inference methods be combined and how should they interact?
- What are the roles of ontologies for Knowledge Graphs?
- How can existing data be ingested into a Knowledge Graph?

In order to answer these questions, the present seminar was aiming at cross-fertilisation between research on different Knowledge Representation mechanisms, and also to help to identify the requirements for Knowledge Representation research originating from the deployment of Knowledge graphs and the discovery of new research problems motivated by applications. We foresee, from the results summarised in the present report, the establishment of a new research direction, which focuses on how to combine the results from knowledge representation research in several subfields for joint use for Knowledge Graphs and Data on the Web.

■ The Seminar

The idea of this seminar emerged when the organisers got together discussing about writing a grant proposal. They all shared, although from different perspectives, the conviction that research on Semantic Web (and its scientific community) reached a critical point: it urged a paradigm shift. After almost two decades of research, the Semantic Web community established a strong identity and achieved important results. Nevertheless,

the technologies resulting from its effort on the one hand have proven the potential of the Semantic web vision, but on the other hand became an impediment; a limiting constraint towards the next major breakthrough. In particular, Semantic Web knowledge representation models are insufficient to face many important challenges such as supporting artificial intelligence systems in showing advanced reasoning capabilities and socially-sound behaviour at scale. The organisers soon realised that a project proposal was not the ideal tool for addressing this problem, which instead needed a confrontation of the Semantic Web scientific community with other relevant actors, in the field. From this discussion, the “knowledge graph” concept emerged as a key unifying ingredient for this new form of knowledge representation – embracing both the Semantic Web, but also other adjacent communities – and it was agreed that a Dagstuhl seminar on “Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web” was a perfect means for the purpose.

The list of invitees to the seminar included scientists from both academia and industry working on knowledge graphs, linked data, knowledge representation, machine learning, automated reasoning, natural language processing, data management, and other relevant areas. Forty people have participated in the seminar, which was very productive. The active discussions during plenary and break out sessions confirmed the complex nature of the proposed challenge. This report is a fair representative of the variety and complexity of the addressed topics.

The method used for organising the seminar deserves further elaboration. The seminar had a five-day agenda. Half of the morning on the first day was devoted to ten short talks (5 minutes each) given by a selection of attendees. The speakers were identified by the organisers as representatives of complimentary topics based on the result of a Survey conducted before the seminar: more than half of the invitees filled a questionnaire that gave them the opportunity to briefly express their perspectives on the topic and to point out relevant challenges that they would put in their future research agenda with the highest priority.



Fig. 6.12
Blackboard with post-its from the open session.

The aim of these short speeches was to ignite the confrontation by sharing the emerging views on the main challenges from

this survey. After the speeches we organised the further discussion in an “Open Space” session that served to collaboratively build the agenda for the rest of the day (and that influenced the agenda of the next days). The open space method consists of giving everyone the opportunity to propose one or more break out topics. To propose a topic, a proposer had to explain in few words what it was about, then write it down on a post-it that was attached on a blackboard (see Figure 6.12). At the end of the session, attendees were invited to sign up for the topics of their interest (by marking the corresponding post-it).

The more popular ones (up to fifteen and having at least three sign ups) were selected to compose the agenda. Each break out session used a one-hour slot during the afternoon. The second day continued with most of the break out sessions with the aim of continuing the discussion started the first day and work towards consolidating a report (finalised on the fourth day). Reports would reflect view and vision emerging from the break out group. On the same day attendees had the opportunity to self-propose to give additional short speeches, addressing missing relevant topics. We used part of the second day’s morning for these speeches. We explicitly asked attendees to avoid speeches on “my research” and to only address relevant challenges that were overlooked so far.

On the third day we started with a plenary discussion and the seminar group agreed on splitting into four groups to discuss “Grand challenges” separately, then share the results before going back to the break out sessions. The aim was to share a common high level vision reference before consolidating the more specific discussions that were ongoing in the break out sessions. On the fourth day, the seminar group split again in break out sessions including a “Grand challenges” one. Each session was assigned to at least two coordinators, who committed to consolidate in a draft report the results from the previous meetings. It was decided to merge a few topics, when appropriate.

Break out sessions had varied level of technical abstraction depending on the nature of the topic, and its level of maturity within the state of the art. To give some examples: the break out session about “Grand challenges” mainly discussed a vision for a future research agenda and maintained a high level of abstraction, while the session on “Human and Social Factors in Knowledge Graphs” provided more concrete insights as it could build on both academic and industrial research results, projects and practical experiences. The session on “Applications of Knowledge Graphs” focused on technical details and issues on two relevant sample applications.

■ Overview of the Report

This report is organised in two main parts: Section 3 includes a list of abstracts providing an overview of the short speeches that we had the first two days. All the other sections are consolidated reports of the emerging vision, research challenges, possible research agenda, and proposed approaches, from break out sessions. When applicable, the reports give an overview of specific relevant research work.

Acknowledgements The editors are grateful to Michael Cochez for his valuable contributions to the planning and organization of this workshop. He was in all respects the fifth organizer.

6.47 Quantum Programming Languages

Organizers: Michele Mosca, Martin Roetteler, and Peter Selinger
Seminar No. 18381

Date: September 16–21, 2018 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.8.9.112
 © Creative Commons BY 3.0 Unported license
 © Michele Mosca, Martin Roetteler, and Peter Selinger



Participants: Matthew Amy, Sébastien Bardin, Xiaoning Bian, Earl Campbell, Andrew Cross, Olivia Di Matteo, Austin G. Fowler, Frank Fu, Vlad Gheorghiu, Sabine Glesner, Robert Glück, Christopher Granade, Markus Grassl, Thomas Häner, Shih-Han Hung, Nader Khammassi, Vadym Kliuchnikov, Sriram Krishnamoorthy, Andrew John Landahl, Albertus Johannes Lindenhovius, Michael W. Mislove, Michele Mosca, Beatrice Nash, Jennifer Paykin, Robert Rand, Mathys Rennela, Francisco Rios, Martin Roetteler, Neil Julien Ross, Bruno Schmitt Antunes, Peter Selinger, Mathias Soeken, Damian Steiger, Rainer Steinwandt, Benoit Valiron, Rodney Van Meter, Michael Walter, Robert Wille, Shigeru Yamashita, Mingsheng Ying, Vladimir Zamdzhiev, Margherita Zorzi, Alwin Zulehner, Paolo Zuliani

This report documents the program and the outcomes of Dagstuhl Seminar 18381 “Quantum Programming Languages”.

The aim of the seminar was to bring together researchers from quantum computing – in particular those focusing on quantum algorithms and quantum error correction – and classical programming languages. Open questions that were of interest to this group include new methods for circuit synthesis and optimization, compiler optimizations and rewriting, embedded languages versus non-embedded languages, implementations of type systems and error reporting for quantum languages, techniques for verifying the correctness of quantum programs, and new techniques for compiling efficient circuits and protocols for fault-tolerant questions and their 2D layout.

Quantum computing is getting real. Several laboratories around the world are implementing hardware platforms. For instance, systems based on superconducting qubits, such as those at IBM, Google, Intel, the University of Maryland, ionQ, and Rigetti are now scaling into the 50-150 qubit range.

While research on the theoretical side of the field addressed fundamental questions such as how to best leverage this new model of computation for algorithmic applications, a topic that has received significantly less attention is how to actually program quantum computers. To take advantage of the immense computing power offered by quantum computers as they come online in the coming years, software tools will be essential. We want these tools to be available, efficient and reliable, so that we can quickly and reliably reap the positive benefits that quantum computers have to offer.

It is clear that quantum programming will require tools for automatically generating large-scale circuits and for synthesizing circuits from elementary fault-tolerant gates which then can be carried out by a future quantum computer. However, it is less clear what the best way will be to go about these challenging

issues. Questions that were discussed at the seminar include the following:

- How can we program a quantum computer? What are the basic structures that a language should support and how can a compiler help a user develop abstract/high-level reasoning about algorithms?
- How do we model the underlying instruction set? As currently the underlying hardware is quickly evolving, how can we best model a fault-tolerant quantum computer?
- How to compile and optimize quantum programs? Automatic translation of high-level programs into circuits will be key to program quantum computers. How to design good tools for this?
- How to we test and verify quantum programs? Given that it is hard for classical computers to simulate the time evolution of a quantum computer, how can we ascertain correctness of a circuit?

The seminar brought together some 44 researchers with diverse skill sets from quantum computing, mathematical foundations of programming languages, implementation of programming languages, and formal verification. The seminar consisted of 23 talks, as well as a number of vibrant discussion sessions and a software demonstration session. The sessions where:

- Wine Cellar discussion, moderated by Sabine Glesner. This was our first discussion session. We discussed the questions raised by Sabine Glesner during her talk: Why do we need quantum programming languages? Which “killer applications” would make quantum programming languages successful? What are appropriate abstractions from quantum hardware? What are theoretical models for quantum computing?
- Discussion session on Debugging, moderated by Rodney Van Meter. This session focused on what are appropriate debug-

ging techniques for quantum computing. The issue arises because the most common classical debugging technique, setting break points and examining the program state, cannot be applied in the context of quantum computing.

- Discussion session on Challenge Problems for Quantum Computing, moderated by Earl Campbell. In this session, we discussed coming up with well-defined problems with some success quantifier for quantum computation, similar to the successful SAT competitions.
- Group survey session on a Bird's Eye View on Quantum Languages, moderated by Robert Rand. In this session, the group compiled a list of all quantum programming languages and toolkits we are currently aware of, and classified them according to various criteria, for example, whether the languages are imperative or functional, whether the computational paradigm is circuit generation or Knill's QRAM model, whether the language is high-level or assembly, whether it supports type-safety and/or verification, etc.
- Group survey session on Tools for Quantum Optimization, moderated by Matthew Amy. In this session, the group compiled a list of available tools for optimization of quantum circuits.
- Group discussion on Opportunities for Education and Outreach, moderated by Rodney Van Meter. The discussion centered on new opportunities for public outreach and education that are enabled by the emergence of new quantum tools.
- Software demonstration session, moderated by Martin Roeteler. In this session, 10 researchers gave rapid demonstrations, of a about 10 minutes each, of various software tools they have designed.

Most of the participants rated the seminar as a success. We managed to connect researchers from different communities, and engaged in a vibrant exchange of novel ideas, and started to tackle important problems such as the analysis of quantum algorithms for real-world computational problems, compiler optimizations, reversible computing, and fault-tolerant quantum computing.

6.48 Algebraic Methods in Computational Complexity

Organizers: Markus Bläser, Valentine Kabanets, Jacobo Torán, and Christopher Umans

Seminar No. 18391

Date: September 23–28, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.9.133

© Creative Commons BY 3.0 Unported license

© Markus Bläser, Valentine Kabanets, Jacobo Torán, and Christopher Umans



Participants: Farid Ablayev, Eric Allender, Josh Alman, Vikraman Arvind, Nikhil Balaji, Markus Bläser, Andrej Bogdanov, Sourav Chakraborty, Stephen A. Fenner, Michael A. Forbes, Lance Fortnow, Anna Gál, William Gasarch, Sevag Gharibian, Frederic Green, Rohit Gurjar, William Hoza, Christian Ikenmeyer, Valentine Kabanets, Neeraj Kayal, Pascal Koiran, Antonina Kolokolova, Arpita Korwar, Michal Koucký, Nutan Limaye, Zhenjian Lu, Vladimir Lysikov, Meena Mahajan, David A. Mix Barrington, Anurag Pandey, Natacha Portier, Noga Ron-Zewi, Chandan Saha, Ramprasad Satharishi, Nitin Saxena, Uwe Schöning, Ronen Shaltiel, Amnon Ta-Shma, Thomas Thierauf, Jacobo Torán, Christopher Umans, Nikolay K. Vereshchagin

The seminar brought together more than 40 researchers covering a wide spectrum of complexity theory. The focus on algebraic methods showed the great importance of such techniques for theoretical computer science. We had 24 talks, most of them lasting about 45 minutes, leaving ample room for discussions. We also had a much appreciated rump session on Tuesday evening in which Antonina Kolokolova, Bill Gasarch, Lance Fortnow, Chandan Saha, William Hoza, Neeraj Kajal and Arpita Korwar presented some open questions. In the following we describe the major topics of discussion in more detail.

Circuit Complexity This is an area of fundamental importance to Complexity. Circuits studied from many different perspectives were one of the main topics in the seminar. *Eric Allender* gave an overview of the Minimum Circuit Size Problem (MCSP): given the truth-table for a Boolean function, what is the size of the minimum circuit computing it? In his talk he mentioned some interesting results proving that some low complexity classes cannot be reduced to the problem of computing superlinear approximations to circuit size.

Arithmetic circuits and formulas are a special computation model that uses $+$ and \times as operators for computing polynomials instead of Boolean operations. *Nutan Limaye* presented a depth hierarchy theorem for this model showing that there is a polynomial computed by a depth $D + 1$ polynomial sized multilinear formula such that any depth D multilinear formula computing the polynomial must have exponential size.

Chandan Saha considered a further restriction to depth three circuits C computing a polynomial $f = T_1 + T_2 + \dots + T_s$, where each T_i is a product of d linear forms in n variables. He presented a randomized algorithm to reconstruct non-degenerate homogeneous depth three circuits, for the case $n > (3d)^2$, given black-box access to f . The algorithm works in polynomial time in n , s and d .

Depth-2 circuits with polynomial size and linear threshold functions were presented by *Meena Mahajan*. She surveyed the landscape below these circuits and present one new result concerning decision lists.

Algebraic Complexity There were also several presentations discussing the complexity of several problems over algebraic structures.

Nitin Saxena considered in his talk the problem of testing whether a set F of polynomials given as algebraic circuits has an algebraic dependence. He showed that this problem can be computed in $\text{AM} \cap \text{coAM}$ thus solving an open question from 2007.

Problems related to the minimum code-word problem and the existence of non-trivial automorphism moving few vertices in graphs or hypergraphs, were presented by *V. Arvind* in his talk. He discuss the parameterized complexity of this and related algebraic problems.

Josh Alman gave an interesting talk on Matrix Multiplication (MM). He surveyed the two main approaches for MM algorithms: the Laser method of Strassen, and the Group theoretic approach of Cohn and Umans and defined a generalization which subsumes these two approaches. He then explained ways to obtain lower bounds for algorithms for MM when using these algorithmic methods.

Rohit Gurjar studied the class of matrices A for which the lattice $L(A)$ formed by all integral vectors v in the null-space of A , has only polynomially many near-shortest vectors. He proved that this is the case when the matrix A is totally unimodular (all sub-determinants are 0, $+1$, or -1). As a consequence he could show a deterministic algorithm for PIT for any polynomial of the form $\det(\sum x_i A_i)$ for rank-1 matrices A_i .

Pseudo-Randomness and Derandomization

Derandomization is an area where there are tight connections between lower bounds and algorithms. Strong enough circuit lower bounds can be used to construct pseudo-random generators that can then be used to deterministically simulate randomized algorithms. A central question in derandomization is whether randomized logspace RL equals deterministic logspace L. To show that $RL = L$, it suffices to construct explicit pseudorandom generators that fool polynomial-size read-once (oblivious) branching programs (roBPs). There were two talks related to this question. *Michael Forbes* presented a method to obtain an explicit PRG with seed-length $O(\log^3 n)$ for polynomial-size roBPs reading their bits in an unknown order. *William Hoza* gave an explicit hitting set generator for read-once branching programs with known variable order. As a corollary of this construction, it follows that every RL algorithm that uses r random bits can be simulated by an NL algorithm that uses only $O(r/\log^c n)$ nondeterministic bits, where c is an arbitrarily large constant. Another consequence of the result is that any RL algorithm with small success probability ϵ can be simulated deterministically in space $O(\log^{3/2} n + \log n \log \log(1/\epsilon))$.

A hitting set is a set of instances such that every non-zero polynomial in the model has a non-root in the set. This would solve the Polynomial Identity Testing problem (PIT) in that model. *Ramprasad Saptharishi* showed that by barely improving the trivial $(s + 1)^n$ size hitting set even for n -variate degree s , size s algebraic circuits, we could get an almost complete derandomization of PIT.

In a second talk, *William Hoza* talked about the possibility of derandomizing an algorithm by using randomness from the input itself. For a language L with a bounded-error randomized algorithm in space S and time $n \cdot \text{poly}(S)$ he gave a randomized algorithm for L with the same time and space resources but using only $O(S)$ random bits; the algorithm has a low failure probability on all but a negligible fraction of inputs of each length.

Andrej Bogdanov considered the problem of extracting true randomness from a set biased dice (Santha-Vazirani sources). He presented a recent result in which he completely classified all non-trivial randomness sources of this type into: non-extractable ones, extractable from polynomially many samples, and extractable from an logarithmically many samples (in the inverse of the error).

Coding Theory Error-correcting codes and other kinds of codes, particularly those constructed from polynomials, i.e. Reed-Solomon codes or Reed-Muller codes, lie at the heart of many significant results in Computational Complexity. This is an area in which the relation between different areas of complexity, like the analysis of algebraic structures or derandomization becomes especially fruitful.

Greatly improving previously known constructions for an odd size alphabet, *Michal Koucký* presented a construction of quasi-Gray codes of dimension n and length 3^n over the ternary alphabet $\{0, 1, 2\}$ with worst-case read complexity $O(\log n)$ and write complexity 2. This generalizes to arbitrary odd-size alphabets. These results were obtained via a novel application of algebraic tools together with the principles of catalytic computation.

Noga Ron-Zewi presented a very recent result showing that Folded Reed-Solomon codes achieve list decoding capacity with constant list sizes, independent of the block length. She explained that multiplicity codes exhibit similar behavior, and used this to obtain capacity achieving locally list decodable codes with query complexity significantly lower than previous constructions.

Binary error correcting code with relative distance $(1 - \epsilon)/2$ and relative rate $\epsilon^{2+o(1)}$ were explained in one of the talks given

by *Amnon Ta-Shma*. Previous explicit constructions had rate about ϵ^3 . The main tool used for this construction are *Parity Samplers*. He explained how to get better explicit parity samplers using a variant of the zig-zag product.

In his second talk, Amnon talked about $(1 - \tau, L)$ erasure list-decodable codes. He presented a recent work where he constructed for the first time an explicit binary $(1 - \tau, L)$ erasure list-decodable code having rate $\tau^{1+\gamma}$ (for any constant $\gamma > 0$ and τ small enough) and list-size $\text{poly}(\log 1/\tau)$, exhibiting an explicit non-linear code that provably beats the best possible linear one. The main ingredient in his construction is a new (and almost-optimal) *unbalanced* two-source extractor.

Quantum Complexity Complexity issues arising in the context of quantum computation are an important area in Complexity Theory since several decades. In this workshop we had one talk on this topic. *Sevag Gharibian* talked about quantum versions of the classical k -SAT problem. He talked about the problem of computing satisfying assignments to k -QSAT instances which have a “matching” or “dimer covering”; this is an NP problem whose decision variant is trivial, but whose search complexity remains open. He presented a parameterized algorithm for k -QSAT instances from a non-trivial class, which allows to obtain exponential speedups over brute force methods.

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!

6.49 Automating Data Science

Organizers: Tijl De Bie, Luc De Raedt, Holger H. Hoos, and Padhraic Smyth
Seminar No. 18401

Date: September 30–October 5, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.9.154

© Creative Commons BY 3.0 Unported license

© Tijl De Bie, Luc De Raedt, Holger H. Hoos, and Padhraic Smyth



Participants: Leman Akoglu, Mitra Baratchi, Michael R. Berthold, Hendrik Blockeel, Pavel Brazdil, Ray G. Butler, Remco Chang, Felipe Leno da Silva, Tijl De Bie, Luc De Raedt, Peter Flach, Paolo Frasconi, Elisa Fromont, Jose Hernandez-Orallo, Holger H. Hoos, Frank Hutter, Tobias Jacobs, Lars Kotthoff, Nada Lavrac, Kevin Leyton-Brown, Jeffrey Lijffijt, Zhengying Liu, Siegfried Nijssen, Andrea Passerini, María Pérez-Ortiz, Bernhard Pfahringer, Kai Puolamäki, Matteo Riondato, Ute Schmid, Marc Schoenauer, Michele Sebag, Padhraic Smyth, Alexandre Termier, Stefano Teso, Heike Trautmann, Isabel Valera, Matthijs van Leeuwen, Joaquin Vanschoren, Jilles Vreeken, Andreas Wierse, Christopher Williams

■ Introduction

Data science is concerned with the extraction of knowledge and insight, and ultimately societal or economic value, from data. It complements traditional statistics in that its object is data as it presents itself *in the wild* (often complex and heterogeneous, noisy, loosely structured, biased, etc.), rather than data well-structured data sampled in carefully designed studies.

Such ‘Big Data’ is increasingly abundant, while the number of skilled data scientists is lagging. This has raised the question as to whether it is possible to automate data science in several contexts. First, from an artificial intelligence perspective, it is related to the issue of “robot scientists”, which are concerned with the automation of scientific processes and which have so far largely focused on the life sciences. It is interesting to investigate whether principles of robot scientists can be applied to data science.

Second, there exist many results in the machine learning community, which has since the early 1980s been applying machine learning at a meta-level, in order to learn which machine learning algorithms, variants and (hyper-)parameter settings should be used on which types of data sets.

In recent years, there have been breakthroughs in this domain, and there now exist effective systems (such as Auto-WEKA and auto-sklearn) that automatically select machine learning methods and configure their hyperparameters in order to maximize the predictive performance on particular datasets.

Third, there are projects such as the Automated Statistician that want to fully automate the process of statistical modeling. Such systems could dramatically simplify *scientific data modeling* tasks, empowering scientists from data-rich scientific disciplines such as bioinformatics, climate data analysis, computational social science, and so on. To ensure success, important challenges not only from a purely modelling perspective, but also in terms of interpretability and the human-computer interface, need to be tackled. For example, the input to the Automated Statistician is

a dataset, and the system produces not only a complex statistical model by means of a search process, but also explains it in natural language.

Fourth, there is an interest in not only automating the model building step in data science, but also various steps that precede it. It is well known in data science that 80% of the effort goes into preprocessing the data, putting it in the right format, and selecting the right features, whereas the model-building step typically only takes 20% of the effort. This has motivated researchers to focus on automated techniques for data wrangling, which is precisely concerned with transforming the given dataset into a format that can be handled by the data analysis component. Here, there are strong connections with inductive programming techniques.

Fifth, as it is often easier for non-expert users to interpret and understand visualisations of data rather than statistical models, work on automatic visualisation of data sets is also very relevant to this Dagstuhl seminar.

Finally, an interesting and challenging research question is whether it is possible to develop an integrated solution that tackles all these issues (as is the topic of the ERC AdG SYNTH).

■ Overview of the seminar

■ Structure of the seminar

The seminar was structured as follows. The mornings were generally dedicated to presentations (short tutorials on day one), whereas the afternoons were generally dedicated to discussions such as plenary discussions, smaller-group breakout sessions, and flex time that was kept open prior to the seminar. The flex time ended up being dedicated to a mix of presentations and breakout sessions.

■ Challenges in automating data science

On day one, a range of challenges for research on automating data science were identified, which can be clustered around the following six themes:

1. **Automating Machine Learning (AutoML)**
Main challenges: computational efficiency; ensuring generalization also for small data; make AutoML faster and more data-efficient using meta-learning; extending ideas from AutoML to exploratory analysis / unsupervised learning.
2. **Exploratory data analysis and visualization**
Main challenges: the fact that there is no single or clearly defined objective; help the user make progress towards an ill-defined goal; (subjective) interestingness of an analysis, a pattern, or a visualization; integrate machine learning and interaction in exploration; exploration of data types beyond simply tabular; veracity of visualizations; how to quantify progress and measure success; the need for benchmarks.
3. **Data wrangling**
Main challenges: extend the scope of AutoML to include data wrangling tasks; user interfaces to provide intuitive input in data wrangling tasks; how to quantify progress and measure success; the need for benchmarks.
4. **Automation and human-centric data science (explainability, privacy, fairness, trust, interaction)**
Main challenges: build-in privacy and fairness constraints in automatic data science systems; the dangers of ignorant usage of automated data science systems; different levels of expertise benefit from different degrees of automation; optimizing the performance of the combined human/machine 'team'; determine when and where the human must be involved; definition or criteria for explainability; risk that automation will reduce explainability and transparency; explainability to whom – a data scientist or layperson?
5. **Facilitating data science by novel querying and programming paradigms**
Main challenges: interactive data models to help users gain intuitive understanding; declarative approaches for data analysis, querying, and visualization; a query language for automated data science.
6. **Evaluation**
Main challenges: robust objective measures for data science processes beyond predictive modelling; subjective measures: measures that depend on the user background and goals; evaluation of the entire data science pipeline versus individual steps; reproducibility in the presence of user interactions.

■ Topics discussed in depth

These identified challenges were then used to determine the program of the rest of the seminar. Talks were held on partial solutions to a range of these challenges. In addition, breakout discussions were held on the following topics:

1. The relation between data-driven techniques and knowledge-based reasoning.
2. Data wrangling.
3. Beyond the black-box: explainability.
4. Automation of exploratory / unsupervised data science tasks, and visualization.
5. Automating data science for human users.

Along with abstracts of the talks, detailed discussions of the main ideas and conclusions of each of these breakout sessions are included in this Dagstuhl report.

■ Discussion and outlook

Automating data science is an area of research that is understudied as such. AutoML, as a subarea of automating data science, is arguably the first subarea where some remarkable successes have been achieved. This seminar identified the main challenges for the field in translating these successes into advances in other subareas of automating data science, most notably in automating exploratory data analysis, data wrangling and related tasks, integrating data and knowledge-driven approaches, and ultimately the data science process as a whole, from data gathering to the creation of insights and value.

Further developing automated data science raises several challenges. A first challenge concerns the evaluation of automated data science methods. Indeed, the possibility to automate is preconditioned on the availability of criteria to optimize. A second key one is how to ensure that automated data science systems remain Human-Centric, viewing humans as useful allies and ultimate beneficiaries. This can be achieved by designing effective user-interaction techniques, by ensuring explainability, and by ensuring privacy is respected and individuals are treated fairly. These are basic requirements for ensuring justified trust in automated data science systems, and thus key drivers to success.

6.50 Progressive Data Analysis and Visualization

Organizers: Jean-Daniel Fekete, Danyel Fisher, Arnab Nandi, and Michael Sedlmair

Seminar No. 18411

Date: October 7–12, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.1

© Creative Commons BY 3.0 Unported license

© Jean-Daniel Fekete, Danyel Fisher and Michael Sedlmair



Participants: Marco Angelini, Michael Aupetit, Sriram Karthik Badam, Carsten Binnig, Remco Chang, Jean-Daniel Fekete, Danyel Fisher, Hans Hagen, Barbara Hammer, Christopher M. Jermaine, Jaemin Jo, Daniel A. Keim, Jörn Kohlhammer, Stefan Manegold, Luana Micallef, Dominik Moritz, Thomas Mühlbacher, Hannes Mühleisen, Themis Palpanas, Adam Perer, Nicola Pezzotti, Gaëlle Richer, Florin Rusu, Giuseppe Santucci, Hans-Jörg Schulz, Michael Sedlmair, Charles D. Stolper, Hendrik Strobelt, Cagatay Turkay, Frank van Ham, Anna Vilanova, Yunhai Wang, Chris Weaver, Emanuel Zraggen

We live in an era where data is abundant and growing rapidly; databases to handle big data are sprawling out past memory and computation limits, and across distributed systems. New hardware and software systems have been built to sustain this growth in terms of storage management and predictive computation. However, these infrastructures, while good for data at scale, do not support data exploration well.

The concept of exploratory data analysis (EDA) was introduced by John Tukey in the 1970's and is now a commonplace in visual analytics. EDA allows users to make sense of data with little or no known model; it is essential in many application domains, from network security and fraud detection to epidemiology and preventive medicine. For most datasets, it is considered a best practice to explore data before beginning to construct formal hypotheses. Data exploration is done through an iterative loop where analysts interact with data through computations that return results, usually shown with visualizations. Analysts reacting to these results and visualizations issue new commands triggering new computations until they answer their questions.

However, due to human cognitive constraints, exploration needs highly responsive system response times (see <https://www.nngroup.com/articles/powers-of-10-time-scales-in-ux/>): at 500 ms, users change their querying behavior; past five or ten seconds, users abandon tasks or lose attention. As datasets grow and computations become more complex, response time suffers. To address this problem, a new computation paradigm has emerged in the last decade under several names: online aggregation in the database community; progressive, incremental, or iterative visualization in other communities. It consists of splitting long computations into a series of approximate results improving with time; the results are then returned at a controlled pace.

This paradigm addresses scalability problems, as analysts can keep their attention on the results of long analyses as they arrive progressively. Initial research has shown promising results in

progressive analysis for both big database queries and for machine learning.

The widespread use of progressive data analysis has been hampered by a chicken-and-egg problem: data visualization researchers do not have online database systems to work against, and database researchers do not have tools that will display the results of their work. As a result, progressive visualization systems are based on simulated systems or early prototypes. In many cases, neither side has currently incentive, skills, or resources to build the components needed.

Recently, data analysis researchers and practitioners have started conversations with their colleagues involved in the data analysis pipeline to combine their efforts. This standard pipeline includes the following core communities: data management, statistics and machine learning and interactive visualization. These initial conversations have led to fruitful evolutions of systems, combining two or three of these communities to complete a pipeline. Database and visualization have collaborated to create systems allowing progressive, approximate query results. Machine-learning and visualization have collaborated to create systems combining progressive multidimensional projections with appropriate scalable visualizations, such as Progressive t-SNE. Most current machine learning algorithms are designed to examine the entirety of a dataset. A major contribution of work like Progressive t-SNE is to have a decomposable algorithm that can compute a meaningful partial result, which then can be passed on to a visual interface for fluent exploration. In these few existing collaborations, the researchers are able to work together and find concrete mechanisms by adapting existing systems for these without re-building them from the ground up. A systematic and widespread linkage between the involved communities, however, is still largely absent.

This Dagstuhl seminar brought the researchers and practitioners who have started this software evolutionary process to

exchange their ideas, experience, and visions. We are convinced that in the forthcoming years, progressive data analysis will become a leading paradigm for data exploration systems, but will require major changes in the algorithms and data structures in use today. The scientific communities involved need to understand the constraints and possibilities from their colleagues to converge faster, with a deeper awareness of the implications of this paradigm shift. The implications are technical, but also human, both perceptual and cognitive, and the seminar will provide a holistic view of the problem by gathering specialists from all the communities.

This summary summarizes the outcomes of our seminar. The seminar focused on

- defining and formalizing the concept of progressive data analysis,
- addressing fundamental issues for progressive data analysis, such as software architecture, management of uncertainty, and human aspects,
- identifying evaluation methods to assess the quality of progressive systems, and threats to research on the topic,

- examining applications in data science, machine learning, and time-series analysis.

As a major result from the seminar, the following problems have been identified:

1. Implementing fully functional progressive systems will be difficult, since the progressive model is incompatible with most of the existing data analysis stack,
2. The human side of progressive data analysis requires further research to investigate how visualization systems and user interfaces should be adapted to help humans cope with progressiveness,
3. The potentials of progressive data analysis are huge, in particular it would reconcile exploratory data analysis with big data and modern machine learning methods,
4. Yet, there are many threats that should be addressed to bring progressive data analysis and visualization mainstream in research and application domains.

6.51 Encouraging Reproducibility in Scientific Research of the Internet

Organizers: Vaibhav Bajpai, Olivier Bonaventure, Kimberly Claffy, and Daniel Karrenberg
Seminar No. 18412

Date: October 7–10, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.41

© Creative Commons BY 3.0 Unported license

© Vaibhav Bajpai, Olivier Bonaventure, Kimberly Claffy, and Daniel Karrenberg



Participants: Vaibhav Bajpai, Steve Bauer, Olivier Bonaventure, Anna Brunström, Kenneth L. Calvert, Georg Carle, Kimberly Claffy, Alberto Dainotti, Anja Feldmann, Ralph Holz, Luigi Iannone, Daniel Karrenberg, Wolfgang Kellerer, Robert Kistelevi, Mirja Kühlewind, Andra Lutu, Matt Mathis, Jörg Ott, Aiko Pras, Damien Saucez, Quirin Scheitle, Jürgen Schönwälder, Henning Schulzrinne, Georgios Smaragdakis, Karen Sollins, Joel Sommers, Brian Trammell, Matthias Wählisch, Klaus Wehrle, John Wroclawski, Thomas Zinner

Reproducibility in scientific research is a means to not only achieve trustworthiness of results, but it also lowers barriers to technology transition [21] and accelerates science by promoting incentives to data sharing. The networking research community however pays limited attention to the importance of reproducibility of results, instead tending to accept papers that appear plausible. Previous studies [10, 15, 22] have shown that a fraction of published papers release artifacts (such as code and datasets) that are needed to reproduce results. In order to encourage reproducibility of research, practitioners continue [7, 11, 14, 17, 20] to do service to educate the community on the need for this change. To provide incentives to authors, vehicles for publication of software and datasets are also emerging. For instance, Elsevier SoftwareX [3] is a new journal designed to specifically publish software contributions. DataCite [13, 19] provides mechanisms for supporting methods to locate and cite datasets. Community Resource for Archiving Wireless Data (CRAWDAD) [23] and Information Marketplace for Policy and Analysis of Cyber-risk & Trust (IMPACT) Cyber Trust [4] provide an index of existing measurement data to not only enable new research but also advance network science by promoting reproducible research. Traditional conferences bestow best dataset awards and actively solicit submissions that reproduce results. SIGCOMM Computer Communication Review (CCR) allows authors to upload artifacts during paper submission to allow reviewers to check for reproducibility, and relaxes page limits for reproducible papers. Association for Computing Machinery (ACM) has recently introduced a new policy [1] on result and artifact review and badging. The policy identifies a terminology to use to assess results and artifacts. ACM has also initiated a new task force on data, software and reproducibility in publication [6] to understand how ACM can effectively promote reproducibility within the computing research community. National Academies of Sciences, Engineering, and Medicine with a goal to move

towards the open science ecosystem has recently (2018) released a report [16] with guidance and concrete recommendations on how to build strategies for achieving open science. The target is to ensure the free availability (and usability) of publications and associated artifacts. The National Science Foundation (NSF) is taking substantial steps [2] in this area whereby submitted proposals are required to provide a results dissemination plan to describe how produced research results are made available to the extent necessary to independently validate the findings. Towards this end, the proposal budget [5] may request funds for the costs of documenting, preparing, publishing or otherwise making available to others the findings and products of the work conducted under the NSF grant. Despite these continued efforts, reproducibility of research exist as an ongoing problem and few papers that reproduce existing research get published [9, 12, 18] in practise.

■ Goals

In this seminar, we discussed challenges to improving reproducibility of scientific Internet research, developed a set of recommendations that we as a community can undertake to initiate a cultural change toward increased reproducibility of our work. The goal of the seminar was to discuss the questions below and to propose recommendations that would improve the state of reproducibility in computer networking research.

■ What are the challenges with reproducibility?

How can researchers (and data providers) navigate concerns with openly sharing datasets? How should we cope with datasets that lack stable ground truth?

The first category of questions tried to identify the challenges with reproducibility [8]. For instance, concerns with openly sharing datasets led to discussions around legal restrictions and

the advantages of researchers keeping data private for their own exclusive future use. Another consideration is double-blind review practices, which require that authors expend effort to obfuscate the source of their data. Would this time be better spent documenting the datasets for sharing to enable reproducibility? A “gap analysis” discussion to understand whether the problem is a lack of appropriate venues or lack of stable ground truth, or more broadly a lack of incentive to reproduce research since publishing (and funding) agents tend to prefer novelty was held. There is also the inherent risk of confirmation bias of existing results; discussion of ideas on how to train young researchers to recognize and counter this tendency was sought.

■ What incentives are needed to encourage reproducibility?

What can publishers do? What can conference organisation committees do? How can we ensure that reviewers consider reproducibility when reviewing papers? How can we manage and scale the evaluation of artifacts during peer review? Do we need new venues that specifically require reproducibility of the submitted research?

The second category of questions is about incentives. Questions about how publishers can promote reproducibility framed discussions on whether publishers can provide storage for authors to upload data artifacts with the associated paper in digital libraries, or whether mechanisms can be developed to highlight reproducible (and reproduced) papers. Questions on how conference organisation committees can inspire ideas for additional incentives (such as best dataset awards or relaxing page limits) for authors to make research reproducible. We identified questions to add to review forms to ensure reviewers pay attention to reproducibility aspects. This further led to discussions on whether committees (in parallel to the regular technical program committee) should evaluate artifacts during the conference review process. Should such a committee be composed of purely young researchers or a blend of young and senior researchers? Questions on the need for specific venues triggered discussions on whether high-impact journals need to establish feature topics on reproducibility or devote a dedicated column for papers that reproduce existing research.

■ What tools and systems are available to facilitate reproducibility?

How effective are emerging interactive lab notebook tools (e.g., Jupyter) at facilitating reproducibility? Should CS course curricula integrate use of these tools for student

projects to help develop skills and habits that enable reproducibility?

The third category of questions attempt to identify and review tools and systems that are available to facilitate reproducibility. Enormous interest has developed recently in tools for recording experimental observations and computational analytics on large data sets. Some researchers now document the entire process for a paper in a Jupyter lab notebook, greatly facilitating reproducibility and extension of the research. The learning curve for these tools may be daunting; we discussed how faculty can evolve CS course curricula to integrate use of these tools for student projects to help develop skills and habits that enable reproducibility.

■ What guidelines or (best practises) are needed to help reproducibility?

How can we ensure authors think about reproducibility? What guidelines would assist reviewers in evaluating artifacts?

The fourth category of questions attempts to develop guidelines (or best practises) to promote reproducibility of research. For instance, we discussed what language could be added to Call for Papers (CFP) to encourage authors to describe reproducibility aspects (of both measurements and results) in their paper submissions.

■ Structure

The seminar lasted 2.5 days. The seminar began with an introductory round where each participant presented one slide to give an overview of their experience that is relevant for the seminar and a set of open questions that the participant wished to discuss during the event. These slides were collected from each participant before the seminar. We had one invited talk (§3.1 of the full report) that we used as a basis for triggering discussions and identifying areas for group work, while a major portion of the seminar time was dedicated to breakout sessions, whereby participants were split into small groups to discuss specific themes and develop ideas with consensus to propose to larger groups. The morning sessions the following day were dedicated to continuing parallel group work with presentations that reported the outcomes of each breakout session from the previous day. In the afternoons, we dedicated some time for seven minute lightning talks to invite ideas for subsequent breakout sessions. One evening, we had a social dinner activity. The afternoon of the third day was spent reviewing and collecting feedback from the participants and to initiating follow up actions identified during the seminar.

■ References

- 1 ACM Artifact Review and Badging. <https://www.acm.org/publications/policies/artifact-review-badging> [Online; last accessed 29-December-2018].
- 2 Computer and Network Systems (CNS): Core Programs. <https://www.nsf.gov/pubs/2018/nsf18569/nsf18569.htm> [Online; last accessed 29-December-2018].
- 3 Elsevier SoftwareX. <https://www.journals.elsevier.com/softwarex> [Online; last accessed 29-December-2018].
- 4 IMPACT Cyber Trust. <https://www.impactcybertrust.org> [Online; last accessed 31-December-2018].
- 5 NSF Proposal & Award Policies & Procedures Guide (PAPPG). <https://www.nsf.gov/pubs/policydocs/pappg18> [Online; last accessed 29-December-2018].
- 6 ACM Task Force on Data, Software and Reproducibility in Publication, 2015. <https://www.acm.org/publications/task-force-on-data-software-and-reproducibility> [Online; last accessed 29-December-2018].
- 7 Vaibhav Bajpai, Arthur W. Berger, Philip Eardley, Jörg Ott, and Jürgen Schönwälder. Global Measurements: Practice and Experience (Report on Dagstuhl Seminar #16012). *Computer Communication Review*, 46(2):32–39, 2016. doi:10.1145/2935634.2935641.
- 8 Vaibhav Bajpai, Mirja Kühlewind, Jörg Ott, Jürgen Schönwälder, Anna Sperotto, and Brian Trammell. Chal-

- lenges with Reproducibility. In *SIGCOMM 2017 Reproducibility Workshop*, pages 1–4, 2017. doi:10.1145/3097766.3097767.
- 9 Bryan Clark, Todd Deshane, Eli M. Dow, Stephen Evanchik, Matthew Finlayson, Jason Herne, and Jeanna Neeffe Matthews. Xen and the Art of Repeated Research. In *USENIX Annual Technical Conference*, pages 135–144, 2004. URL: <http://www.usenix.org/publications/library/proceedings/usenix04/tech/freenix/clark.html>.
 - 10 Christian S. Collberg and Todd A. Proebsting. Repeatability in Computer Systems Research. In <http://doi.acm.org/10.1145/2812803>, Communications of the ACM, 2016. doi:10.1145/2812803.
 - 11 Philip Eardley, Marco Mellia, Jörg Ott, Jürgen Schönwälder, and Henning Schulzrinne. Global Measurement Framework (Dagstuhl Seminar 13472). *Dagstuhl Reports*, 3(11):144–153, 2013. doi:10.4230/DagRep.3.11.144.
 - 12 Heidi Howard, Malte Schwarzkopf, Anil Madhavapeddy, and Jon Crowcroft. Raft Refloated: Do We Have Consensus? In *Operating Systems Review*, pages 12–21, 2015. doi:10.1145/2723872.2723876.
 - 13 Catherine Jones, Brian Matthews, and Ian Gent. Software reuse, repurposing and reproducibility. In *Proceedings of the 12th International Conference on Digital Preservation, iPRES 2015, Chapel Hill, North Carolina, USA, November 2-6, 2015*, 2015. URL: <http://hdl.handle.net/11353/10.429590>.
 - 14 Balachander Krishnamurthy, Walter Willinger, Phillipa Gill, and Martin F. Arlitt. A Socratic Method for Validation of Measurement-based Networking Research. In *Computer Communications*, pages 43–53, 2011. doi:10.1016/j.comcom.2010.09.014.
 - 15 Stuart Kurkowski, Tracy Camp, and Michael Colagrosso. MANET Simulation Studies: The Incredibles. In *Mobile Computing and Communications Review*, pages 50–61, 2005. doi:10.1145/1096166.1096174.
 - 16 National Academies of Sciences, Engineering, and Medicine. *Open Science by Design: Realizing a Vision for 21st Century Research*. The National Academies Press, 2018. doi:10.17226/25116.
 - 17 Vern Paxson. Strategies for Sound Internet Measurement. In *Internet Measurement Conference, (IMC)*, pages 263–271, 2004. doi:10.1145/1028788.1028824.
 - 18 Diana Andreea Popescu and Andrew W. Moore. Reproducing Network Experiments in a Time-controlled Emulation Environment. In *Traffic Monitoring and Analysis (TMA)*, 2016. URL: <http://tma.ifip.org/2016/papers/tma2016-final10.pdf>.
 - 19 Laura Rueda, Martin Fenner, and Patricia Cruse. Datacite: Lessons learned on persistent identifiers for research data. *IJDC*, 11(2):39–47, 2016. doi:10.2218/ijdc.v11i2.421.
 - 20 Geir Kjetil Sandve, Anton Nekrutenko, James Taylor, and Eivind Hovig. Ten Simple Rules for Reproducible Computational Research. *PLOS Computational Biology Journal*, 2013. doi:10.1371/journal.pcbi.1003285.
 - 21 Henning Schulzrinne. Networking Research - A Reflection in the Middle Years. In *Computer Communications*, 2018. doi:10.1016/j.comcom.2018.07.001.
 - 22 Patrick Vandewalle, Jelena Kovacevic, and Martin Vetterli. Reproducible Research in Signal Processing. *IEEE Signal Processing Magazine*, 2009. doi:10.1109/MSP.2009.932122.
 - 23 Jihwang Yeo, David Kotz, and Tristan Henderson. CRAWDAD: A Community Resource for Archiving Wireless Data at Dartmouth. *SIGCOMM Computer Communication Review*, pages 21–22, 2006. doi:10.1145/1129582.1129588.



Fig. 6.13

“A little bit of Dagstuhl history. Originally when Dagstuhl opened in 1990 they prided themselves on being isolated from the world. They bought this radio to keep up with the Iraq war. Now they have WiFi throughout.”

Twitter post by 18391 Dagstuhl Seminar participant Lance Fortnow.

<https://twitter.com/fortnow/status/1044248721414860801>. Photo courtesy of Lance Fortnow.

6.52 Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative

Organizers: Henning Fernau, Petr A. Golovach, and Marie-France Sagot
Seminar No. 18421

Date: October 14–19, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.63

© Creative Commons BY 3.0 Unported license

© Marie-France Sagot, Dieter Kratsch, Henning Fernau, and Petr A. Golovach



Participants: Faisal N. Abu-Khzam, Ricardo Andrade, Cristina Bazgan, Daniel Binkele-Raible, Andreas Björklund, Oscar Defrain, Arnaud Durand, Martin Dyer, Khaled M. Elbassioni, Michael R. Fellows, Henning Fernau, Fedor V. Fomin, Serge Gaspers, Petr A. Golovach, Benjamin Gras, Anne-Sophie Himmel, Leandro Ishi Soares de Lima, Giuseppe F. Italiano, Mamadou Moustapha Kanté, Petteri Kaski, Mathieu Liedloff, Daniel Lokshtanov, Kazuhisa Makino, Alberto Marchetti-Spaccamela, Andrea Marino, João Marques-Silva, Arnaud Mary, Lhouari Nourine, Michal Pilipczuk, Jean-Florent Raymond, Frances A. Rosamond, Günter Rote, Marie-France Sagot, Saket Saurabh, Mohamed Yosri Sayadi, Martin Schirneck, Blerina Sinimeri, Leen Stougie, Yann Strozecki, Takeaki Uno, Heribert Vollmer, Kunihiko Wasa, Gerhard J. Woeginger, Petra Wolf

About fifty years ago, NP-completeness became the lens through which Computer Science views computationally hard (decision and optimization) problems. In the last decades various new approaches to solve NP-hard problems exactly have attracted a lot of attention, among them parameterized and exact exponential-time algorithms, typically dealing with decision and optimization problems.

While optimization is ubiquitous in computer science and many application areas, relatively little is known about enumeration within the “Algorithms and Complexity” community. Fortunately there has been important algorithmic research dedicated to enumeration problems in various fields of Computer Science, as, e.g., Artificial Intelligence and Data Mining, in Natural Sciences Engineering and Social Sciences.

Enumeration problems require to list all wanted objects of the input as, e.g., particular subsets of the vertex or edge set of a given graph or particular satisfying assignments of logical expressions. Contrary to decision, optimization and even counting problems, the output length of enumeration problems is often exponential in the size of the input and cannot be neglected. This motivates the classical approach in enumeration, now called output-sensitive, which measures running time in (input and) output length, and asks for output-polynomial algorithms and algorithms of polynomial delay. This approach has been studied since a long time and has produced its own important open questions, among them the question whether the minimal transversals of a hypergraph can be enumerated in output-polynomial time. This longstanding and challenging question has triggered a lot of research. It is open for more than fifty years and most likely the best known open problem in algorithmic enumeration.

Recently as a natural extension on research in exact exponential-time algorithms, a new approach, called input-sensitive, which measures the running time in the input length, has found

growing interest. Due to the number of objects to enumerate (in the worst case), the corresponding algorithms have exponential running time. So far branching algorithms are a major tool. Input-sensitive enumeration is strongly related to lower and upper combinatorial bounds on the maximum number of objects to be enumerated for an input of given size. Such bounds can be achieved via input-sensitive enumeration algorithms but also by the use of combinatorial (non-algorithmic) means.

The area of algorithmic enumeration is in a nascent state, though it has a huge potential due to theoretical challenges and practical applications. While output-sensitive enumeration has a long history, input-sensitive enumeration has been initiated only recently. Natural and promising approaches like using parameterized or approximative approaches have not been explored yet in their full capacities.

The principal goals of our Dagstuhl seminar were to increase the visibility of algorithmic enumeration within (Theoretical) Computer Science and to contribute to establishing it as an area of “Algorithms and Complexity”. The seminar brought together researchers within the algorithms community, other fields of Computer Science and Computer Engineering, as well as researchers working on enumeration problems in other application areas, in particular, in Biology. Besides the people already working with enumeration, researchers from other fields of Computer Science were invited. In particular, researchers who are interested in Parameterized Complexity and different aspects of counting problems were participating in the seminar. The aim was to accelerate developments and discuss new directions including algorithmic tools and hardness proofs.

The seminar collected 44 participants from 13 countries. The participants presented their recent results in 18 invited and contributed talks. Open problems were discussed in several open problem and discussion sessions.

6.53 Shape Analysis: Euclidean, Discrete and Algebraic Geometric Methods

6

Organizers: Michael Breuß, Alfred M. Bruckstein, Christer Oscar Kiselman, and Petros Maragos

Seminar No. 18422

Date: October 14–19, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.87

© Creative Commons BY 3.0 Unported license

© Michael Breuß, Alfred M. Bruckstein, Christer Oscar Kiselman, and Petros Maragos

Participants: Marianne Akian, Jesús Angulo, Péter Balázs, Michael Breuß, Alfred M. Bruckstein, Sara Brunetti, Vasileios Charisopoulos, Robert Dachsel, Leila De Floriani, Panagiotis Filntisis, Federico Iuricich, Tao Ju, Bert Jüttler, Joe Kileel, Ron Kimmel, Christer Oscar Kiselman, Zorah Lähner, Petros Maragos, Roberto Mecca, Gautam Pai, Frank R. Schmidt, William Smith, Sibel Tari, Silvia Tozza, Martin Welk, Laurent Younes



Dagstuhl seminar 18422 Shape Analysis: *Euclidean, Discrete and Algebraic Geometric Methods* took place October 14–19, 2018. 26 researchers from North America and Europe discussed state-of-the-art, current challenges, and promising future research directions in the areas of 2-D and 3-D shape analysis from a cross-disciplinary point of view. Participants included international experts from the fields of continuous-scale shape analysis, discrete shape analysis, tropical geometry and numerical computing. The seminar consisted of an opening and getting to know session and 26 scientific presentation sessions. Furthermore, there was time for extensive discussions both between the talks and in the evenings.

The topics in our seminar addressed the sketched challenges and developments that will be useful for shape analysis. Especially we aimed to discuss the possibilities of combining fields like tropical geometry with more classical techniques as for instance from mathematical morphology. We discussed possibilities of applying machine learning methods in this context and consid-

ered recent advances from more classical fields like differential geometry and partial differential equations that can be useful for setting up and understanding shape analysis methods in all of these approaches.

The purpose of this seminar was to address these challenges with the latest tools related to geometric, algorithmic and numerical concepts. To do so, we brought together researchers working on shape analysis topics from different perspectives. The purpose in bringing together researchers from those different communities sharing substantial interest in shape analysis was to explore the benefits of a cross-disciplinary point of view.

Promising new ways to combine the latest techniques from these different fields were identified during in-depth discussions in small groups. Some especially promising research directions in the areas of deep learning, mathematical morphology, shape from shading, modelling deformable shapes, and tropical geometry were discussed in small groups between the talks and in the evenings.

6.54 Computational Aspects of Fabrication

Organizers: Bernd Bickel, Marc Alexa, Jessica K. Hodgins, and Kristina Shea
Seminar No. 18431

Date: October 21–26, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.104

© Creative Commons BY 3.0 Unported license

© Marc Alexa, Jessica K. Hodgins, and Kristina Shea



Participants: Marc Alexa, Thomas Auzinger, Vahid Babaei, Moritz Bächer, Amit Haim Bermano, Sabine Demey, Laura Devendorf, Georges Fadel, Madeline Gannon, Serena Graziosi, Jessica K. Hodgins, Scott Hudson, Alec Jacobson, Lin Sebastian Kayser, Leif Kobbelt, Manfred Lau, Sylvain Lefebvre, David I. W. Levin, Lin Lu, Jonas Martinez-Bayona, Sara McMains, Radomir Mech, Ankur Mehta, Shuhei Miyashita, Caitlin Mueller, Julian Panetta, Mark Pauly, Nadya Peek, Nico Pietroni, Thijs Roumen, Ryan Schmidt, Carolyn C. Seepersad, Kristina Shea, Melina Skouras, Bernhard Thomaszewski, Nobuyuki Umetani, Philipp Urban, Charlie Wang, Tim Weyrich, Christopher Bryant Williams, Maria Yablouina

As manufacturing goes digital, we are facing a fundamental change in the workflow of fabrication. While access to advanced digital fabrication and 3D-printing technology becomes ubiquitous and provides new possibilities for fabricating complex, functional, multi-material objects with unconventional properties, its potential impact is currently limited by the lack of efficient and intuitive methods for content creation. Existing tools are usually restricted to expert users, have been developed based on the capabilities of traditional manufacturing processes, and do not sufficiently take fabrication constraints into account. Scientifically, we are facing the fundamental challenge that existing simulation techniques and design approaches for predicting the physical properties of materials and objects at the resolution of modern 3D printers fail to scale well with possible object complexity.

To achieve significant progress, we need a deep understanding of interdisciplinary fundamentals: Shape, Appearance of Shape and Materials, Validated Simulation, and Engineering Design. The purpose of this Dagstuhl Seminar is to bring together leading experts from academia and industry in the area of computer graphics, geometry processing, mechanical engineering, human-

computer interaction, material science, and robotics. The goal is to address fundamental questions and issues related to computational aspects of fabrication, build bridges between related fields, and further pioneer this area.



There has been a considerable growth in the number of articles treating aspects of computational fabrication, scattered across multiple disciplines and journals. In this seminar we gathered together these various threads and described the computational accomplishments and outstanding challenges. Researchers from different communities analyzed which existing fabrication workflows could benefit most from computation and identify novel application domains, with the aim of cross-fertilizing ideas between disciplines. The main goal of this seminar was identifying and reporting common grand challenges and developing a roadmap for addressing them. Additionally, the seminar sought to discuss and establish standards and best practices for sharing research results, code, and hardware prototypes, facilitating reproducibility and reusability of results among disciplines. An important aspect of this was to analyze teaching and learning needs for new students in the field, and coordinating the development of teaching material.



BLOCKCHAIN SECURITY AT SCALE


Leibniz

BEN-OR'S FRAMEWORK




IT WORKS | ? | NOPE!
 → ADDS RESILIENCE TO SECURITY

MANIPULATING INCENTIVES




PROOF OF PERSONHOOD

ARE WE DECENTRALIZED YET. COM



REDESIGNING BITCOIN'S FEE MARKET




PROBLEM: BLOCKSIZE AFFECTS REVENUE

PAYMENT PRIVACY


MIXING: COIN JOIN / COINSHUFFLE
 STANDALONE: ZCASH, MONERO

→ SET ANONYMITY → DP SCHEME?
 SMALL ← | → USER-BASE SIZE


PERON VIRTUAL CHANNELS



METABOLIC SCALING



STARK INTERACTIVE PROOFS



INCENTIVE COMPATIBILITY (IC)

- 1) DEFINE "INTENDED" BEHAVIOR
 - 2) PROVE IT'S AN "EQUILIBRIUM"
- 1) PLAYERS
 2) STRATEGIES } GAME
 3) PAYOFFS
- DOMINANT STRATEGY (DS): ALWAYS BEST RESPONSE
 → DSIC

SCHLOSS DAGSTUHL

Fig. 6.14
 "Here is a summary of my notes as #sketchnotes after a week @dagstuhl on #Blockchain Security at Scale (<http://dagstuhl.de/18461>); I probably took more notes during this week than the past year combined."
 Twitter post by 18461 Dagstuhl Seminar participant Florian Tschorsch.
<https://twitter.com/flotschorsch/status/1063503170222075904>. Photo courtesy of Florian Tschorsch.

6.55 Data Physicalization

Organizers: Jason Alexander, Petra Isenberg, Yvonne Jansen, Bernice E. Rogowitz, and Andrew Vande Moere
Seminar No. 18441

Date: October 28–November 2, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.127

© Creative Commons BY 3.0 Unported license

© Jason Alexander, Petra Isenberg, Yvonne Jansen, Bernice E. Rogowitz, Andrew Vande Moere



Participants: Jason Alexander, Stephen Barrass, Stephen Brewster, Sheelagh Carpendale, Tim Dwyer, Leanne Elias, Yuri Engelhardt, Sean Follmer, Denton Fredrickson, Christian Freksa, Pauline Gourlet, Ian Gwilt, Sarah Hayes, Uta Hinrichs, Trevor Hogan, Eva Hornecker, Samuel Huron, Jörn Hurtienne, Petra Isenberg, Daniel Keefe, David Kirsh, Roberta Klatzky, Till Nagel, Bettina Nissen, Lora Oehlberg, Dietmar Offenhuber, Laura Perovich, Bernice E. Rogowitz, Daniel K. Schneider, Volker Schweisfurth, Adrien Segal, Aurélien Tabard, Barbara Tversky, Brygg Ullmer, Andrew Vande Moere, Karin von Ompeda, Yun Wang, Wesley J. Willett

This executive summary gives a brief overview of the goals and agenda for our seminar.

■ Introduction

Among the earliest man-made artifacts are physical representations of semantic concepts. These “physicalizations” externalize numerical and abstract concepts, providing physical metaphors that allow us to reason, remember, and communicate. Physicalizations are created for many different purposes, or “intents.” The sundial, for example, transforms shadows into a readable representation of time of day, the mercury thermometer transforms temperature into a displacement along a number line, and a scatterplot transforms the values of two variables into a form that allows the reader to interpret correlation. Our first data representations were based on natural objects, such as charcoal scraped onto walls or built from clay, or later, ink on paper. With the advent of computers, we’ve substituted physical representations with pixels on a computer screen. The resurgence of physicalization asks what we have lost in this transformation. Certainly, a computer-based visualization allows us to zoom an image, transform variables in real time, and to zoom through virtual computer-based world. However, these representations can sever the relationship to the natural world, depriving us of the touch, feel, and emotion that comes from interacting with real objects.

This Dagstuhl seminar brought together a diverse group of researchers and practitioners to explore the benefits and challenges of physicalization – computer scientists trained in visualization, virtual reality and human-computer interaction; architects of virtual and augmented systems; perceptual and cognitive scientists; and artists and designers. Through interactive discussions and demonstrations, we explored physicalization, as a set of methodologies for representing data, for engaging

audiences, and for artistic expression. There were no formal paper presentations. Instead, the work was done through interactive discussions, hands-on workshops, and interactive demonstrations outside the lecture hall. Fig. 6.15 shows two examples of data physicalizations exhibited by our participants.

■ The Week at a Glance

Monday. After explaining various organization matters and formalities, the first day began with the physical construction of custom name badges. The hands-on session offered the first occasion to get to know each other, also by comparing the name badges themselves, then and during the whole duration of the seminar. After this session, each of the 40 participants gave a two-minute presentation of their achievements and interests in the field of physicalization. Some described their artistic creations based on data, some showed how mapping data onto physical dimensions enhanced data analysis. Others showed how scientists and artists collaborated to explain biological and physical principles through physicalization, or how principles of perception and cognition could be used to guide how data are effectively mapped onto visual dimensions. The electricity in the air was palpable; everyone realized that we were on the threshold of a new discipline. This energy drove the first deliverable for this seminar: Dagstuhl monograph, and perhaps a book, highlighting extended versions of these 40 contributions.

To allow a more focused approach towards yet unexplored research topics, the first activity focused on defining “pillars” of physicalization, crystallizing the learning and background of the different intellectual communities identified in our proposal. In break-out sessions, each group ideated around and then synthesized the most fundamental papers, examples, principles and challenges for their pillar, relative to physicalization. The Perception and Cognition Group focused on sensory processing,

especially touch perception and embodied cognition, which deals with the way we learn about the world through our motor interactions with it. The Evaluation group shared their experiences in measuring human responses, and how methods might be extended to physicalization. The Design group explored artistic and design approaches to data physicalization. The Applications group identified existing and future application areas. The technology group surveyed the range of materials and devices for physicalization. In the reports-back, the group was encouraged to explore challenges and limitations for each of these core areas, to set the stage for the cross-disciplinary discussions beginning on Tuesday.

Tuesday. At the start the second day, each participant generated three questions for each of the five pillar areas. A massive exercise was undertaken wherein all these questions were organized and grouped into emergent categories. Spontaneous discussions started around the meaning and validity of the themes that emerged, as these would align the next round of thematic discussions (Fig. 6.16). The “Design Patterns” group focused on identifying general templates for characterizing data physicalization. The “Emotion” group explored unique affordances that touching data enable. The Vis vs. Phys Group delved into the ways vision and touch were different, and what unique advantages that might enable. The “Critical” group explored the range of ethical and critical matters that could be related to physicalization practice in particular.

Wednesday. Wednesday morning was devoted to hands-on workshops where practitioners engaged small groups in interactive activities. Samuel Huron’s group explored a set of physicalization examples, to identify common principles; Daniel Keefe demonstrated a virtual reality exploration of the human heart; Till Nagel, Laura Perovich, and Dietmar Offenhuber led a group into the forest to collect natural objects which they used to create physicalizations; Robert Friska provided surprising problem solving examples that drove a discussion about physical reasoning; Barbara Tversky showed examples of how physical gestures we make contribute to problem solving; Daniel Schneider led a workshop on computational embroidery, where the data are represented in yarn color and texture. By the time we piled into the bus for Trier, the large and diverse group of 40 had transformed into a dynamic community of scholars. Listening into the conversations, you could hear artists explaining how sculpture is taught in Art School, perceptual psychologists describing touch perception; and engineers revealing expertise on control systems for autonomous micro-robots.

Thursday. Thursday was the major work day. Building on the explorations of Tuesday, we divided into groups and worked on a variety of topics. One group, for example, focused on how to teach visualization and how to use data physicalization to teach other areas. Another group worked on a white paper on emotion and physicalization. Another worked on ideas for using physicalization for environmentally-situated projects. Another group focused on categorizing the critical considerations that emerged on Tuesday.

Friday. The final day aimed into transforming the progress of the past days into concrete contributions. For example, plans for a physicalization contest were discussed, ideas for the book of 40 contributions were pushed further, as well as on a special issue on physicalization for the Journal of Perceptual Imaging. Another group discussed the topic of scientific funding, among which the experience of past H2020 research proposals, and the potential of future research proposal initiatives that relate to physicalization.

■ Hands-on Demonstrations

The area outside the main conference room was populated with demonstrations provided by the participants. For example, Andrew Vande Moere brought a wireless, networked display that allowed participants to vote about particular topics that were shown; Volker Schweinsfurth, Daniel Schneider, and many others brought along and exhibited 3-D printed, hand-crafted, or even embroidered physicalizations. Bernice Rogowitz ran participants through several experiments around the topics of touch perception and touch/vision interactions. Section 4 of the full report details several of the hands-on experiences seminar participants and organizers could engage in; while Fig. 6.17 shows the diversity of the physical artifacts that participants exhibited.



Fig. 6.15 Examples of data physicalizations exhibited by our seminar participants. Left: Data objects for visualizing the experience, expression, and description of chronic pain. Right: Acoustic Sonifications whose sounds convey useful information about a dataset.



Fig. 6.16 A collaborative affinity diagramming session on Tuesday, to organize the large amount of thematic research questions into cross-disciplinary themes of research interest.



Fig. 6.17
Exhibition tables outside our seminar room.

6.56 Visualization and Processing of Anisotropy in Imaging, Geometry, and Astronomy

Organizers: Andrea Fuster, Evren Özarslan, Thomas Schultz, and Eugene Zhang
Seminar No. 18442

Date: October 28–November 2, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.10.148

© Creative Commons BY 3.0 Unported license

© Andrea Fuster, Evren Özarslan, Thomas Schultz, Eugene Zhang



Participants: David Bommers, Bernhard Burgeth, Tom Dela Haije, Ali Can Demiralp, Aasa Feragen, Luc Florack, Andrea Fuster, Shekoufeh Gorgi Zadeh, Hans Hagen, Magnus Herberthson, Ingrid Hotz, Jana Hutter, Daniel Jörgens, Andreas Kleefeld, Chunlei Liu, Timothy Luciani, Evren Özarslan, Renato Pajarola, Reynier Peletier, Marco Pizzolato, Gerik Scheuermann, Thomas Schultz, Chantal Tax, Bei Wang, Donggang Wang, Carl-Fredrik Westin, Hsiang-Yun Wu, Eugene Zhang, Yue Zhang, Jovisa Zunic

■ Topics and Motivation

Directional preferences or anisotropies are encountered across many different disciplines and spatial scales. For example, local anisotropies are imprinted in the cosmic microwave background radiation, the human brain contains elongated nerve fibers, etc. Such anisotropies lead to (physical) orientation-dependent quantities, i.e., quantities that take on different values when considered along different directions. Compared to scalar or vector-valued data, it is much more challenging to model, process, and visualize anisotropic quantities. Suitable mathematical models often involve tensors and other higher-order descriptors, and pose specific research challenges in several areas of computer science, such as visualization, image analysis, and geometry processing.

In order to explore synergies between different fields, to inform computer scientists about open application challenges, and domain experts about existing solutions, this seminar brought together researchers from three different disciplines:

- Medical **imaging**, where several modalities are now available to probe anisotropic behavior. In particular, Diffusion Weighted Magnetic Resonance Imaging (DW-MRI) is based on measuring anisotropic diffusion. It makes it possible to visualize and quantify microstructural information in fibrous tissues such as white-matter and muscles, and to infer larger-scale structures, such as fiber tracts in the human brain.
- Computer graphics and **geometry processing**, where tensor fields have a wide range of applications, such as quadrangular and hexahedral geometry remeshing, street network modeling, geometry synthesis, computational architecture, and path planning for environment scans.
- Cosmology and **astronomy**, where anisotropy plays a crucial role. For example, anisotropies in the cosmic microwave background (CMB) consist of small temperature fluctuations in the blackbody radiation left over from the Big Bang.

Anisotropies are also found in the CMB in the form of a polarization tensor field, and they arise in the field of “cosmography”, where efforts are united to map (parts of) the cosmos, e.g. the large-scale distribution of matter in the Universe or cosmic web.

■ Organization of the Seminar

This seminar was the seventh in a series of Dagstuhl seminars that was started in 2004, and has been devoted to the visualization and processing of tensor fields and higher-order descriptors. This particular instance of the seminar series focused on anisotropy in the fields of imaging, geometry, and astronomy.

To ensure a steady inflow of new ideas and challenges, we put an emphasis on inviting researchers who previously did not have the opportunity to attend one of the meetings in this series. This was true for almost half the attendees in the final list of participants.

The seminar itself started with a round of introductions, in which all participants presented their area of work within 100 seconds with help of a single slide. This helped to create a basis for discussion early on during the week, and was particularly useful since participants came from different scientific communities, backgrounds, and countries.

A substantial part of the week was devoted to presentations by 29 participants, who spent 20 minutes each on presenting recent advances, ongoing work, or open challenges, followed by ten minutes of discussion in the plenary, as well as in-depth discussions in the breaks and over lunch. Abstracts of the presentations are collected in this report. On Wednesday we held the traditional social event which was joined by almost all

participants, and offered additional welcome opportunities for interaction.

A total of six breakout sessions were organized in the afternoons of Monday and Tuesday. Moderators summarized the respective discussion in the plenary on Thursday afternoon. The organizers came up with initial suggestions for session topics, which were refined further after discussion with the seminar participants. The session topics were as follows:

- Astronomy
- Time-varying anisotropy
- Theoretical tools
- Visualization
- Diffusion MRI
- Geometry

Notes were taken during all sessions, and the main points are summarized later in this report.

■ Outcomes

The participants all agreed that the meeting was inspiring and successful. It also stimulated new scientific collaborations and joint grant proposals. In addition we plan to publish another Springer book documenting the results of the meeting. Participants have pre-registered seventeen chapters already during the seminar, and we are in the process of collecting additional contributions both from participants and from researchers working on closely related topics who could not attend the meeting. We expect that the book will be ready for publication in 2020.

■ Acknowledgment

The organizers thank all the attendees for their contributions and extend special thanks to the moderators of the breakout sessions and the team of Schloss Dagstuhl for helping to make this seminar a success. As always, we enjoyed the warm atmosphere, which supports both formal presentations as well as informal exchanges of ideas.

6.57 Genomics, Pattern Avoidance, and Statistical Mechanics

Organizers: Michael Albert, David Bevan, Miklós Bóna, and István Miklós
Seminar No. 18451

Date: November 4–9, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.11.1

© Creative Commons BY 3.0 Unported license

© Miklós Bóna



Participants: Michael Albert, David Bevan, Miklós Bóna, Mathilde Bouvel, Marília Braga, Robert Brignall, Cedric Chauve, Anders Claesson, Michael W. Deem, Sergi Elizalde, Andrew Elvey Price, Péter L. Erdős, Guillaume Fertin, Yoong Kuan Goh, Torin Greenwood, Sylvie Hamel, Christine E. Heitsch, E. J. Janse van Rensburg, László Kozma, Anthony Labarre, Olya Mandelshtam, István Miklós, Alois Panholzer, Jay Pantone, Seth Pettie, Yann Ponty, Svetlana Poznanovik, Thomas Prellberg, Pijus Simonaitis, Fiona Skerman, Jakub Sliacan, Heather Smith, Jason Smith, Rebecca Smith, Einar Steingrímsson, Jens Stoye, Jessica Striker, Krister Swenson, Vincent Vatter, Stéphane Vialette

This report documents the program and the outcomes of Dagstuhl Seminar 18451 “Genomics, Pattern Avoidance, and Statistical Mechanics”.

The workshop took place from November 4, 2018 to November 9, 2018. It had 40 participants, who were researchers in theoretical computer science, combinatorics, statistical mechanics and molecular biology. It was a geographically diverse group, with participants coming from the US, Canada, Brazil, Germany, Iceland, the United Kingdom, Sweden, France, Switzerland, Hungary, Australia, and New Zealand. The workshop featured 21 talks, three of which were hourlong talks, and an open problem session.

Several collaborative projects have been started. For example, Jay Pantone, Michael Albert, Robert Brignall, Seth Pettie, and Vince Vatter started exploring the topic of 1324-avoiding permutations with a bounded number of descents, disproving a 2005 conjecture of Elder, Rechnitzer, and Zabrocki related to Davenport-Schinzel sequences. Had the conjecture been affirmed, it would have implied that the generating function for 1324-avoiding permutations is non-D-finite.

At the open problem session, Yann Ponty raised the following question: what is the number of independent sets in restricted families of trees, like caterpillars or complete binary plane trees? The main motivation for this question relates to a deep connection between such independent sets and RNA designs. This question led to a new collaborative effort by Mathilde Bouvel, Robert Brignall, Yann Ponty and Andrew Elvey Price.

Sergi Elizalde and Miklós Bóna have started working on Dyck paths that have a unique maximal peak. That collaboration since extended to the area of probabilistic methods, involving a researcher working in that field, Douglas Rizzolo.

Numerous participants expressed their pleasure with the workshop and its sequence of talks. The prevailing view was that while the participants came from three different fields, they were all open to the other two fields, and therefore, they all learned about results that they would not have learned otherwise. Therefore, we have all the reasons to believe that the workshop was a success, and we would like to repeat it some time in the future.

6.58 Blockchain Security at Scale

Organizers: Rainer Böhme, Joseph Bonneau, Ittay Eyal
Seminar No. 18461

Date: November 11–16, 2018 | Dagstuhl Seminar
 Full report – DOI: 10.4230/DagRep.8.11.21
 © Creative Commons BY 3.0 Unported license
 © Rainer Böhme

Participants: Svetlana Abramova, Sarah Azouvi, Foteini Baldimtsi, Eli Ben-Sasson, Alex Biryukov, Rainer Böhme, Joseph Bonneau, Mic Bowman, Dominic Breuker, Christian Cachin, Nicolas Christin, Lisa Eckey, Ittay Eyal, Bryan Ford, Christina Garman, Arthur Gervais, Philipp Jovanovic, Aljosha Judmayer, Ghassan Karame, Assimakis Agamemnon Kattis, Stefan Katzenbeisser, Patrik Keller, Ron Lavi, Patrick McCorry, Ian Miers, Malte Möser, Tyler W. Moore, Neha Narula, Tim Roughgarden, Tim Ruffing, Emin Gün Sirer, Yonatan Sompolinsky, Itay Tsabary, Florian Tschorsch, Marie Vasek, Roger Wattenhofer, Edgar Weippl, Aviv Zohar



The security of blockchain-based systems has attracted great interest in the research community following the initial financial success of Bitcoin. Several security notions for blockchain-based systems have been proposed, varying in degree of formality and applicability to real-world systems. However, a major blind spot remains about the environment surrounding blockchain-based systems. This environment is typically assumed to be static (irresponsible to activities of the blockchain system). This is a sound starting point for security analysis while the stakes involved are small compared to the environment (i. e., the global economic and political system). However, if blockchain-based systems truly offer compelling advantages over legacy systems, they may eventually become the dominant form of organizing certain social choice problems. This “scale change” challenges the assumption that the blockchain-based system remains below the threshold of relevance for the parts of its environment that are vital for its security. One instance where this may already occur is the influence of mining puzzles on hardware design and electricity prices.

The purpose of the seminar was to bring together researchers with expertise in various subfields of blockchain-based systems to jointly revisit security foundations. The primary goal was to incorporate explicit consideration of reciprocity effects between properties of cryptocurrency protocols and their environment.

The primary intended outcome of this seminar was proposing a new design principle, viewing security as a key scalability property to consider in addition to performance and efficiency. Second, the seminar aimed to converge on standard terminology for security notions that are robust to scale. Third, we applied this new methodology to Bitcoin specifically as a test case, producing

a sort-of “break glass in case of rampant runaway growth” security plan.

Specific questions were:

1. **From micro-level to macro-level incentives** Bitcoin’s ecosystem remains small relative to large multinational corporations. What happens to incentives when a cryptocurrency reaches a scale similar to large national economies?
2. **Cryptographic agility** How does the ability to upgrade cryptographic algorithms might change in the future as cryptocurrency protocols become widely embedded in hardware and/or codified in the law.
3. **Reciprocity effects on hardware design** How will the hardware industry be affected by the increasing importance of superior hardware for mining, and possibly trusted execution environments (TEE) in the future?
4. **Mining economics at scale** How will mining economics change in the future, in particular, dynamics between miners at large-scale power consumption levels, with mass availability of cheap commodity mining hardware (including TEE-based), and with different incentives, e.g., in a high-valued fee-only revenue model.
5. **Reconsidering non-monetary incentives** Can cryptocurrencies be resilient to disruptive nation-level attacks that are not due to monetary incentives?
6. **Governance at scale** To date, cryptocurrencies largely rely on informal leadership from a small group of influential software developers. Can this be translated into a more democratic model? What does democratic control mean for a cryptocurrency when the *demos* is not clearly defined?

6.59 Provenance and Logging for Sense Making

Organizers: Jean-Daniel Fekete, T. J. Jankun-Kelly, Melanie Tory, and Kai Xu
Seminar No. 18462

Date: November 11–16, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.11.35

© Creative Commons BY 3.0 Unported license

© Kai Xu, Jean-Daniel Fekete, T. J. Jankun-Kelly, and Melanie Tory



Participants: Sara Alspaugh, Daniel Archambault, Simon Attfield, Leilani Battle, Christian Bors, Remco Chang, Christopher Collins, Michelle Dowling, Alex Endert, Jean-Daniel Fekete, Melanie Herschel, T. J. Jankun-Kelly, Andreas Kerren, Steffen Koch, Robert Kosara, Olga A. Kulyk, Robert S. Laramée, Sérgio Lifschitz, Aran Lunzer, Phong H. Nguyen, William Pike, Ali Sarvghad, Claudio T. Silva, Holger Stitz, Melanie Tory, John Wenskovich, William Wong, Kai Xu, Michelle X. Zhou

Sense making is one of the biggest challenges in data analysis faced by both industry and research community. It involves understanding the data and uncovering its model, generating hypothesis, select analysis methods, creating novel solutions, designing evaluation, and the critical thinking and learning wherever needed. Recently many techniques and software tools have become available to address the challenges of so-called ‘Big Data’. However, these mostly target lower-level sense making tasks such as storage and search. There is limited support for the higher-level sense making tasks mentioned earlier. As a result, these tasks are often performed manually and the limited human cognition capability becomes the bottleneck, negatively impacting data analysis and decision making. This applies to both industry and academia. Scientific research is a sense making process as well: it includes all the sense making tasks mentioned earlier, with an emphasis on the generation of novel solutions. Similar to data analysis, most of these are conducted manually and considerably limit the progress of scientific discovery.

Visual Analytics is a fast-growing field that specifically targets sense making [6]. It achieves this by integrating interactive visualization with data analytics such as *Machine Learning*. It follows a human-centered principle: instead of replacing human thinking and expertise with algorithms and models, it enables the two to work together to achieve the best sense making result. Fast progress has been made in the last decade or so, which is evidenced by the publications in the Visual Analytics conferences such as IEEE VAST (part of IEEE VIS) and the increasing popularity of visual approaches in many other fields such as Machine Learning, Information Retrieval, and Databases.

One recent advance in Visual Analytics research is the capture, visualization, and analysis of *provenance* information. Provenance is the history and context of sense making, including

the “7W” (Who, When, What, Why, Where, Which, and How) of data used and the users’ critical thinking process. The concept of provenance is not entirely new. In 1996, Shneiderman recognized the importance of provenance by classifying *history* as one of the seven fundamental tasks in data visualization [4]. History allows users to review previous actions during visual exploration, which is typically long and complex. Provenance can provide an overview of what has been examined and reveal the gap of unexplored data or solutions. Provenance can also support collaborative sense making and communication by sharing the rich context of what others have accomplished [7].

The topic of provenance has been studied in many other fields, such as Human-Computer Interaction (HCI), WWW, Database, and Reproducible Science. The HCI research community heavily relies on user information, such as logging and observation, in their study. These closely relate to provenance and share the common goal of making sense of user behavior and thinking. The collaboration between the two fields can potential create novel solutions for some long-standing research challenges. For instance, it has been shown that provenance information can be used to semi-automate part of the qualitative analysis of user evaluation data [3], which is notoriously time-consuming.

The WWW and Database research community has been actively working on provenance for the last decade or so, with a particular focus on tracking data collection and processing. This has led to the recent publication of the W3C reference model on provenance⁴⁰. A important part of these efforts is to make sense of the source and quality of the data and the analyses base on them, which has a significant impact on their uncertainty and trustworthiness [1]. Similarly, there is a fast growing Reproducible Science community, whose interest in provenance is “improving the reliability and efficiency of scientific research

⁴⁰ <https://www.w3.org/TR/prov-overview/>

... increase the credibility of the published scientific literature and accelerate discovery” [2].

There is a trend of cross-community collaboration on provenance-related research, which has led to some exciting outcomes such as the work integrating visualization with reproducible science [5, 8]. However, there are still many challenging research questions and many provenance-related research efforts remain disconnected. This seminar brought together researchers from the diverse fields that relate to provenance. Shared challenges were identified and progress has been made towards developing novel solutions.

The main research question that this seminar aims to address is: **How to collect, analyze, and summarize provenance information to support the design and evaluation of novel techniques for sense making across related fields.** The week-long seminar started with a day of self-introduction, lightning talks, and research topic brain storming. The self-introduction allowed

attendees to know each other better, and the lightning talks covered the latest work in the research fields related to provenance. Each participant proposed several research questions, which were then collated and voted on to form the breakout groups. The following are the research areas chosen by the participants:

- Storytelling and narrative;
- Provenance standard and system integration;
- Task abstraction for provenance analysis;
- Machine learning and provenance;
- User modeling and intent.

The rest of the week was breakout session, and each participant had the option to change group halfway. The seminar finished with a presentation from each group and discussions on the next steps to continue the collaboration. Many interesting problems were identified, and progress was made towards new solutions. Please refer to the rest of the report for the details on the identified research questions and the progress made by the end of week.

■ References

- 1 Melanie Herschel and Marcel Hlawatsch. Provenance: On and Behind the Screens. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD '16*, pages 2213–2217, New York, NY, USA, 2016. ACM.
- 2 Marcus R. Munafò, Brian A. Nosek, Dorothy V. M. Bishop, Katherine S. Button, Christopher D. Chambers, Nathalie Percie du Sert, Uri Simonsohn, Eric-Jan Wagenmakers, Jennifer J. Ware, and John P. A. Ioannidis. A manifesto for reproducible science. *Nature Human Behaviour*, 1(1):0021, January 2017.
- 3 P.H. Nguyen, K. Xu, A. Wheat, B.L.W. Wong, S. Attfield, and B. Fields. SensePath: Understanding the Sensemaking Process Through Analytic Provenance. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):41–50, January 2016.
- 4 Ben Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages*, pages 336–, Washington, DC, USA, 1996. IEEE Computer Society.
- 5 C. T. Silva, J. Freire, and S. P. Callahan. Provenance for Visualizations: Reproducibility and Beyond. *Computing in Science Engineering*, 9(5):82–89, September 2007.
- 6 James J. Thomas and Kristin A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. National Visualization and Analytics Centre, 2005.
- 7 Kai Xu, Simon Attfield, T.J. Jankun-Kelly, Ashley Wheat, Phong H. Nguyen, and Nallini Selvaraj. Analytic provenance for sensemaking: A research agenda. *IEEE Computer Graphics and Applications*, 35(3):56–64, 2015.
- 8 Zheguang Zhao, Lorenzo De Stefani, Emanuel Zraggen, Carsten Binnig, Eli Upfal, and Tim Kraska. Controlling false discoveries during interactive data exploration. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD '17*, pages 527–540, New York, NY, USA, 2017. ACM.

6.60 Next Generation Domain Specific Conceptual Modeling: Principles and Methods

Organizers: Heinrich C. Mayr, Sudha Ram, Wolfgang Reisig, and Markus Stumptner
Seminar No. 18471

Date: November 18–23, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.11.63

© Creative Commons BY 3.0 Unported license

© Heinrich C. Mayr



Participants: João Paulo Almeida, João Araújo, Fernanda Baião, Ladjel Bellatreche, Dominik Bork, Isabelle Comyn-Wattiau, Rocco De Nicola, Felicita Di Giandomenico, Hans-Georg Fill, Frederik Gailly, Aditya K. Ghose, Martin Glinz, Giancarlo Guizzardi, Kamalakar Karlapalem, Julio Cesar Leite, Stephen W. Liddle, Wolfgang Maaß, Heinrich C. Mayr, Judith Michael, Daniel Moldt, Oscar Pastor Lopez, Barbara Pernici, Jolita Ralyte, Wolfgang Reisig, Marcela Ruiz Carmona, Bernhard Rumpe, Heinz W. Schmidt, Michael Schrefl, Sibylle Schupp, Vladimir Shekhovtsov, Prina Soffer, Markus Stumptner, Ketil Stølen, Bernhard Thalheim, Pia Wilsdorf, Manuel Wimmer

Models are the basic human tools for managing complexity and understanding and therefore play a key role in all scientific and engineering disciplines as well as in everyday life. Many modeling paradigms have evolved over time into a wide variety of modeling languages, methods and tools that have come and gone. This is particularly true for Informatics, which is a modeling discipline in itself.

Since the 1970s, special attention has been paid to conceptual modeling. This approach essentially uses a formal language whose concepts are linked to a semantic interpretation (e.g. by the grounding in an ontology) and a more or less transparent graphic or textual representation (which supports efficient linguistic perception). Normally, such a language is embedded in a model/meta model hierarchy. The dimensions of conceptual modeling languages are structure, dynamics (behavior) and functionality.

Despite all efforts, however, there is still no comprehensive and consistent use of conceptual modeling in practice. Often conceptual models are only used as prescriptive documents, which – e.g. in the area of software development or business process management – are rarely synchronized with the developed artifact, so that reality and model diverge step by step. This observation motivated us to promote and conduct this seminar by focusing on domain-specific conceptual modeling, as this promises a methodology that is more tailored to the needs of each user group.

In view of the highly committed discussions during the seminar, the intensive discussions in the working groups and the very positive results of the participant survey, we can say without exaggeration that the seminar was a complete success. Almost all participants wished for a continuation, which we will probably apply for in 2021, when the already decided projects (cooperation and publications of subgroups) are on their way.

Since, with one exception, every participant wanted to present their ideas in a talk, the programme was tightly packed: 36 talks

and 2 full evening sessions in working groups, the results of which were reported on the next morning, made the seminar week a very intensive but also highly inspiring experience.

First results are already tangible:

- The workshop “Conceptual Modeling for Multi-Agent Problem Solving” at the IJCAI 2019 in Macao: The idea was born during the seminar and implemented afterwards: <http://austria.omilab.org/psm/content/cmmaps19/info>.
- A working group is currently writing a summary paper on the results of the working groups on which agreement was reached in the plenary discussions.
- Questions that were discussed during the seminar will be included in contributions to the Summer School “Next Generation Enterprise Modelling in the Digital Transformation” in Vienna (July 15-26, Vienna).
- The seminar organizers are currently writing a somewhat more popular scientific column to be submitted to CACM.

A number of open questions and “grand challenges” that also could be topics of future relevant conferences have been identified, among others:

- Business Transformations in the age of digitalization as “Models are driving the Digital Transformation”
- Social Aspects of Conceptual Modeling
- Explanatory Models for Neural Networks and Big Data
- Conceptual Modeling for validation purposes in simulation
- Modeling of Ultra Large Scale Architectures
- Privacy Modeling
- Modeling of Behavior Goals for Assistive Systems and Emotions
- Better integration into teaching at universities of applied sciences and universities
- Tools and Technical Infrastructures for Conceptual Modeling, in particular for “multi-metamodeling frameworks”

- Involvement of researchers and practitioners from other fields: “go beyond the obvious”.

The biggest challenge for a follow-up seminar will be to encourage more practitioners to participate. For this purpose, we will propose to dedicate two consecutive seminar days to this and the discussion with them, as practitioners usually cannot spend more time.

6.61 Implementing FAIR Data Infrastructures

Organizers: Natalia Manola, Peter Mutschke, Guido Scherp, Klaus Tochtermann, and Peter Wittenburg

Seminar No. 18472

Date: November 18–21, 2018 | Dagstuhl Perspectives Workshop

Full report – DOI: 10.4230/DagRep.8.11.91

© Creative Commons BY 3.0 Unported license

© Natalia Manola, Peter Mutschke, Guido Scherp, Klaus Tochtermann, and Peter Wittenburg



Participants: Marcel R. Ackermann, Luiz Dlavo Bonino da Silva Santos, Timothy W. Clark, Ron Dekker, Kees den Heijer, Michel Dumontier, Marie Farge, Sascha Friesike, Carole Goble, Kathleen Gregory, Gregor Hagedorn, Wilhelm Hasselbring, Oliver Kohlbacher, Paolo Manghi, Natalia Manola, Daniel Mietchen, Peter Mutschke, Heike Neuroth, Andreas Rauber, Marc Rittberger, Raphael Ritz, Guido Scherp, Birgit Schmidt, Achim Streit, Klaus Tochtermann, Dieter Van Uytvanck, Tobias Weigel, Mark D. Wilkinson, Peter Wittenburg

The Dagstuhl Perspectives Workshop on “Implementing FAIR Data Infrastructure” aimed at bringing together computer scientists and digital infrastructure experts from different domains to discuss challenges, open issues, and technical approaches for implementing the so-called FAIR Data Principles in research data infrastructures. Moreover, the workshop aimed to shape the role of and to develop a vision for computer science for the next years in this field, and to work out the potentials of computer science in advancing Open Science practices.

In the context of Open Science, and the European Open Science Cloud (EOSC) in particular, the FAIR principles seem to become a common and widely accepted conceptual basis for future research data infrastructures. The principles consist of the four core facets that data must be **F**indable, **A**ccessible, **I**nteroperable, and **R**eusable in order to advance the discoverability, reuse and reproducibility of research results. However, the FAIR principles are neither a specific standard nor do they suggest specific technologies or implementations. They describe the core characteristics of data use. Thus, the FAIR principles cover a broad range of implementation solutions. This certainly incorporates the risk of having a highly fragmented set of solutions at the end of the day.

Given this, and in view of the “need for a fast track implementation initiative [of the EOSC]”⁴¹, it is strongly needed to turn the principles into practice. Therefore, the workshop took the recommendations of the European Commission Expert Group on FAIR Data “Turning FAIR into reality” as a starting point and discussed what can be done next from the perspective of computer science to enable data providers to make their data FAIR.

The workshop started with three ignition talks on the wider background and context of the FAIR principles (given by Peter

Wittenburg), the relationship of FAIR to Open Data (given by Natalia Manola) and the role of the principles within the EOSC (given by Klaus Tochtermann). Based on these talks as well as inputs from all participants in the forefront of the workshop, we have split the discussion into three working groups addressing, for each of the four principles, the main key challenges for implementing FAIR and the question what and how computer science can contribute to these key challenges. Based on the results of these three initial working groups we furthermore split into more focused groups addressing the problem of licenses w.r.t. data use, (self)improvement of FAIRification, and the relation of FAIR and data intensive science.

Finally, we identified three major areas to be addressed in the manifesto which we discussed in three further working groups:

1. **Infrastructures & Services Aspects:** This group focused on the question by which technical means research data infrastructures and data services can be advanced to better address and fulfil the FAIR principles.
2. **Computer Science Research Topics:** The working group discussed the relationship of research areas in computer science and topics relevant to implement FAIR data infrastructures.
3. **FAIR Computer Science Research:** While the other two groups mainly focused on the contribution of computer science to implement FAIR, this working group addressed the question how the FAIR principles are currently adopted by computer science research itself and what should be improved.

The participants will continue their work in the aforementioned issues, and a manifesto is foreseen to be ready by mid May 2019.

⁴¹ <https://www.dtls.nl/wp-content/uploads/2017/05/DE-NL-Joint-Paper-FINAL.pdf>

Is next next to is and is next next to or and is next next to and and is next next to next and is next next to to or is or next to is and is or next to next and is or next to and and is or next to to or is and next to is and is and next to next and is and next to or and is and next to and and is and next to to or is to next to or and is to next to and and is to next to next and is to next to to and is to next to is or is is next to and and is is next to or and is is next to next and is is next to to and is is next to is?

Fig. 6.18

“#Dagstuhl poetry. Inspired by “Between”. Generated. @dagstuhl #mxdagstuhl”

Twitter post by 18491 Dagstuhl Seminar participant Vadim Zaytsev.

<https://twitter.com/grammarware/status/1070081279083847681>. Photo courtesy of Vadim Zaytsev.




Fig. 6.19

Dagstuhl Impression: Bridge – Between

6.62 High Throughput Connectomics

Organizers: Moritz Helmstaedter, Jeff Lichtman, and Nir Shavit
Seminar No. 18481

Date: November 25–30, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.11.112

© Creative Commons BY 3.0 Unported license

© Nir Shavit



Participants: Daniel R. Berger, Davi Bock, Kevin Briggman, Julia Buhmann, Albert Cardona, Forrest Collman, Nuno Maçarico da Costa, Winfried Denk, Eva Dyer, Rainer W. Friedrich, Jan Funke, Christel Genoud, Stephan Gerhard, William Gray Roncal, Moritz Helmstaedter, Michal Januszewski, Joergen Kornfeld, Anna Kreshuk, Julia Kuhl, Wei-Chung Allen Lee, Jeff Lichtman, Jeremy Maitin-Shepard, Yaron Meirovitch, Josh Morgan, R. Clay Reid, Kerriane Ryan, Stephan Saalfeld, Aravinthan D.T. Samuel, Louis Scheffer, Nir Shavit, Jochen Triesch, Xueying Wang, Adrian Wanner, Casimir Wierzynski, Dirk Zeidler

Our workshop brought together experts in the computational aspects of connectomics. A week of lectures and work-group meetings in a lively and collegial environment led to a collection of interesting conclusions. One big idea that was put forth in the meeting was the gargantuan effort of reconstructing a complete mouse brain. Another was to completely map the white matter connectivity of a mammalian brain. We also discussed which techniques/pipelines we should continue to pursue as a community. In that vein one big conclusion was that you have to have both the engineers and software working on a pipeline; distributing software only is not sufficient (you need dedicated engineers to run the software, it can't be based just on grad students). Zeiss reported on a multibeam 331 beam microscope that was in the making. There were also discussions on quality measures and metrics for connectomics reconstruction, and on developing standardized datasets for segmentation training and comparison of algorithms (scaling up from current day small datasets). Finally, there were discussions on the ethics and policies in the area going forward – Should we rely more on industrial partners to provide compute power and storage, or is it better to keep most of the research in universities and non-for-profit research institutes.

■ Introduction

The sheer complexity of the brain means that sooner or later the data describing brains must transition from something that is rather easily managed to something far less tractable. This transition appears to now be underway. The accumulation of ever-bigger brain data is a byproduct of the development of a number of new technologies that provide digitized information about the structural organization (anatomy) and the function of neural tissue. These new collection approaches bring novel data into neuroscience that potentially bears on many poorly understood aspects of the nervous system. Fundamental scientific

questions such as the way learned information is instantiated in the brain and how brains change over the course of development and aging are likely to be usefully addressed in the coming decades as large data sets mapping networks of neurons at high resolution become available.

Mapping networks of neurons at the level of synaptic connections, a field called connectomics, began in the 1970s with a study of the small nervous system of a worm and has recently garnered general interest thanks to technical and computational advances that automate the collection of electron-microscopy data and offer the possibility of mapping even large mammalian brains. However, modern connectomics produces 'big data', unprecedented quantities of digital information at unprecedented rates, and will require, as with genomics at the time, breakthrough algorithmic and computational solutions.

Unfortunately the generation of large data sets is actually the easy part. Our experience in the nascent field of connectomics indicates that there are many challenges associated with the steps after data acquisition, that is, the process of turning the data into a mineable commodity. This workshop will focus on addressing these challenges by bringing together researchers developing algorithms and deploying software systems that enable high-throughput analysis of connectomic data.

While high-throughput connectomics must tackle many of the problems that occur in big data science and engineering, tremendous differences in data size, computational complexity, and the problem domain will require novel computational frameworks, algorithms, and systems. Input image data in connectomics is reaching, even in its initial stages, petabytes in size at a terabytes-per-hour rate, and currently requires millions of cycles of computation per pixel. Such data is not easily moved or stored, and so on-the-fly analysis of the data as it comes off the microscope is the most likely future solution. Achieving the kind of throughput that will allow us to process the data at the rate at which it is being generated necessitates a three orders of

magnitude reduction in cycles per pixel, compared to the status quo. Furthermore, there is locality to the data. Unlike other big data problems, which can often be represented as independent key-value pairs spread across many machines, reconstruction of neural circuits requires frequent data exchanges across adjacent image regions. Buffering all the data in machine memory is infeasible, as is data replication on multiple servers. That means one cannot rely on Moore's law and parallelism across data centers to solve this problem—we need to be smarter.

In a nutshell, a connectomics data set is a collection of images taken on a volume of brain tissue that has been sectioned into many thousands of small slices, each only a few tens of nanometers thick. These slices are then imaged using custom electron microscopes to produce an image stack that will in the near future reach petabytes in size. Using one of the standard electron microscopy pipeline approaches, the key computational problems that must be addressed in order to turn the raw acquired digitized images into a useful form of “connectivity graph” are stitching, alignment, neuron reconstruction, and synapse detection. Each digitized image tile needs to be stitched together with neighboring tiles to form a composite image of a slice. Then, the stitched slice image is aligned with the previous and subsequent slice images. Despite being mostly similar, image alignment is challenging because typically a conveyor belt collects the slices and each may rotate a few degrees, or stretch depending on its thickness. Fortunately, because of the high image resolution, alignment is practical, as axons and dendrites are readily visible in cross-section and can be traced from one section to the next. A second challenge is that, once the image data is aligned, the sectioned objects must be individuated. In these data sets, the objects are neurons and other cellular entities that are interwoven in the three-dimensional space of the sample tissue. The reconstruction of neural processes as they pass from one section to the next is directly related to the computer vision problem of obtaining a segmentation of an image series, that is, the labeling of pixels in the images according to which cell they belong to.

Although considerable progress has been achieved in computer-based image segmentation in the last few years, reliable automatic image segmentation is still an open problem. Automating the segmentation of connectomic data is challenging because the shapes of neural objects are irregular, branching, non-repeating and intertwined. Moreover, the actual number of different objects and their synaptic interconnections in a volume of brain tissue is unknown and, at the moment, even difficult to estimate or bound. Segmentation of a standard electron microscopy image is further complicated by the fact that the range of pixel intensity values of cell membranes overlaps with that of other organelles. Thus, simple thresholding to find cell boundaries does not work.

In the eyes of many, the term big data is synonymous with the storage and analysis of massive collections of digital information. The “big” refers to the size of the input sets, typically ranging in the tens or even hundreds of terabytes, and arriving at rates of several tens or hundreds of gigabytes per second. In connectomics, the size of the input set is at the high end of the big data range, and possibly among the largest data ever acquired. Images at several nanometers resolution are needed to accurately reconstruct the very fine axons, dendrites, and synaptic connections. At this resolution, a cubic mm is about 2 petabytes of data. A complete rat cortex including some white matter might require 500 cubic mm and thus would produce about an exabyte (1000 petabytes) of data. This amount is far beyond the scope of storage that can be handled by any system today (as a reference point, consider that Walmart or Aldi's database systems manage a few petabytes of data). A complete human cortex, 1000-times that of a rodent, will require a zetabyte (1000 exabytes) of data, an amount of data approaching that of all the information

recorded in the world today. Obviously this means that the goal of connectomics will not be to acquire complete human brains and that for the near future one must consider reconstructions of neuronal substructures as opposed to whole brains. Moreover, it is clear that as we go beyond a few millimeters, one cannot store the raw data: it must be analyzed on the fly as it comes off the microscope and then discarded, keeping the physical tissue sample for re-imaging if needed.

What is this on-the-fly acquisition rate? The new multi-beam electron microscopes currently produced by Carl Zeiss LLC have a staggering throughput approaching 400 sections per day or a terabyte of data per hour, placing them at the far end of the big data rate spectrum. This rate, if it can be matched with appropriate reconstruction algorithms, will allow researchers to process a cubic mm of rodent brain, that is, 2 petabytes of data, in about 6 months operating 24 hours a day, 7 days a week. Whatever computational pipeline is used to extract the connectomics graph from the image data, it will eventually have to work on the fly, at the pace of the microscope that generates this data.

The algorithms and computational techniques for developing such high throughput connectomics pipelines are the target of this workshop. The massive amounts of storage and computation require expertise not only in computational neurobiology, machine learning, and alignment techniques, but also in parallel computation, distributed systems, and storage systems. There are several groups of researchers around the world that specialize in collecting the electron microscopy datasets, and several that engage in developing matching computational pipelines. Our aim is to bring these researchers together for an extended 5-day brainstorming session. We will also invite some top researchers in related fields such as machine learning, computer vision, distributed systems, and parallel computing. Our goal for this meeting is to both build an understanding of the state of the art in high-throughput connectomics pipelines, and to brainstorm on how to move the field forward so that high throughput connectomics systems become widely available to neurobiology labs around the world.

Concretely, we would like to come out of this workshop with a hierarchical plan for future connectomics systems that solve existing systems' problems. We will begin the workshop by having workgroups discuss these problems in existing systems and then dedicate the latter part to collectively working out solutions. We will consider three levels:

1. The system layer: how data is stored, moved around and computed on in a distributed and parallel fashion.
 2. The pipeline layer: how processing progresses from stitching through alignment and reconstruction.
 3. The algorithm layer: the specific machine learning and error detection and correction algorithms used in various pipeline stages to bring the datasets to analyzable connectivity graphs.
- Our plan is to discuss each of these in detail, with the hope of concluding the workshop with a coherent plan on how to proceed.

■ Relation to previous Dagstuhl seminars

To the best of our knowledge there have been no similar Dagstuhl seminars in the past. The field of connectomics is a young cutting edge big data research area that will have important implications on both computation in the sciences (and in particular on the use of large scale machine learning in the sciences) and on artificial intelligence (through the development of new neural network models based on the neurobiological discoveries this research may lead to). We believe it is important for modern computer science to engage in such interdisciplinary applications of computing and algorithms and we are therefore eager to initiate this new seminar direction.

6.63 Network Visualization in the Humanities

Organizers: Katy Börner, Oyvind Eide, Tamara Mchedlidze, Malte Rehbein, and Gerik Scheuermann

Seminar No. 18482

Date: November 25–30, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.11.139

© Creative Commons BY 3.0 Unported license

© Katy Börner, Oyvind Eide, Tamara Mchedlidze, Malte Rehbein, and Gerik Scheuermann



Participants: Gregor Betz, Katy Börner, Ulrik Brandes, Nicole Coleman, Melanie Conroy, Marten Düring, Tim Dwyer, Oyvind Eide, Kimmo Elo, Sara Irina Fabrikant, Christina Gillmann, Hans Hagen, Francis Harvey, Gerhard Heyer, Stefan Jänicke, Fotis Jannidis, Andreas Kerren, Stephen G. Kobourov, Tamara Mchedlidze, Martin Nöllenburg, Malte Rehbein, Gerik Scheuermann, Antonis Symvonis, Timothy Tangherlini, Christopher Warren, Scott Weingart, Florian Windhager

■ Seminar Goals

The application of computer-based methods by scholars of the Humanities has a tradition that goes back to the mid 20th century. Labelled “Digital Humanities” some 15 years ago, it has seen a significant growth since then [1]. An important part of Digital Humanities methodology is to establish data sets [2] based on cultural artefacts such as fiction texts, paintings, musical scores and recordings, and historical sources in all media. This is done in a number of different ways and includes some sort of extraction of data from sources structured in different, less explicit ways than what is needed for operationalisation and computer assisted analysis and visualization. When this process works well, it supports scholars’ endeavours to answer existing research questions and to generate new insights and novel research questions. A significant part of the data collected can be modeled as networks.

Existing network analysis and visualization techniques have already proven themselves immensely useful in analyzing Digital Humanities data and providing new discoveries [3]. The central goal of research on network visualization for digital humanities scholars is to develop visualization techniques and algorithms that empower scholars to use those effectively as part of their research process and for communicating study results to readers. While network science approaches are widely used in other research areas, the power of a network mindset and approach has not yet been fully exploited within the Humanities.

The seminar aimed to enhance the development of network visualization algorithms and tools centered around humanities research. In particular, its goals were as follows:

- **Interdisciplinary Exchange:** to discuss existing network visualization methods and algorithms in perspective of their potential application within the Humanities;
- **Terminology Gap:** Bridging the gap in terminology between

Digital Humanities on one side and computer scientists in Network Visualization and Graph Drawing on the other side;

- **Data:** to discuss Humanities’ data sources and their nature, research questions, use cases, and specific application profiles in perspective of their potential support by network visualisation.
- **Research Agenda:** Formulation of research agenda on “Network Visualization in the Digital Humanities”. Creation of interdisciplinary teams of researchers that address specific scientific challenges of the agenda;

■ Seminar Program

The seminar brought together 27 researchers from Network Visualization and Digital Humanities communities. The initial two days of the weeklong event were devoted to bring together the different communities and to develop a mutual understanding. Researchers informed each other about their scholarly background through short, five-minute talks. In addition, there were eight long, 45 minutes, presentations in which digital humanities scholars discussed network and network visualization challenges and opportunities in their field of expertise. This was complemented by surveys on network visualization and successful examples of cooperation between visualization and digital humanities researchers.

During both days the participants were asked to post questions and issues they would like to discuss in the remaining three days of the seminar. After a voting, four research areas most interesting the participant were identified. All four met the guiding principles in that they describe both: highly relevant applications within the Humanities as well as innovative research challenges for Network Visualisation. They are as follows:

- Complex networks, in particular multivariate, multilayered, and multilevel networks;
- Linked networks;
- Temporal networks;
- Uncertainty, incompleteness, and ambiguity of data.

Four groups were formed to work on those four topics over the remaining three days. There were several opportunities for joint discussions and progress reports across the groups. Summaries of the group discussions can be found in Section 4 of the full report.

■ Future Plans

During the seminar the participants decided to proceed with a publication of a manifesto, outlining a research agenda for “network visualisation in the Humanities”. It was also planned to publish an edited volume on specific aspects of the overarching topic, possibly along the four major research areas identified by the seminar. The volume will be submitted as a special issue to “Historical Network Research”, an Open Access Journal.

■ Evaluation

The feedback provided by the participants in form of a survey collected by Schloss Dagstuhl was highly positive and in most aspects above the average collected over the last 60 seminars. The participants agreed that the seminar inspired new ideas, collaborations, joint publications and brought insight from neighboring fields. There was a number of positive comments by the participants on the structure and organization of the seminar as well as several useful suggestions for the future seminars.

■ Acknowledgment

As an organizing committee of the seminar we would like to thank the scientific and administration staff of Schloss Dagstuhl for the excellent support they provided, both in the preparation phase and during the seminar. On behalf of all participants, we would also like to thank Dagstuhl for the high quality facilities provided, for excellent rooms for work and socializing, for the tasty meals, and of course also for the excellent wine cellar. The organizers of the seminar would also like to thank Ray Siemens and Dan Edelstein for their contributions to the initial Dagstuhl proposal. Finally, we thank Christina Gillmann for taking the responsibility for this report.

■ References

- 1 Thaller, Manfred (2017), *Geschichte der Digital Humanities*. In: Fotis Jannidis, Hubertus Kohle und Malte Rehbein (Hg.): *Digital Humanities. Eine Einführung*. Stuttgart, J.B. Metzler Verlag, S. 3–12.
- 2 Schöch, Christof (2013) *Big? Smart? Clean? Messy? Data in the Humanities*. In: *Journal of Digital Humanities* 2 (3).
- 3 Rehbein, Malte (forthcoming) *Historical Network Research, Digital History, and Digital Humanities*. In: Marten Düring, Florian Kerschbaumer, Linda von Keyserlingk und Martin Stark (Hg.): *The Power of Networks. Prospects of Historical Network Research*, Routledge.

6.64 Multidirectional Transformations and Synchronisations

Organizers: Anthony Cleve, Ekkart Kindler, Perdita Stevens, and Vadim Zaytsev
Seminar No. 18491

Date: December 2–7, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.12.1

© Creative Commons BY 3.0 Unported license

© Perdita Stevens, Ekkart Kindler



Participants: Anthony Anjorin, Gabor Bergmann, Dominique Blouin, James Cheney, Anthony Cleve, Sebastian Copei, Davide Di Ruscio, Zinovy Diskin, Jeremy Gibbons, Holger Giese, Martin Gogolla, Soichiro Hidaka, Michael Johnson, Tsushima Kanae, Gabor Karsai, Ekkart Kindler, Heiko Klare, Hsiang-Shang Ko, Harald König, Ralf Lämmel, Yngve Lamo, Théo Le Calvar, Nuno Macedo, Kazutaka Matsuda, James McKinna, Fiona A. C. Polack, Blake S. Pollard, Robert Rosebrugh, Bernhard Rumpe, Andy Schürr, Bran V. Selic, Friedrich Steimann, Perdita Stevens, Matthias Tichy, Frank Trollmann, Jens Holger Weber, Nils Weidmann, Bernhard Westfechtel, Vadim Zaytsev, Albert Zündorf

The Dagstuhl Seminar on “Multidirectional Transformations and Synchronisations” was the latest on a sequence of events [1] (coordinated by the Bx Steering Committee [2]) on bidirectional transformations (abbreviated **bx**) and related topics. Broadly speaking, the concern of the growing community interested in this topic is the *maintenance of consistency* between multiple data sources, in the presence of change that may affect any of them. The focus of this Dagstuhl meeting, in particular, was the special issues that arise when one considers more than two data sources at one time. Technical definitions of **bx** have always allowed for there to be more than two, but in practice, most work to date has focused on maintaining consistency between two data sources. We abbreviate “multidirectional transformation”, hereinafter and generally, to **multx** or **mx**⁴².

We began the week with a presentation of case studies and introductory tutorials, and towards the end of the week we had several plenary talks, in cases where someone was able to discuss work that seemed likely to be of general interest. Overall, though, this was a highly interactive Dagstuhl: most of our time during the week was spent either in working groups or, synergistically, discussing in plenary session the outcomes of those working groups and what else needed to be addressed. Reports from each of those working groups, and descriptions of the case studies, are found in the body of this report. Here we briefly introduce them. Inevitably, the topics of the groups overlapped, and some topics that were proposed for working groups were not reached during the week. We encouraged participants to move freely between groups to foster cross-fertilisation of ideas. The names given are those of the authors of the brief reports. Collectively, these topics comprise a research roadmap for the subject.

- **WG1: Whether Networks of Bidirectional Transformations**

Suffice for Multidirectional Transformations.

This group began what turned out to be a recurring theme of the week: see below.

- **WG2: Partial Consistency Notions.**
This relates to handling situations in which consistency is not perfectly restored, but only improved to some extent.
- **WG3: Semantics of Multidirectional Transformations.**
This group raised questions about definitions of syntactic and semantic consistencies, vertical and horizontal propagation, etc. After creating enough awareness of the importance of this topic, the working group dissolved and its efforts were merged into others, in particular WG4, WG8 and WG12.
- **WG4: Multiple Interacting Bidirectional Transformations.**
This group started with an intention of providing a good example of a “truly” multidirectional transformation and defined scenarios where several **multx** and **bx** work together towards restoring consistency.
- **WG5: Mathematical Backgrounds for Multidirectional Transformations.**
Following on from WG1, this group considered, from a theoretical perspective, handling **multx** by the use of a common “federated” supermodel related by spans of asymmetric lenses.
- **WG6: Synchronisation Policy.**
Separate from the issue of what the mechanism is to restore consistency, when should the mechanism be used, and whose decides that?
- **WG7: Use Cases and the Definition of Multidirectional Transformations.**
When are **multx** really necessary in practice, and how?
- **WG8: Human Factors: Interests of Transformation Developers and Users.**

⁴² Consensus on just one of those two options was not achieved by this workshop!

Sometimes in our focus on technical aspects we lose the human element – who are the humans involved and what do they need from whatever languages, tools and techniques are developed for multx?

- **WG9: Provenance in Multidirectional Transformations.**
Information about what changed and why – provenance and traceability – are crucial to trust in multx; how can that information be provided and handled?
- **WG10: Living in the Feet of the Span.**
Following on from WG1 and WG5: what happens when one conceptually uses a common supermodel, but does not wish to materialise it?
- **WG11: Programming Languages for Multidirectional Transformations.**
This group discussed the challenges that need to be met to produce such languages, with a focus on their type systems.
- **WG12: Verification and Validation of Multidirectional Transformations.**
What needs to be verified or validated about multx, and in what ways do these needs challenge the state of the art in verification and validation?

A recurring theme of the week, turning up in one guise or another in most of the working groups, was the question of the extent to which excellent solutions to the two-source bx problems would, or would not, automatically solve the multx problems. Do problems involving multx really introduce new issues, or are they just more complicated than problems involving bx, perhaps organised in networks? The bx problem is far from solved – we do not yet, for example, have widely adopted and well-supported specialist bx languages – and so there was some feeling that we lacked a firm foundation on which to address multx. More positively, considering multx has the potential to help bx research make progress, by motivating areas that still require more study in order to support the multx case. For example, heterogeneity of

the languages in which the data sources and the changes to them are expressed clearly points to a need for bx, and hence multx, approaches that do not need to materialise edit histories for all the sources, nor a common supermodel for them – even if the theory that addresses them might still call on such things conceptually.

The following case study descriptions are also included in the report:

- Multidirectional Transformations for Microservices
- Multidirectionality in Compiler Testing
- Bringing Harmony to the Web
- A Health Informatics Scenario

Perhaps the most important observation from this Dagstuhl meeting, though, is how broad the scope of the research necessary to address multx concerns is. The issues and examples discussed by participants went far beyond multidirectional versions of issues and examples already raised in earlier bx meetings. For example, microservices, the focus of the case study presented by Albert Zündorf (§4.1 of the full report), would not traditionally have fallen under the bx umbrella, yet is clearly related.

This widening of scope is natural. As IT systems become more interdependent and more important to our everyday lives, it is inevitable that data, and the (often separately developed) behaviour it supports, reside in many places. They are coupled, in the sense that changes in one place may mean that changes in another place are necessary, in order to maintain all of these systems in useful operation. Making all such changes manually does not scale: some degree of automated maintenance of consistency is inevitably required. Multx thus subsumes much of software engineering and inherits its concerns.

Readers of this document may wish to join the Bx community by subscribing to its mailing list and/or consulting the Bx wiki: see <http://bx-community.wikidot.com/start>. There is also a catalogue of examples of bx and multx, including some that were discussed at this Dagstuhl [3].

References

- 1 Bidirectional Transformations Wiki: *Bx Events*, <http://bx-community.wikidot.com/bx-events>, 2012.
- 2 Bidirectional Transformations Wiki: *Bx Steering Committee*, <http://bx-community.wikidot.com/bx-steering-committee>, 2012.
- 3 Bidirectional Transformations Wiki: *List of pages tagged with multxdagstuhl*, <http://bx-community.wikidot.com/system:page-tags/tag/multxdagstuhl>, 2018.

6.65 Algebraic Coding Theory for Networks, Storage, and Security

Organizers: Eimear Byrne, Martin Bossert, and Antonia Wachter-Zeh

Seminar No. 18511

Date: December 16–21, 2018 | Dagstuhl Seminar

Full report – DOI: 10.4230/DagRep.8.12.49

© Creative Commons BY 3.0 Unported license

© Eimear Byrne, Martin Bossert, and Antonia Wachter-Zeh



Participants: Gianira Nicoletta Alfarano, Iryna Andriyanova, Daniel Augot, Allison Beemer, Rawad Bitar, Jessalyn Bolkema, Martin Bossert, Eimear Byrne, Giuseppe Cotardo, Marwa Dammak, Salim El Rouayheb, Tuvi Etzion, Alexey Frolov, Philippe Gaborit, Heide Gluesing-Luerssen, Oliver Wilhelm Gnilke, Marcus Greferath, Daniel Heinlein, Lukas Holzbaur, Anna-Lena Horlemann-Trautmann, Sidharth Jaggi, Jörg Kliewer, Stanislav Kruglik, Margreta Kuijper, Sascha Kurz, Julien Lavauzelle, Hunter Lehmann, Andreas Lenz, Julia Lieb, Pierre Loidreau, Felice Manganiello, Carolyn Mayer, Sihem Mesnager, Alessandro Neri, Cornelia Ott, Mario Osvin Pavcevic, Sven Puchinger, Alberto Ravagnani, Julian Renner, Joachim Rosenthal, Ronny Roth, John Sheekey, Carmen Sippel, Emina Soljanin, Razane Tajeddine, Violetta Weger, Eitan Yaakobi, Yiwei Zhang, Jens Zumbärgel

Algebraic Coding Theory for Networks, Storage, and Security was the fourth in a series of seminars exploring applications of coding theory in modern communications theory (see also Dagstuhl Seminars 16321 (2016), 13351 (2013) and 11461 (2011)). The seminar brought together 50 mathematicians, engineers and computer scientists with expertise in coding theory, network coding, storage coding, cryptography and code-based security to participate in dissemination and collaboration within the seminar themes.

The main focus of this workshop was to explore novel results in coding theory for application in data storage management, cryptography and privacy. The impact of novel coding techniques across these domains was discussed and explored. Particular emphasis was placed on new applications of coding theory in public key cryptography, coding techniques for privacy in distributed storage and on practical schemes using coding theory for content delivery. These novel coding applications continue to have a significant impact on changing focus and broadening of coding theory fundamentals.

Overview talks were given by Philippe Gaborit (*Recent Results for Cryptography Based on Rank Metric*), Emina Soljanin, (*Service Rates of Codes*), Eitan Yaakobi, (*Private Proximity Retrieval*), Sacha Kurz, (*Multisets of Subspaces and Divisible Codes*), Heide Gluesing-Luerssen (*On Ferrers Diagram Codes*) and Salim El Rouayheb (*GASP Codes for Secure Distributed Matrix Multiplication*). In addition, several short talks were given by other participants based on current research interests with a view to stimulating collaboration. Presentations on system cybersecurity, private information retrieval, locally recoverable codes, adversarial channels and various aspects of rank metric codes were given. The remaining seminar time was allocated to discussion groups, including those in code-based cryptography, private computation, service rates of codes, algebraic geometry codes and adversarial channels. Aside from the working group

discussions, participants took the opportunity to engage in specific collaborations with co-authors.

We summarize some of the content of the working group discussions below. It has been well documented that redundancy is a basic requirement for stability of distributed data storage systems. Algebraic codes have been identified as having applications in providing efficiency in this domain far exceeding replication. Coding theory methods allow information retrieval minimizing disc access, storage size, local recoverability, data repair and data retrieval. Consequently, the area of storage coding has seen an exponential growth. An important aspect of user access in distributed storage is privacy of information retrieval so that users who are remotely accessing files can do so without storage servers knowing what they have accessed. Attempts to efficiently solve this problem come from coding theory.

An important application of secret sharing schemes is distributed storage of private data, where each party is a storage node and all parties wish to store a secret securely and reliably. Secret sharing is a fundamental cryptographic primitive and is used as a building block in numerous secure protocols. In our discussions we focussed on secret sharing schemes for the threshold access structure and on secret sharing with errors/attacks in a broader context. Fuzzy vaults and secret sharing over networks were discussed. A motivation for this area is for example biometric authentication in the presence of adversaries.

Another aspect of distributed storage is the service rate of codes. Emerging applications, such as distributed learning and fog computing, add yet another use for coding. In these applications, the goal is to maximize the number of users that can be simultaneously served by the system. One such service is simultaneous download of different jointly coded data blocks by many users competing for the system's resources. Here, coding affects the rates at which users can be served. The achievable service rate region is the set of request rates for each file that

can be supported by the system. A variety of approaches to open problems about service rate were discussed. In particular, we addressed the question of code constructions that serve all requests for fixed rate constraints on file and the problem of how to determine the achievable service rate region for certain families of codes.

Privacy and security present formidable challenges in our modern connected world. Public-key cryptography is the foundation of multi-party communication as well as for key exchange of symmetric cryptosystems. With the increasing likelihood of a capable quantum computer, post-quantum secure systems have recently turned into the research focus, especially for devices that are hard to update and have very long life cycles. Code-based cryptography provides post-quantum secure public-key systems.

The working group on code-based-cryptography discussed the importance of security reduction arguments and went through several examples of these in relation to coding theory in cryptography. This discussion was a great benefit to the participants, many of whom have expertise in coding theory and keen to broaden their understanding of cryptography. The group also focussed on McEliece-like systems based on quasi-cyclic moderate density parity-check (QC-MDPC) codes and on low-rank parity-check (LRPC) codes. Distinguisher attacks were discussed, as well as

possible modifications to the broken Gabidulin based cryptosystem.

Reliable communication across a channel in the presence of an adversary is a very general channel model that arises in many applications. Coding strategies for data transmission and authentication across the arbitrarily varying channel (where an adversary may alter the channel statistics) and for covert communication were discussed. A framework for linear systems under attack, such as the scenario where a restricted number of sensor measurements is vulnerable to adversarial attacks, was introduced and coding theoretic arguments used for attack detection and correction strategies.

There were about 20 PhD and postdoctoral researchers in attendance, who reported a very positive experience and satisfaction at being given the opportunity to explore new collaborations with more senior researchers and to get exposure to new problems in coding theory. All participants welcomed the time made available to them to take part in discussion groups and in more focussed collaborations. All were very pleased with the quality of the facilities and administrative support offered by staff at Schloss Dagstuhl, which made for a very productive meeting. Andreas Lenz and Rawad Bitar organised an afternoon excursion to Trier for the group. Giuseppe Cotardo collected and compiled data for the final published report.

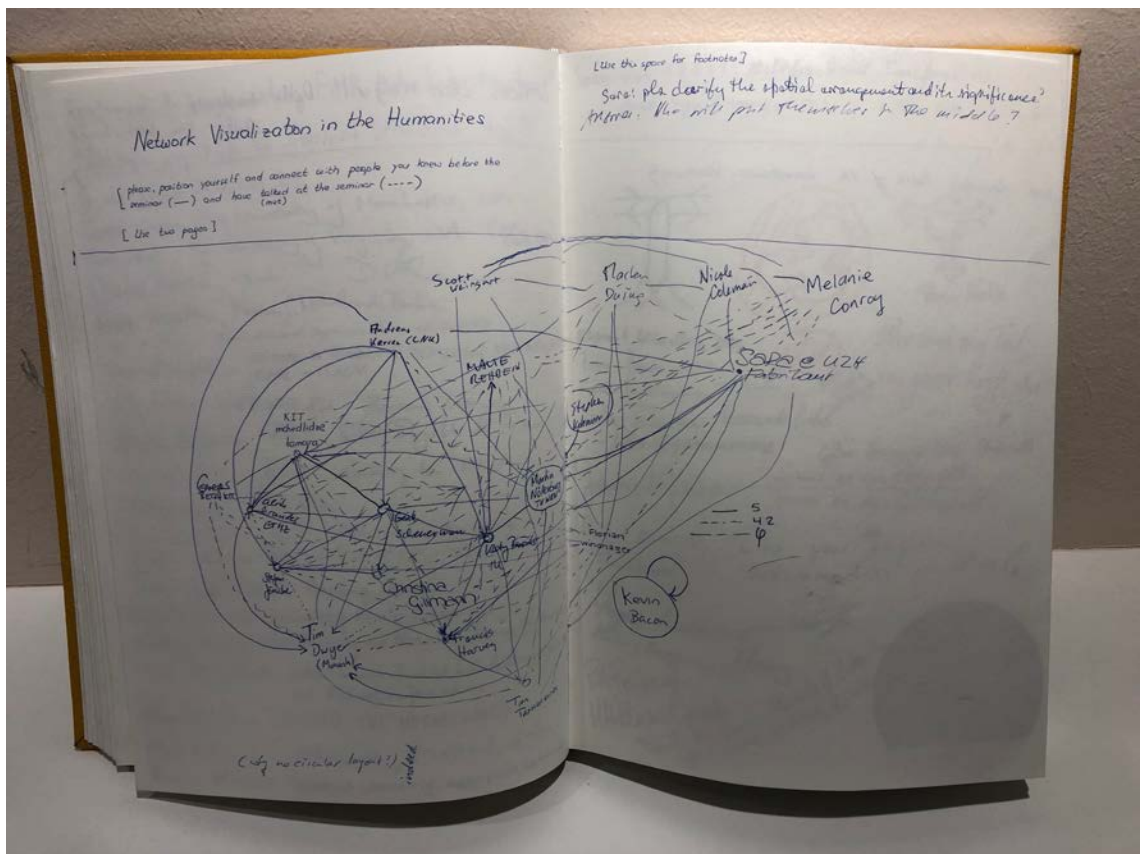


Fig. 6.20
 Dagstuhl Tradition.

7 **Öffentlichkeitsarbeit** *Public Relations and Outreach*

Pressemitteilungen und Medienarbeit

7.1

Press Releases and Media Work

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Vermittlung des Konzepts von Schloss Dagstuhl ist dabei ebenfalls ein Thema. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl⁴³ abrufbar.

Schloss Dagstuhl hat sich zur allgemeinen Anlaufstelle für Journalisten etabliert, die über bestimmte Informatikthemen, aber auch über Schloss Dagstuhl berichten möchten. Durch Unterstützung des Saarländischen Rundfunks steht Schloss Dagstuhl ein professionelles Reporterset zur Verfügung, welches Rundfunkjournalisten erlaubt, vor Ort mit Seminarteilnehmern Interviews in digitaler, verlustfreier Audioqualität zu führen.

Schloss Dagstuhl verbreitet Neuigkeiten rund um sein Programm auch über soziale Netzwerkdienste wie Twitter und LinkedIn. Über Twitter-Nutzer @dagstuhl werden Programmankündigungen, die Publikation von neuen Tagungsbänden aber auch andere relevante Neuigkeiten an aktuell ca. 1 700 Abonnenten verbreitet. Zunehmend nutzen aber auch Seminarteilnehmer den Dienst, um ihre Eindrücke vom Seminar mitzuteilen. Darüber hinaus werden über den Twitter-Nutzer @dblp_org Informationen über die Bibliographiedatenbank dblp an ca. 770 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ gepflegt (derzeit über 630 Mitglieder), mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Weiterhin werden dort interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben.

Regular press releases showcase and disseminate information about current computer science 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 website⁴³.

Schloss Dagstuhl has become a port of call for journalists seeking to report on specific computer science topics and/or on Schloss Dagstuhl itself. 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.

News on the program of Schloss Dagstuhl are also disseminated via social networks such as Twitter and LinkedIn. The Twitter handle @dagstuhl is used to disseminate program announcements, publication announcements, and other relevant news to about 1,700 followers, but is also increasingly used by Dagstuhl Seminar participants to share their impressions. Additionally, information about the dblp computer science bibliography is sent using the Twitter account @dblp_org, having about 770 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with more than 630 members), which supports the networking of participants in Dagstuhl Seminars. Additionally, interesting news about Schloss Dagstuhl are announced there.

Fortbildung

7.2

Educational Training

Lehrerfortbildung

Seit nunmehr über 25 Jahren engagiert sich Schloss Dagstuhl im schulischen Bereich durch die Organisation einer jährlichen Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Diese beiden Institute unterstützen die Fortbildung auch finanziell, indem sie die Kosten der Referenten tragen.

Jede Lehrerfortbildung dauert drei Tage; an jedem Tag werden in jeweils 3-stündigen Vorträgen zwei Informatikthemen vorgestellt. Die intensive Fortbildung richtet sich

Teacher training

Since more than 25 years, Schloss Dagstuhl hosts an annual teacher training workshop specifically designed for teachers of upper 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). These two institutes support the event also financially by assuming the costs of speakers.

Each workshop lasts three days; each day two computer science topics are presented in a three hour presentation each. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, since 2011 up to five teachers of other federal states can

⁴³ <https://www.dagstuhl.de/about-dagstuhl/press/>

zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch können seit 2011 bis zu fünf Lehrer aus anderen Bundesländern teilnehmen. Mehr Informationen zur Lehrerfortbildung 2018 gibt es auf der Webseite der Veranstaltung⁴⁴.

participate. Details on the workshop in 2018 are available at the event webpage⁴⁴.

„Dagstuhler Gespräche“

7.3

“Dagstuhler Gespräche”

Um die Türen des Schlosses etwas weiter für die Allgemeinheit und die Region zu öffnen, hat Schloss Dagstuhl zusammen mit der Stadt Wadern eine neue Veranstaltungsreihe ins Leben gerufen: die *Dagstuhler Gespräche*. Der interessierten Öffentlichkeit werden hier Themen aus dem breiten Spektrum der Informatik sowie ihre praktische Anwendung im Alltag oder in wirtschaftlichen Prozessen anschaulich in Form eines Impulsvortrages näher gebracht, um danach in einen gemeinsamen Dialog einzusteigen. An den Dagstuhler Gesprächen nehmen Entscheider und Gestalter aus Wirtschaft, Politik und der Informatik teil, aber auch Interessierte aus der Bevölkerung sind herzlich eingeladen.

Für den 18. Mai 2018 konnte der ehemalige Präsident der Gesellschaft für Informatik e.V. und frühere Leiter des Fraunhofer Instituts für Rechnerarchitektur und Softwaretechnik (FIRST), Prof. Dr. Stefan Jähnichen, als Vortragender gewonnen werden. Unter dem Titel „Digitalisierung – (k)ein Ende in Sicht?“ erklärte er aus seinem Erfahrungsschatz die Digitalisierung und hob spezielle Problemstellungen zu Sicherheit und Akzeptanz ihres Einsatzes hervor, die – wenn nicht zufriedenstellend gelöst – dem Trend der zunehmenden Digitalisierung entgegenstehen können. Speziell erörterte er, wie man sich in der digitalen Welt ausweisen kann, wie also digitale Identitäten gestaltet werden können, um einen vertrauensvollen Umgang mit Daten zu ermöglichen.

Der Vortrag fand regen Anklang und löste angeregte Diskussionen aus. Die Reihe wird im kommenden Jahr gewiss fortgesetzt werden.

In order to open its doors a bit further for the general public and the local region, Schloss Dagstuhl, together with the town of Wadern, initiated a new series of events: the *Dagstuhler Gespräche* (“Dagstuhl conversations”). The interested public will be introduced to a broad spectrum of topics from computer science, as well as to practical applications of those topics in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

The talk on May 18, 2018 was given by the former President of the Gesellschaft für Informatik e.V. (German Informatics Society) and former head of the Fraunhofer FIRST institute, Prof. Dr. Stefan Jähnichen. Under the title “Digitalisierung – (k)ein Ende in Sicht?” (Digitalization – (no) end in sight?), he explained the digital transformation from his rich experience and highlighted specific problems concerning security and acceptance, which could impede the growing trend of digitalization if no satisfying solutions are found. He discussed how one can prove one’s identity in the digital world and how digital identities can be designed such as to foster trust in handling of data.

The talk was well received, and the discussion was lively. The Dagstuhler Gespräche will certainly see a continuation in the next year.

⁴⁴ <https://www.dagstuhl.de/18503>

8

Einrichtungen

Facilities

Das Zentrum verfügt über drei Standorte; der Hauptstandort ist Schloss Dagstuhl in Wadern. Die Geschäftsstelle mit Sachbearbeitungsteam und wissenschaftlichen Mitarbeitern, die für die Dagstuhl-Seminare und Perspektiven-Workshops verantwortlich sind, befinden sich auf dem Campus der Universität des Saarlandes in Saarbrücken, während der Bibliographiedienst durch wissenschaftliche Mitarbeiter in Räumlichkeiten der Universität Trier betreut wird. Der Dagstuhl-Verlagsdienst befindet sich in Saarbrücken und Wadern.

The institution operates from three sites: the main site is Schloss Dagstuhl in Wadern. The administrative office and the scientific staff operating the Dagstuhl Seminars and Perspectives Workshops are located on the campus of Saarland University in Saarbrücken, while the scientific staff operating the Bibliographic Services are located in offices on the campus of the University of Trier. Dagstuhl Publishing is located in Saarbrücken and Wadern.

Hauptstandort in Wadern

8.1

Main Site in Wadern

Der Hauptstandort in Wadern umfasst das historische Schloss (gebaut um 1760) mit einem Anbau aus den 1970ern, einem 1993 fertiggestellten Erweiterungsbau, in dem sich Forschungsbibliothek, Hörsäle, Gästezimmer, Büros und Infrastruktur befinden, und ein 2012 fertiggestelltes Gästehaus mit Gästezimmern, einem Konferenzraum und Räumlichkeiten der Gebäudeverwaltung. Alle Einrichtungen in Wadern sind ganzjährig in Betrieb, abgesehen von je zwei Wochen im Sommer und Winter, die für größere Instandhaltungsarbeiten genutzt werden.

The main site in Wadern comprises the historic manor house (built around 1760) with an extension from the 1970s, a facility completed in 1993, which is housing a research library, lecture halls, guest rooms, offices and infrastructure, and a guest house completed in 2012 with guest rooms, a conference room, and garages for facility management. All facilities at Wadern are operated all year round except for two weeks each in summer and winter when larger maintenance tasks are scheduled.

Die Kapazitäten von Dienstleistungen und Räumlichkeiten zur Veranstaltung von Seminaren sind genau aufeinander abgestimmt: Das Zentrum hat 71 Gästezimmer, davon sind 18 Doppelzimmer, sodass insgesamt 89 Teilnehmer übernachten können. Bei Normalbetrieb finden parallel zwei Seminare mit jeweils 30 und 45 Teilnehmern statt, wobei jedem Seminar ein Hörsaal für 35 bzw. 60 Personen zur Verfügung steht. Obwohl so eine Gesamtsumme von 75 Teilnehmern entsteht, ist es nur selten notwendig, Seminargäste in Doppelzimmern oder einem nahegelegenen Hotel unterzubringen. Die Obergrenze von 71 Zimmern wird regelmäßig erreicht, weshalb es wohl kaum Möglichkeiten gibt, die Nutzung unserer Einrichtungen weiter auszubauen.

The capacities of services and facilities for hosting seminars at the main site are well coordinated: the site has 71 rooms, including 18 double rooms, for a total capacity of 89 participants staying overnight. During routine operation two seminars with nominally 30 and 45 participants are hosted in parallel, each using a lecture hall with 35 and 60 seats, respectively. Even though this sums up to 75 seminar participants it is rarely necessary to book seminar guests into double rooms or a nearby hotel. The maximum capacity of 71 rooms is reached regularly and hence there is hardly a way to increase utilization of facilities further.

■ Tagungsräume

Schloss Dagstuhl bietet drei Hörsäle für jeweils 25 bis 60 Personen. Alle Hörsäle sind mit einem Beamer, einem MS-Windows-Laptop und einer Audioanlage einschließlich Mikrofonen ausgestattet. Durch diese Technik werden Vorträge, Präsentationen und Live-Vorführungen optimal unterstützt. Mittels eines Presenters können Vortragende ihre vorbereiteten Materialien präsentieren, ohne zum Laptop oder Arbeitsplatz zurückkehren zu müssen.

■ Conference Facilities

Schloss Dagstuhl has three lecture halls with a seating capacity of 25 to 60 each. All lecture halls are equipped with a projector, an MS-Windows notebook, 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 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.

2017 wurde ein zweiter großer Hörsaal fertiggestellt. Dazu wurden der kleinste Hörsaal und ein benachbarter Computerraum zu einem neuen großen Saal zusammengelegt, um den heutigen Anforderungen bezüglich Raumangebot und technischer Ausstattung gerecht zu werden.

The construction works for a second large lecture hall has been finished in 2017. Schloss Dagstuhl's smallest lecture hall and an adjacent computer room have been merged into a new large lecture hall meeting current requirements, both in terms of size and technical equipment.

Neben den Hörsälen gibt es im Zentrum sechs Seminarräume. Davon sind zwei mit modernen Beamern ausgestattet, während in einem ein großes Plasmadisplay montiert ist. Fünf Beamer auf Rollwagen stehen zusätzlich zur flexiblen Benutzung in allen Räumen zur Verfügung.

In addition to the lecture halls, the center has six meeting rooms. Two are equipped with up-to-date projectors and one has a large plasma display on the wall. Five mobile projectors are available for use in all of the rooms.

Alle Hörsäle und andere Tagungsräume sind Tafeln und/oder Whiteboards ausgestattet.

Daneben gibt es über das ganze Zentrum verteilt weitere Räume, in denen Gäste sich in entspannter Atmosphäre treffen und diskutieren können. Insbesondere 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 in angenehmer Atmosphäre.

■ Dagstuhl's Küche

Die Mahlzeiten sind ein wichtiger Bestandteil des wissenschaftlichen Programms von Schloss Dagstuhl. Die Sitzordnung wird absichtlich stets zufällig gemischt, um eingefahrene Gruppen aufzuteilen und Gäste zu ermuntern, während ihres Aufenthalts möglichst viele verschiedene Kollegen kennenzulernen. Große Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten.

Dagstuhl's Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden jeden Tag von unseren Mitarbeitern der Küche und unseren Auszubildenden frisch zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dies steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber der Mehrheit der internationalen Gäste des Zentrums durchaus entgegen.

Sowohl die Zutaten als auch die Gerichte wechseln saisonal. An warmen Sommerabenden wird auf Anfrage auf der Terrasse vor dem Speisesaal gegrillt, unter anderem saarländische Schwenker, eine lokale Variante des Grillsteaks, die unter dauerndem Schwenken des Grillrostes zubereitet wird. In den kalten Monaten steht einmal wöchentlich ein schmackhafter Eintopf auf dem Speiseplan. Über das Jahr hinweg wird eine ausgewogene Mischung an regionalen und internationalen Spezialitäten aus neuen sowie bewährten und beliebten Rezepten angeboten. Im Allgemeinen sind die angebotenen Gerichte im Sommer etwas leichter und im Winter ein wenig schwerer. Die Küche arbeitet nach dem HACCP-Konzept (Hazard Analysis and Critical Points Concept) und hält sich an die Kennzeichnungspflicht von Allergenen, zu der alle lebensmittelverarbeitenden Betriebe verpflichtet sind. Des Weiteren achten wir auf deklarationsfreie Zusatz- und Konservierungsstoffe.

Alle Gäste, die aus medizinischen oder ethischen Gründen Einschränkungen bei der Speiseauswahl haben, können sich vor dem Seminar bei Schloss Dagstuhl melden. Unsere Küchenmitarbeiter erarbeiten gerne individuelle Lösungen für jeden Gast, soweit es irgend möglich ist. Gäste, die koscheres Essen benötigen, haben die Möglichkeit, mitgebrachte abgepackte Speisen selbst zu erhitzen.

Um unseren Gästen trotz eines begrenzten Budgets eine ausgewogene Qualität anbieten zu können, bietet unsere Küche ein Frühstücksbuffet, dienstags bis donnerstags abhängig von den personellen Kapazitäten ein Mittagsbuffet sowie ein Menü am Abend an. Montags und freitags wird aus logistischen Gründen auch am Mittag ein

All lecture halls and meeting rooms are equipped with blackboards and/or whiteboards.

The center also offers a variety of other spaces where guests can sit and work together in a relaxed atmosphere. Particularly in the evening, guests gravitate towards the wine cellar and upstairs café, two of the coziest places in the house and great places for continuing a productive discussion in a comfortable atmosphere.

■ Dagstuhl's Kitchen

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are mixed deliberately in order to break up cliques and encourage guests to talk to as many different people as possible during the course of their stay. Large tables in the dining hall promote collaborative interaction during meals.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy, and tasty meals. Everything is freshly prepared each day by the kitchen's staff and the apprentices in training. The focus is on lighter fare during the day in order to aid scientists' concentration, and on a warm meal in the evening, breaking with the German tradition of a cold evening meal while matching the internationality of the center's guests.

Both ingredients and dishes vary with the seasons. On warm summer evenings, guests are invited on demand to partake of grilled *Schwenker* (the local variant of barbecued steak) on the outdoor patio adjacent to the dining hall. During the cold winter months, warm soups appear on the menu weekly. In general, the kitchen tries to keep meals lighter in the summertime and heavier in the winter, offering a blend of regional and international dishes year-round that include some new recipes and many tried-and-true Dagstuhl favorites. The kitchen works in accordance with the HACCP Concept (Hazard Analysis and Critical Points Concept) and adheres to the mandatory labeling of allergens, which is required of all food processing establishments. Food additives and conservatives for which labeling is non-mandatory are also carefully monitored.

All guests with special dietary requirements due to ethical or health reasons can announce their needs previous to the events. Our kitchen staff will then work out individual solutions if at all possible. Guests who need kosher meals can heat up ready-to-eat meals for themselves.

To accomplish all of this within a reasonable budget, the center offers a buffet-style breakfast and a set evening meal served by the kitchen's friendly and dedicated staff. From Tuesday to Thursday the kitchen offers a buffet-style lunch depending on the staff capacities. Due to logistical reasons, a set meal is served at lunch on Mondays and Fridays. The large dining-hall, seating up to 80 persons, opens onto the castle garden and patio, and offers a relaxed, familiar atmosphere.

Small and late-morning breaks punctuate the daily routine. During the small coffee break during the morning hot drinks are served outside the lecture halls. During the longer coffee break in the afternoon, hot drinks together with freshly baked cake are served in the dining

Menü serviert. Unser Restaurant mit den großen Fenstern zum Garten des Hauptgebäudes bietet ca. 80 Personen Platz. Hier herrscht eine entspannte und fast familiäre Atmosphäre, was nicht zuletzt auf unsere freundlichen und engagierten Mitarbeiter zurückzuführen ist.

Kleine und große Pausen unterbrechen auf angenehme Weise die tägliche Routine und anstrengenden Diskussionen. In der kleinen Kaffeepause am Vormittag stehen vor den Vortragsräumen heiße Getränke auf einem Kaffeewagen bereit. In der großen Kaffeepause am Nachmittag wird den Gästen im Speiseraum neben heißen Getränken auch frisch gebackener Kuchen angeboten. Darüber hinaus gibt es im Gästehaus, der „alten“ Cafeteria und dem Weinkeller jeweils einen Kaffeevollautomaten zur Zubereitung von Kaffee, Kakao und Tee. Im Kiosk vor der Cafeteria können Gäste Snacks erwerben. Abends gibt es in der Cafeteria und im sogenannten Weinkeller einen Gruß aus der Küche, bestehend aus Brot und einer Käseauswahl.

■ Kinderbetreuung

Schloss Dagstuhl bietet Teilnehmern, die mit Kindern anreisen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser Service kann gegen ein geringes Entgelt im Voraus gebucht werden. Alternativ ist es Eltern auch möglich, eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Schloss Dagstuhl kommt für die Unterkunft und Verpflegung der Kinder auf. Wenn statt Inanspruchnahme der Kinderbetreuung von Schloss Dagstuhl eine Betreuungsperson mitreist, hat diese ebenfalls freien Aufenthalt.

Dagstuhls Angebot der Kinderbetreuung für Eltern wird weiterhin gut genutzt. Im Jahr 2018 wurden 17 Kinder durch eine Tagesmutter und 12 weitere durch Verwandte und 4 durch die Eltern selbst betreut. Insgesamt beherbergte Schloss Dagstuhl 33 Kinder von Teilnehmern an 19 Veranstaltungen während 18 Wochen.

■ Freizeit und Ambiente

Die Freizeitanlagen auf Schloss Dagstuhl wurden so gestaltet, dass sie auf unterschiedliche Art und Weise sowohl tagsüber als auch abends die Kommunikation zwischen den Seminarteilnehmern fördern. Die Mischung aus Arbeit und Freizeit in entspannter, familiärer Atmosphäre ist ein wichtiger Bestandteil des Dagstuhl-Konzepts. Gäste leben und arbeiten zusammen in einem Komplex aus drei Gebäuden, im Zentrum das historische Schloss, wo sie rund um die Uhr freien Zugang zu den zahlreichen Freizeiträumen und -anlagen haben. Musikalische Gäste können ihre Fertigkeiten im barocken Musiksaal zu Gehör bringen, wo ein Flügel und diverse andere Instrumente wie z. B. zwei Konzertgitarren zur Verfügung stehen. Unser Zentrum verfügt außerdem über eine Sauna, einen Billardtisch, Tischfußball, Mountainbikes, eine Dartscheibe, einen Freizeitraum mit Fitnessgeräten und Tischtennis sowie einen Außenbereich mit Volleyballnetz.

In addition, there are self-service bean-to-cup coffee machines in the guest house, at the “old” café, and in the wine cellar. Guests can buy small snacks at the kiosk in front of the cafe. Bread and cheese is served in the cafe and the wine cellar every night.

■ Childcare

Schloss Dagstuhl gladly offers to organize childcare with a certified nanny for participants who need to visit our center with young children. The service, which supports families and particularly women computer scientists, can be booked for a small recompense prior to the seminar.

Parents also have the option to bring along their own “nanny,” usually a spouse or relative. In the case of seminar participants the costs for room and board are absorbed by the center for the children. If an own nanny takes care for the children instead of Dagstuhl's childcare service, also the cost for the accompanying person for room and board are absorbed by Dagstuhl.

Guests make good use of Dagstuhl's childcare offer for parents. In 2018, Dagstuhl hosted 33 children, 17 of whom were cared for by a nanny on site, 12 by relatives, and 4 by their parents. Participants of 19 events in 18 weeks were thus able to attend although they were traveling with their children.

■ Leisure Facilities

Leisure facilities at Schloss Dagstuhl are designed to encourage and support communication among seminar participants in different settings throughout the day and evening. This work/life continuum within a relaxed, informal setting is an important part of the Dagstuhl concept. Guests live and work together in a complex of three buildings, the historical manor house (“Schloss”) in the middle, and enjoy full access to the center's many unique rooms and facilities around the clock. Musically talented guests are welcome to exercise their skills in the baroque music room on the upper floor of the historical main building, which features a grand piano and various other instruments, e.g., two concert guitars. Schloss Dagstuhl also has a full sauna, a pool table, table football facilities, mountain bikes, a dartboard, and a recreation room with gym equipment and table tennis as well as outdoor sports grounds featuring a volleyball net.

Geschäftsstelle in Saarbrücken

8.2

Dagstuhl Office at Saarbrücken

8

Die Geschäftsstelle in Saarbrücken befindet sich auf dem Campus der Universität des Saarlandes im Gebäude E11. Die Räumlichkeiten werden vom Sachbearbeitungsteam und von einem Teil des wissenschaftlichen Stabs genutzt. Es hat sich gezeigt, dass ein überraschend großer Teil unserer Tätigkeit enge Zusammenarbeit zwischen dem wissenschaftlichen Stab und dem Sachbearbeitungsteam erfordert. Darüber hinaus profitiert der wissenschaftliche Stab davon, dass sich auf dem Campus in Saarbrücken viele Informatiker in unmittelbarer Nähe befinden.

The Dagstuhl Office in Saarbrücken is located on the campus of Saarland University in building E11. The site houses some administrative staff and a part of the scientific staff. By now it is clear that a surprisingly big part of our work requires close interaction between scientific and administrative staff. The scientific staff benefit from the availability of a very large number of computer scientists on the Saarbrücken campus.

Dagstuhl an der Universität Trier

8.3

Dagstuhl at University of Trier

Die für die Bibliographiedatenbank dblp zuständigen Mitarbeiter haben ihren Standort an der Universität Trier. Die Ende 2010 zunächst auf Basis zweier Projekte gestartete Zusammenarbeit zwischen Schloss Dagstuhl und der Universität Trier wurde im November 2018 in eine offizielle und permanente Außenstelle von Schloss Dagstuhl auf dem Campus der Universität Trier überführt. Dabei profitiert das dblp-Team von der engen Zusammenarbeit mit der Abteilung Informatikwissenschaften und als externer Partner im Digital Research and Bibliographic Meta Data Lab des Center for Informatics Research and Technology (CIRT).

The scientific and editorial staff working on the *dblp computer science bibliography* is located at the Dagstuhl offices at the University of Trier. Initially based on a project-based cooperation between Schloss Dagstuhl and the University of Trier which was first established in 2010, in November 2018, an official and permanent Schloss Dagstuhl branch office has been established on the campus of the University of Trier. In Trier, the dblp team benefits from the close cooperation with the University's department of computer sciences, and as an external partner in the Center for Informatics Research and Technology (CIRT) lab for Digital Research and Bibliographic Meta Data.

9 **Zentrale Dienste** ***Central Services***

Schloss Dagstuhl verfügt über zwei zentrale Dienste: die IT-Abteilung und eine Forschungsbibliothek. Beide Einrichtungen befinden sich am Hauptstandort in Wadern.

Schloss Dagstuhl has two central services: The IT service and a research library, which are both located at the main site in Wadern.

Bibliothek

9.1

Research Library

Schloss Dagstuhl unterhält eine hervorragend bestückte Spezialbibliothek für Informatik, die an zahlreichen nationalen und überregionalen Bibliotheksverbänden teilnimmt. Die Bibliothek ist für Wissenschaftler vor Ort rund um die Uhr und für externe Wissenschaftler nach Absprache zugänglich. Der Bibliothekskatalog kann online durchsucht werden.

Für jedes Seminar wird eine individuelle Buchausstellung zusammengestellt, bestehend aus Büchern, die von Seminarteilnehmern verfasst oder herausgegeben wurden. Die anwesenden Autoren werden gleichzeitig gebeten, ihre Bücher zu signieren. Außerdem wird der Name eines jeden Seminarteilnehmers in der Online-Teilnehmerliste mit seinen oder ihren in der dblp-Literaturdatenbank erfassten Veröffentlichungen verlinkt. Diese beiden Maßnahmen ermöglichen den Seminarteilnehmern einfachen und schnellen Zugriff auf seminarrelevante Literatur.

Die Bibliothek verfügt über eine umfangreiche Sammlung an Büchern, Konferenzbänden und Zeitschriften:

- der Buchbestand wird durch das Seminarprogramm bestimmt. Bei Neuanschaffungen liegt der Fokus auf Büchern, die einen Bezug zu Dagstuhl-Seminaren oder Perspektiven-Workshops haben oder von Seminarorganisatoren oder -teilnehmern verfasst wurden. Außerdem erhält die Bibliothek zahlreiche Bücher als Spenden von Verlagen und Autoren. Aktuell verfügt die Bibliothek über etwa 34 000 Informatikbücher.
- Beiträge in Konferenzbänden verkörpern den wichtigsten Teil der Literatur in der Informatik. Die Bibliothek hat die kompletten ACM- und IEEE-Proceedings elektronisch abonniert; ältere Bände stehen auch in Druckform zur Verfügung. Die Verlagsgruppe SpringerNature spendet der Bibliothek alle Bände der Reihe Lecture Notes in Computer Science (LNCS) sowohl in Druckform als auch elektronisch. Die Bibliothek verfügt somit über Druckexemplare aller veröffentlichten Bände ab Band 1.
- Fachzeitschriften leisten einen wichtigen Beitrag zur langfristigen Dokumentation. Häufig werden in Zeitschriften erweiterte Fassungen von Ergebnissen veröffentlicht, die zuvor in Konferenzbänden publiziert wurden. Die Bibliothek bietet Zugriff auf über 1 000 elektronische Fachzeitschriften. Die meisten sind in Zeitschriftenpaketen enthalten, die in Zusammenarbeit mit deutschlandweiten Konsortien lizenziert sind, beispielsweise DFG-geförderte National- und Allianzlizenzen sowie von der Leibniz-Gesellschaft geförderte Konsortiallizenzen.
- Die Bibliothek ermöglicht den Online-Zugriff auf über 7 000 deutschlandweite und internationale Zeitungen und Magazine aus über 120 Ländern.

Schloss Dagstuhl maintains a very well equipped research library for computer science which is part of the national network of libraries. The library is permanently open for researchers on site and accessible upon request for outside users. The library catalogue can be searched online.

For each seminar, the library prepares a special book exhibition with books authored or edited by participants. The attendant authors are asked to autograph them. In the online list of participants, each participant is also linked to his or her publications as they are recorded in the dblp literature database. Together, these services provide quick access to relevant literature for seminar participants.

The library maintains a large collection of books, conference proceedings, and journals:

- The collection of books is guided by the seminar program. New textbooks relevant to Dagstuhl Seminars and Perspectives Workshops or written by seminar organizers and participants are prioritized when purchasing new volumes. In addition, the library receives numerous books as donations from publishers and authors. Currently, the library holds about 34,000 books on computer science.
- Papers in conference proceedings represent the most important literature in computer science. The library subscribes to all relevant ACM and IEEE conference proceedings electronically. Back volumes are still available in print. The publisher SpringerNature donates all volumes of its Lecture Notes in Computer Science series (LNCS) both as printed and electronic copies to the library. The library holds printed copies of all published volumes since LNCS volume 1.
- Journals in computer science are important for keeping long-term records. Journals often publish extended versions of results previously published at conferences. The library provides access to over 1,000 scientific electronic journals. Most of them are included in journals packages that are licensed in cooperation with national initiatives, e.g., nationwide DFG-funded national and alliance licenses and consortia licenses supported by the Leibniz Association.
- The library provides online access to more than 7,000 national and international newspapers and magazines from more than 120 countries.

■ Zusammenarbeit

Schloss Dagstuhl Fachbibliothek ist an zahlreichen Bibliotheksdatenbanken beteiligt. Der komplette Zeitschriftenbestand (ältere Ausgaben in Druckform und aktuelle Abonnements ausschließlich online) sind in der Zeitschriftendatenbank (ZDB) aufgeführt. Zusätzlich ist der Bestand an elektronischen Zeitschriften in der Elektronischen Zeitschriftenbibliothek (EZB) erfasst. Diese Datenbanken bilden die Grundlage für den deutschlandweiten und internationalen Leihverkehr der Bibliotheken und ermöglichen uns, unseren Forschungsgästen auch Literatur zur Verfügung zu stellen, die in unserem Bestand nicht vorhanden ist.

Darüber hinaus ist der aktuelle Buchbestand im Katalog des Südwestdeutschen Bibliotheksverbundes (SWB) aufgeführt und somit für alle wissenschaftlichen Bibliotheken durchsuchbar, z.B. über den Karlsruher Virtuellen Katalog. Die Bibliothek war auch Mitglied bei LITexpress, der virtuellen Bibliothek für Rheinland-Pfalz, das Saarland und die deutschsprachige Gemeinschaft in Belgien, ein Medienverleihservice für die Einwohner dieser Regionen. LITexpress wurde jedoch zum 31.12.2018 eingestellt. Außerdem besteht eine enge Zusammenarbeit zwischen Schloss Dagstuhl und der Saarländischen Universitäts- und Landesbibliothek (SULB), der Campusbibliothek für Informatik und Angewandte Mathematik an der Universität des Saarlandes sowie der Bibliothek des Leibniz-Instituts für Neue Materialien (INM), die sich alle in Saarbrücken befinden.

Schloss Dagstuhl Fachbibliothek ist institutionelles Mitglied des Deutschen Bibliotheksverbandes (DBV).

■ Spenden an die Bibliothek

Die Bibliothek von Schloss Dagstuhl profitiert von zahlreichen Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2018 Buchspenden von den Verlagen, die in Fig. 9.1 aufgeführt sind. Auch viele Seminarteilnehmer spenden der Bibliothek ihre Bücher. Autorenexemplare werden ebenso dankbar entgegengenommen. Insgesamt erhielt das Zentrum im Berichtszeitraum 661 Bände als Spenden von Verlagen und Seminarteilnehmern.

■ Collaboration

The research library of Schloss Dagstuhl participates in numerous library databases. The complete journal holdings (back volumes in print and current subscriptions online only) are listed in the German union catalogue of serials (Zeitschriftendatenbank, ZDB). In addition, the electronic journal holdings are recorded in the Electronic Journal Library (EZB). These databases are the basis on which national and international online lending libraries deliver copies of articles and allow us to procure non-existent literature for our research guests.

In addition, the current book inventory is listed in the catalogue of the Southwestern German Library Network (SWB) and hence searchable for all academic libraries, e.g., through the Karlsruhe Virtual Catalogue. The library was also 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 regions. However, LITexpress was discontinued on 31.12.2018. Furthermore, Schloss Dagstuhl closely cooperates with the Saarland University and State Library (SULB), the Campus Library for Computer Science and Mathematics at Saarland University, and the library of the Leibniz Institute for New Materials (INM), all based in Saarbrücken.

The Schloss Dagstuhl research library has an institutional membership in the German Library Association (DBV).

■ Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. In 2018, 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. The center received a total of 661 volumes during the year 2018 as donations from publishing houses and seminar participants.

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.

IT-Service

9.2

IT Service

Die IT-Abteilung bietet umfassenden Support für sämtliche internen Vorgänge an allen drei Standorten. Darüber hinaus betreut sie die IT-Infrastruktur und -Dienste und bietet Unterstützung für alle Gäste bei Dagstuhl-Veranstaltungen.

Der IT-Service umfasst u.a.:

- Internetzugang über Ethernet und WLAN in allen Räumen. Für den WLAN-Zugang bietet Schloss Dagstuhl persönliche Accounts an und ist auch an der *eduroam*-Initiative beteiligt (eine praktische Alternative für Gäste, die bereits einen *eduroam*-Account haben). Innerhalb sämtlicher Einrichtungen stellt Schloss Dagstuhl ein weitläufiges Netzwerk von Zugangspunkten zum Drahtlosnetzwerk zur Verfügung, das aktiv überwacht und regelmäßig erweitert wird. Die Verbindung zum (externen) Internet wird durch zwei redundante 100 Mbit/s-Leitungen sichergestellt, die durch den DFN e.V. (Deutsches Forschungsnetz) betrieben werden. Die Übertragungsrate der beiden Leitungen wurde im Juni 2018 auf jeweils 375 Mbit/s erhöht.
- Fahrbare ebenso wie fest montierte Präsentationsmöglichkeiten in den Tagungsräumen. In den größeren Tagungsräumen können Vortragende den vorhandenen oder den eigenen Laptop verwenden.
- Zugang zu Netzwerkfarbdruckern, einem Scanner und einem Kopierer.
- Zugang zu gemeinschaftlich genutzten Computern mit den Betriebssystemen Microsoft Windows, Apple Mac OS X und Linux.
- Technischen Support für Seminarteilnehmer und Mitarbeiter von Schloss Dagstuhl.

Der IT-Service verwaltet (virtuelle) Server für alle Abteilungen, z.B.

- einen Webserver, auf dem sich Schloss Dagstuhls Internetpräsenz befindet (<https://www.dagstuhl.de>), die Informationen für Teilnehmer, zum Seminarprogramm usw. enthält,
- einen Server, auf dem sich DROPS befindet, Schloss Dagstuhls Publikationsplattform (<http://drops.dagstuhl.de>),
- den dblp-Server (<https://dblp.dagstuhl.de> und <https://dblp.org>).

Darüber hinaus stellt der IT-Service Tools für das gemeinschaftliche Arbeitsumfelds zur Verfügung und hält sie in Stand, z.B. *Sihot* (eine Software zur Organisation von Gastdaten), MySQL-Datenbanken, ownCloud (ein Cloud-basiertes Speichersystem) und weitere.

The IT service provides comprehensive support for all internal operations at all three sites. Moreover, it provides IT infrastructure, services, and support for all guests of Dagstuhl events.

This service includes – among others – the following:

- Internet access via Ethernet and Wi-Fi throughout all rooms. For Wi-Fi access Schloss Dagstuhl offers personal accounts and also takes part in the *eduroam* service⁴⁵ (which is a comfortable option for guests with existing *eduroam* accounts). Within its facilities, Schloss Dagstuhl provides a generous network of professional-grade wireless network access points that is actively monitored and extended regularly. External internet access for Schloss Dagstuhl is provided through two redundant 100 Mbit/s connections that are managed by DFN e.V. (National Science Network). In June 2018, the capacity of both connections has been increased to 375 Mbit/s each.
- Mobile and stationary presentation facilities in meeting rooms. In large meeting rooms presenters can use either a provided laptop or their own.
- Access to network color printers, a scanner, and a copier.
- Access to shared computers with operating systems Microsoft Windows, Apple Mac OS X, and Linux.
- Technical support for both seminar participants and Dagstuhl staff.

The IT service manages (virtualized) servers for Schloss Dagstuhl's divisions, such as

- a web-server hosting Schloss Dagstuhl's web page at <https://www.dagstuhl.de>, providing information for participants, information about the seminar program, etc.,
- a server hosting DROPS at <http://drops.dagstuhl.de>, Schloss Dagstuhl's publishing platform,
- the dblp server at <https://dblp.dagstuhl.de> and at <https://dblp.org>.

Furthermore, for internal work procedures, the IT service provides and maintains tools for a collaborative work environment, such as *Sihot* (a software for organizing guest data), MySQL data bases, ownCloud (a cloud-based storage system), and several others.

⁴⁵ *eduroam* (education roaming) is a world-wide roaming access service developed for the international research and education community, see <https://www.eduroam.org>.

10 Kunst *Art*

Dagstuhl als Galerie

10.1

Dagstuhl as Art Gallery

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des Flurs sowie die hervorragende Ausleuchtung mit starken Kontrasten zwischen Tag und Nacht bieten den Künstlern sehr gute Möglichkeiten, ihre Werke darzustellen. Die Kunstwerke an den Wänden des schmalen Gangs durchbrechen die Nüchternheit des Neubaus in anregender und angenehmer Weise. Die wechselnden Ausstellungen bieten einen erfrischenden und dynamischen Kontrast zu der ständigen Kunstsammlung von Schloss Dagstuhl.

Prof. Reinhard Wilhelm, ehemaliger wissenschaftlicher Direktor des Zentrums, fungierte nach seinem Eintritt in den Ruhestand im April 2014 weiterhin als Betreuer der Ausstellungsaktivitäten von Schloss Dagstuhl. Das Zentrum veranstaltet jährlich etwa drei bis vier Kunstausstellungen für jeweils zwei bis drei Monate.

Waren es bisher Künstler und einzelne Sammler, die ihre Werke ausstellten, so kam seit 2016 durch die Zusammenarbeit zwischen der Saarland-Sporttoto GmbH (kurz Saartoto), der Hochschule der Bildenden Künste Saar (kurz HBKsaar) und Schloss Dagstuhl die Sammlung von Saartoto als Reservoir für eine Ausstellungsserie hinzu. Als bedeutender Förderer von Künstlern besitzt Saartoto einen großen Bestand an Kunstwerken. Im Rahmen der Zusammenarbeit wird diese Kunstsammlung durch die HBKsaar erfasst und dokumentiert. Gleichzeitig wurden und sollen auch in Zukunft aus dem Saartoto-Fundus Ausstellungen für Schloss Dagstuhl zusammengestellt werden. Dabei werden die Kunstwerke aktuellen Werken von Künstlern der HBKsaar und aus der Großregion Saar-Lor-Lux gegenübergestellt. Die Galerie MediArt aus Luxemburg unterstützte das Projekt durch die Leihgabe von Bildern der Künstler aus der Großregion. Schloss Dagstuhl möchte an dieser Stelle allen beteiligten Personen danken, namentlich insbesondere Michael Burkert, Peter Jacoby und Josef Gros (Saartoto); Matthias Winzen und Nadine Brettar (HBKsaar); Paul Bertemes (MediArt); sowie Reinhard Wilhelm und Angelika Mueller-von Brochowski (Schloss Dagstuhl).

Die fünf Ausstellungen (siehe Fig. 10.1), die im Jahr 2018 stattfanden, sind nachfolgend beschrieben. Die jeweils aktuellen Ausstellungen sind nach Anmeldung auch für die interessierte Öffentlichkeit zugänglich.

■ »KONKRETE Funde«

Konkrete Kunst aus den Saartoto-Beständen bildet die konzeptuelle Basis der Ausstellung. Das Aufspüren konkreter Bildsprachen und Herangehensweisen in Positionen junger Kunstschafter der HBKsaar beleuchtet die universelle Wirksamkeit und zeitlose Gültigkeit und zeigt zugleich neue Spielarten dieses fundamentalen Gestaltungsrepertoires.

Art exhibitions are regularly organized in the so-called cloister of the new building. The spacious surroundings, excellent lighting, and dramatic day-to-night contrast offer artists a unique exhibition space. Arranged along the corridor walls, the artworks offset the otherwise ascetic nature of the new building. These temporary exhibits offer a fresh and dynamic counterpoint to the center's permanent collection, which can be found scattered throughout the three buildings.

Prof. Reinhard Wilhelm has continued to supervise the Schloss Dagstuhl art exhibitions following his retirement as the center's Scientific Director in April 2014. The center holds approximately three to four art exhibits per year, with each exhibit generally running for two to three months.

Until now, the exhibitions were organized by artists and individual collectors. The year 2016, however, saw the establishment of a cooperation between Saarland-Sporttoto GmbH (Saartoto for short), Hochschule für Bildende Künste Saar (university of art and design; HBKsaar for short), and Schloss Dagstuhl, which makes Saartoto's collection accessible to Schloss Dagstuhl for a series of exhibitions. Being a major art sponsor, Saartoto is in possession of a substantial art collection. In the context of this collaboration, HBKsaar takes stock of and documents Saartoto's art collection. At the same time, there were, and will be, exhibitions at Schloss Dagstuhl where Saartoto artworks are contrasted with recent works by HBKsaar artists and artists from the greater region Saar-Lor-Lux (Saarland, Lorraine, and Luxembourg). The Luxembourg-based art gallery MediArt supported the project by loaning several paintings by artists from the greater region. Schloss Dagstuhl would like to thank everyone involved, especially Michael Burkert, Peter Jacoby, and Josef Gros (Saartoto); Matthias Winzen and Nadine Brettar (HBKsaar); Paul Bertemes (MediArt); as well as Reinhard Wilhelm and Angelika Mueller-von Brochowski (Schloss Dagstuhl).

The five exhibitions (cf. Fig. 10.1) hosted by Schloss Dagstuhl in 2018 are described below. Current exhibitions are open to the interested public upon request.

■ »KONKRETE Funde«⁴⁶

Concrete art from the Saartoto collections forms the conceptual basis of the exhibition. The search for concrete visual languages and approaches in the positions of young artists of the HBKsaar illuminates the universal effectiveness and timeless validity and at the same time shows new varieties of this fundamental design repertoire.

⁴⁶ engl. Concrete finds

Linienzüge und Farbüberlagerungen, elementares Formenspiel, poetische und typografische Textvariationen sowie Serialität und Transformation bilden die Schnittmenge zwischen etablierter freier und junger angewandter Kunst.

Konstellationen des Begründers der konkreten Poesie Eugen Gomringer, konkrete Fotografie von Kilian Breier, Linienmodulationen in Weiß von Leo Erb, Farbnuancierungen der HBK-Absolventin Julia Lutz und noch mehr Konkretes aus den Archivbeständen trifft auf Designexperimente mit Wort, Text und Schriftarten und auf das Wiederentdecken elementarer Formen in fotografischen Architektur- und Raumausschnitten.

Klarheit, die Tendenz zur Vereinfachung und die Verlangsamung der Lesegeschwindigkeit führen in der konkreten Kunst zum Schärfen der Sinne und zum Verfeinern der Wahrnehmung. Nicht nur um die Alltagstauglichkeit und Lesbarkeit ihrer Designkonzepte zu steigern, greifen junge Gestaltende diesen visuellen Baukasten auf, sondern sie wecken im innovativen Spiel mit dessen vielfältigen Möglichkeiten das Bewusstsein für unsere Lebensumwelt, unser Konsumverhalten und aktuelle Werbestrategien.

Die kombinierte Präsentation von Kunst und Design lässt die zeitlose Relevanz des Konkretes als Ruhe- und Orientierungspol innerhalb der visuellen Fülle heutiger Lebens- und Konsumwelten anklingen.

■ »Schülerkunscht«

Am 26.04.2018, fand im Rahmen einer Vernissage, die Eröffnung der Kunstausstellung „Schülerkunscht“ im Schloss Dagstuhl statt. Das Besondere, die Künstler, waren dieses Mal Schüler des Hochwald-Gymnasiums. Normalerweise werden in der namhaften Galerie des Schloss Dagstuhls nur renommierte Künstler ausgestellt.

Umso grösser war die Freude der Schule und der SchülerInnen in diesem internationalen Rahmen, Schülerarbeiten präsentieren zu können. Die Arbeiten entstanden im Unterricht unter der Aufsicht des jeweiligen Kunstlehrers.

Lines and color overlays, elementary playing with shapes, poetic and typographic text variations as well as serial imagery and transformation form the intersection between established fine art and young applied art.

Constellations of the founder of concrete poetry Eugen Gomringer, concrete photography by Kilian Breier, line modulations in white by Leo Erb, color nuances by HBK graduate Julia Lutz, and even more concrete art works from the archive collections meet design experiments with words, text, and fonts, and the rediscovery of elementary forms in photographic architectural and spatial segments.

Clarity, the tendency towards simplification and slowing down the reading speed lead in concrete art to sharpening the senses and refining perception. Not only in order to increase the suitability for everyday use and readability of their design concepts, young designers take up this visual construction kit, but also, in an innovative play with its diverse capabilities, they awaken awareness for our living environment, our consumer habits and current marketing strategies.

The combined presentation of art and design evokes the timeless relevance of the concrete art as a resting and orientation pole within the visual abundance of today's lifestyles and consumerism.

■ »Schülerkunscht«⁴⁷

On April 26, 2018, the opening of the art exhibition "Schülerkunscht" took place at Schloss Dagstuhl during a vernissage. What was special was that, this time, the artists were pupils of the Hochwald-Gymnasium. Normally, only the works of renowned artists are exhibited in the renowned gallery of Schloss Dagstuhl.

All the greater was the joy of the school and the pupils to be able to present their work in this international forum. The works were created in the classroom under the supervision of the respective art teachers. So now

⁴⁷ engl. Students Art

»KONKRETE Funde« Works from the art collection of Saartoto, of the artists: Eugen Gomringer, Kilian Breier, Leo Erb, and Julia Lutz curated by Nadine Brettar February 19 to April 20, 2018
»Schülerkunscht« Works from the students at Hochwald-Gymnasium April 26 to June 1, 2018
»Let There Be More Light« Works from Gerd Borkelmann curated by Beate Kolodziej, M.A. June 7 to July 20, 2018
»a f f i n – Bilder von Claudia Vogel und Dirk Rausch « Works from Claudia Vogel and Dirk Rausch August 13 to October 5, 2018
»Isabell Gawron und Michael Mahren« Works from Isabell Gawron and Michael Mahren October 15 to December 14, 2018

Fig. 10.1
Art exhibitions in 2018.

So hängen nun viele interessante, völlig unterschiedliche Arbeiten, die sich zu einem harmonischen Ganzen zusammen fügen.

Gezeigt werden Themen des Lehrplans, die unterschiedlich umgesetzt sind. Ausgesucht wurden die Arbeiten nach künstlerischen Aspekten z. B. Raum und Fläche, gesellschaftskritische Porträts bis hin zur abstrakten Landschaften. Ein besonderes „Schmankerl“ bietet die Reihe „Comic und Pop-Art“, die schon großen Anklang fand bei den Besuchern.

■ »Let There Be More Light «

Im Werk des Malers Gerd Borkelmann bestimmt seit einiger Zeit die Farbe Grau seine Leinwände und Arbeiten auf Papier. In vielfältigen Nuancen verwendet er die unterschiedlichen Mischverhältnisse von Weiß und Schwarz. Es entstehen ebenso sanfte wie kraftvolle Farbkompositionen, bei denen mal der Bildträger durchschimmert, mal die Fläche kompakt verschlossen wird. Es ist Malerei, in die man sich vollends vertiefen kann.

■ »a f f i n – Bilder von Claudia Vogel und Dirk Rausch «

Die anschauliche Thematisierung der Korrelation von Form und Farbe innerhalb des bildkompositionellen Gefüges ist eines der zentralen künstlerischen Anliegen von Dirk Rausch. Dabei arbeitet er mit zunächst schlicht anmutenden Elementen wie balkenartigen Formationen, die sich in ihrer verschiedenfarbigen Erscheinung so überlagern, dass nuancenreiche Farbdurchdringungen entstehen. Die Positionierung der formalen Aspekte innerhalb des Bildfeldes zeugt von einer ebenso experimentellen wie auch genau durchdachten kompositionellen Intention, bei der Bildfläche und Bildformat, ebenso wie der lichte Bildgrund, als originäre Bestandteile der gesamten Gestaltung gelten. Die Arbeiten von Dirk Rausch changieren zwischen Disziplinierung der sehr überlegt eingesetzten Form und bewusster Zufälligkeit im Gestaltungsprozess, sodass, trotz der minimalistischen Tendenz, in der Interaktion von Form, Farbe und Fläche lebendige Bildsituationen entstehen, die Farb-Form-Kompositionen als sinnliches Erlebnis von hoher ästhetischer Qualität in besonderer Weise erfahrbar machen.

Die Künstlerin Claudia Vogel stellt sich in ihren Werkstücken die Frage nach der Arbeit mit bzw. am Bild. So widmet sie sich innerhalb eines Werkzyklus intensiv dem Thema der Bildbegrenzung bzw. der Sichtbarmachung der das Bild konstituierenden Elemente, beispielsweise der Rahmenkonstruktion. Zudem zeugen die Arbeiten von einer starken Konzentration auf die Farbigkeit und auf eine sorgfältige Auswahl der materiellen Bildmittel, welche den Bildern gleichermaßen ein Höchstmaß an Stofflichkeit und Plastizität verleihen. Insbesondere die Ergründung der Bildtiefe, welche die Künstlerin in ihren Werken durch die sukzessive Erfassung verschiedener Bildebenen formuliert, ist kennzeichnend für die Arbeiten.

Die Künstlerin arbeitet parallel an verschiedenen Bildlösungen. Trotz aller Gegensätzlichkeit und unterschiedlichster Bildaussagen ist den verschiedenen Werkgruppen

many interesting, completely different works are on display, which form a harmonious whole.

Topics of the curriculum are shown, which are implemented differently. The works were selected according to artistic aspects, e.g. space and surface, and from socially critical portraits to the abstract landscapes. This "Comic and Pop Art" series is a special treat that has already resonated very well with visitors.

■ »Let There Be More Light «

In the work of the painter Gerd Borkelmann, the color grey has dominated his canvases and works on paper for some time now. In a variety of nuances, he uses the different mixtures of white and black. Gentle as well as powerful color compositions are created, in which sometimes the canvas shimmers through, at other times the surface is compactly covered. It is painting in which one can fully immerse.

■ »a f f i n – Paintings from Claudia Vogel and Dirk Rausch «

The vivid thematization of the correlation of form and color within the compositional structure of the picture is one of Dirk Rausch's central artistic objectives. He works with seemingly simple elements such as beam-like formations, which overlap in their different colors to create nuanced color penetrations. The positioning of the formal aspects within the image field attests an experimental yet precisely thought-out compositional intention, in which the image surface and image format, as well as the light background, are regarded as original components of the overall design. Dirk Rausch's works alternate between disciplining the very deliberately used form and conscious randomness in the design process, so that, despite the minimalist tendency, lively pictorial situations arise in the interaction of form, color and surface, which make color-form compositions a sensual experience of high aesthetic quality that can be experienced in a special way.

The artist Claudia Vogel explores the question of working with or on the picture in her works. Within a work cycle, she devotes herself intensively to the subject of image limitation or the visualization of the elements constituting the image, for example the frame construction. In addition, the works bear witness to a strong concentration on color and a careful selection of the physical pictorial means, which lend the pictures a maximum of materiality and plasticity. In particular, the exploration of the depth of the image, which the artist formulates in her works by successively capturing different pictorial levels, is characteristic of her art.

The artist works in parallel on various visual techniques. In spite of all their contradictions and the most diverse visual statements, the various groups of works share a close intellectual connection: the works testify to Claudia Vogel's intense interest in colors and structures, formulated between calculated structurality and free visualization.

ein enger gedanklicher Zusammenhang gemein: Die Arbeiten zeugen vom intensiven Interesse Claudia Vogels an Farben und Strukturen, formuliert zwischen kalkulierter Strukturhaftigkeit und freier Bildwerdung.

■ »Isabell Gawron und Michael Mahren«

Intuitiv, spontan und mit surrealer Inspiration bewegt sich Isabell Gawron auf dem Gebiet der Abstraktion. Laufspuren, transparente Malschichten, die das Unterliegende nicht ganz verbergen sowie das Aufkratzen der oberen Farbebene und sich mehrfach überlagernde Schraffuren mit dem Grafitstift weisen auf den Malprozess an sich hin. Völlig losgelöst von Vorlage und Gegenstand lässt Sie Freiraum für unterschiedliche Interpretationen. Formlosigkeit und die Spontaneität in der künstlerischen Produktion sind die Merkmale ihrer Bilder.

Unmittelbar vor Ort entstanden die Zeichnungen Michael Mahrens (1948 – 2018), die fast notizenartig aufgenommen wurden. Während bei den Arbeiten nach Rhythmus und Bewegung der Impuls von außen wesentlich war, waren es bei den realistischen Zeichnungen vor Ort der spontane Duktus und der eigene Impuls, der die Zeichnung belebte.

■ »Isabell Gawron and Michael Mahren«

Intuitive, spontaneous, and with surreal inspiration, Isabell Gawron navigates the field of abstraction. Traces of running, transparent layers of paint that do not completely hide the surface beneath as well as the scratching of the upper layer of paint and multiple overlapping hatchings with the graphite pen reveal the painting process in itself. Completely detached from the model and the object, she leaves room for different interpretations. Formlessness and spontaneity in artistic production are the characteristics of her paintings.

The sketches by Michael Mahren (1948 – 2018) were made in situ and were drawn almost like notes. While in the works based on rhythm and movement the impulse from outside was essential, in the realistic drawings made in situ it was the spontaneous flow and one's own impulse that animated the drawing.

Kunstankauf durch Spenden

10.2

Art Sponsorship and Donations

Das Internetangebot von Schloss Dagstuhl enthält eine Seite, die es Teilnehmern, Einzelpersonen und Gruppen ermöglicht, Kunst für Dagstuhl zu stiften. Die Kunstobjekte werden über das Internet angeboten, dabei wird der Preis in kostengünstige Anteile aufgeteilt. Sobald alle Anteile eines Bilds gezeichnet sind, werden die Teilnehmer aufgefordert, den Gegenwert der bestellten Anteile als Spende einzuzahlen, wodurch dann das Objekt angekauft werden kann. Die Stifter werden sowohl in der virtuellen Internet-Galerie von Schloss Dagstuhl als auch an dem realen Objekt genannt. Dadurch ist es Schloss Dagstuhl möglich, Werke von Künstlern, die im Zentrum ausgestellt haben, anzukaufen und permanent auszustellen.

Im Jahr 2018 erhielt Schloss Dagstuhl insgesamt 257 € von verschiedenen Spendern. Wir möchten diese Stelle nutzen, allen Spendern, die 2018 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Nähere Informationen und aktuelle Neuigkeiten finden sich auf der [Kunst-Webseite](https://www.dagstuhl.de/art/)⁴⁸ von Dagstuhl.

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 online and donations are made by acquiring shares at affordable prices. Donors pay the value of their pledged shares as soon as a piece is fully subscribed for, thus allowing it to be purchased. Donors' names appear in Dagstuhl's online art gallery and also next to the art items themselves. In this way, Schloss Dagstuhl is able to purchase works of art from those who exhibit at the center, and add these works to its permanent art exhibition.

In 2018, Schloss Dagstuhl received a total of 257 € from various donors. We would like to thank all donors who contributed to Dagstuhl's art collection in 2018.

For further information and current news about Dagstuhl's art program, please visit Dagstuhl's art webpage⁴⁸.

Dagstuhls permanente Kunstaussstellung

10.3

Dagstuhl's Permanent Art Exhibition

Die von Gästen immer wieder positiv hervorgehobene Kunstsammlung geht auf den Gründungsdirektor Professor Wilhelm zurück. Seine Idee war es, den 1995

The art collection, continually praised by guests, was initiated by Founding Director Professor Wilhelm. It was his idea to use works of art in order to enliven the

⁴⁸ <https://www.dagstuhl.de/art/>

neueröffneten Speisesaal und den etwa ein Jahr älteren Neubau, durch Kunstwerke zu beleben. Dazu startete er die oben beschriebenen Kunstausstellungen. Unter Mitwirkung der Künstler wird aus jeder Ausstellung ein Werk ausgewählt, für das dann Spender gesucht werden. In den letzten 25 Jahren kamen so ungefähr 180 Kunstwerke zusammen. Auch durch diese Initiative angeregt und verstärkt erhielt Dagstuhl in den vergangenen Jahren weitere Spenden von Künstlern und Mäzenen. Die Arbeiten kommen in den Räumen des Zentrums in Wadern sowie in der Geschäftsstelle in Saarbrücken sehr gut zur Geltung.

New Building as well as the dining room opened in 1994 and 1995, respectively. To this end, Professor Wilhelm launched the exhibitions described above. Assisted by the artists, one picture from each exhibition was chosen and donors were drummed up. Thus, approximately 180 works of art could be acquired over the last 25 years. Additionally, this initiative has increasingly encouraged artists and patrons to make donations. All of the pictures adorn the rooms of Schloss Dagstuhl in Wadern as well as the Dagstuhl Office in Saarbrücken.

11

Struktur der Gesellschaft *Structure of the Company*

Gründung und Gesellschafter

11.1

Formation and Shareholders

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit elf Gesellschaftern (siehe Fig. 11.1) organisiert. Dies sind die vier Gesellschafter, die Schloss Dagstuhl gegründet haben, nämlich die Gesellschaft für Informatik e. V. (GI), die Universität des Saarlandes, die Technische Universität Kaiserslautern und das Karlsruher Institut für Technologie (KIT). Als vier weitere Gesellschafter wurden 1994 die Technische Universität Darmstadt, die Johann Wolfgang Goethe-Universität Frankfurt am Main, die Universität Stuttgart und die Universität Trier aufgenommen. Drei international renommierte Forschungsinstitute, das Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), das Centrum Wiskunde & Informatica (CWI, Niederlande) und die Max-Planck-Gesellschaft (MPG, Deutschland) wurden 2005/2006 als weitere Gesellschafter aufgenommen.

Aufgrund eines Beschlusses der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (heute Gemeinsame Wissenschaftskonferenz) wurde das Zentrum mit Wirkung zum 1. Januar 2006 als Serviceeinrichtung für die Forschung in die gemeinsame Forschungsförderung von Bund und Ländern aufgenommen. Es ist seit 2005 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.

Schloss Dagstuhl wurde im Juli 2009 erstmals durch die Leibniz-Gemeinschaft evaluiert. Die Stellungnahme der Evaluierungs-Kommission vom März 2010 war sehr positiv: Schloss Dagstuhl widme sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen. Schloss Dagstuhl wurde 2016 erneut mit hervorragendem Ergebnis evaluiert. In der Stellungnahme des Senats der Leibniz-Gemeinschaft wurde das Veranstaltungsprogramm und die Beteiligung an der Literaturdatenbank dblp als „exzellent“ bewertet, während der Bereich Open Access (Publishing) als „sehr gut“ bewertet wurde.

Schloss Dagstuhl is operated as a non-profit organization by eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e. V.⁴⁹ (GI), the Universität des Saarlandes, the Technische Universität Kaiserslautern, and the Karlsruher Institut für Technologie (KIT). In 1994, the organization was extended to include four new associates: the Technische Universität Darmstadt, the Johann Wolfgang Goethe-Universität Frankfurt am Main, the Universität Stuttgart and the Universität Trier. Finally, in 2005 and 2006, three internationally renowned research institutes joined the association: the Institut National de Recherche en Informatique et en Automatique (INRIA, France), the Centrum Wiskunde & Informatica (CWI, Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung⁵⁰ (today Joint Science Conference) the center has been classified as a research service institution for joint funding by the German federal and state governments since January 2006. Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association and changed its name accordingly from “Internationales Begegnungs- und Forschungszentrum für Informatik”⁵¹ to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”⁵² in 2008.

In July 2009, Schloss Dagstuhl was evaluated for the first time by the Leibniz Association. The March 2010 findings of the evaluation commission were very positive, and established that the center has shown outstanding commitment to its designated task of supporting the international computer science research community by providing a seminar center for academic events. In 2016, Schloss Dagstuhl has been evaluated again, with excellent results. In the Leibniz Association Senate report, the seminar program and the cooperation with the computer science bibliography dblp were rated as “excellent” whereas the Open Access Publishing was rated “very good.”

Organe der Gesellschaft

11.2

Dagstuhl Organs

Die drei Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, die stellvertretend für die Gesellschaft als juristische Person handeln, sind die folgenden:

- Gesellschafterversammlung
- Aufsichtsrat
- Geschäftsführung

Details zu den Organen sind den folgenden Abschnitten zu entnehmen.

The three organs of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, which act for the company as a legal entity, are the following:

- Shareholders' Meeting
- Supervisory Board
- Management

Detailed information is given in the sections below.

⁴⁹ engl.: German Informatics Society

⁵⁰ engl.: Federal/State Government Commission for Educational Planning and Research Promotion

⁵¹ engl.: International Conference and Research Center for Computer Science

⁵² engl.: Schloss Dagstuhl – Leibniz Center for Informatics

■ Die Gesellschafterversammlung

Die Gesellschafter beschließen über alle Änderungen an der Gesellschaft, insbesondere über die Aufnahme weiterer Gesellschafter, über die Änderung des Gesellschaftsvertrags und über ihre Auflösung. Die Gesellschafter bestätigen unter anderem auch die von Gesellschaftern neu entsandten Mitglieder in den Aufsichtsrat sowie die Berufung und Abberufung der Geschäftsführer. Derzeit haben anteilig nach der Höhe der Geschäftsanteile alle Gesellschafter die gleiche Anzahl von Stimmen, außer der Gesellschaft für Informatik, die die dreifache Anzahl besitzt. Beschlüsse werden entweder in der mindestens einmal jährlichen stattfindenden Gesellschafterversammlung gefasst oder durch schriftliche Stimmabgabe.

■ 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 (siehe Fig. 11.2) setzen sich aus vier Repräsentanten der Gesellschaft für Informatik, je einem Vertreter der drei Gründungsuniversitäten, zwei Vertretern der später hinzugekommenen vier Universitäten und je einem Vertreter des Bundes und der beiden Bundesländer Saarland und Rheinland-Pfalz, in denen Schloss Dagstuhl formal seinen Sitz hat, zusammen. Die reguläre Amtszeit der Aufsichtsratsmitglieder beträgt mindestens vier volle, abgeschlossene Geschäftsjahre und endet mit der Entlastung für das vierte Geschäftsjahr. Die Vertreter der Universitäten in Darmstadt und Stuttgart wechseln im Allgemeinen Amtszeit für Amtszeit mit denen der Universitäten in Frankfurt und Trier ab.

Der Aufsichtsrat entscheidet über die Berufung und Abberufung der Geschäftsführer sowie der Mitglieder des Wissenschaftlichen Direktoriums, des Wissenschaftlichen Beirates und des Kuratoriums. Alle Beschlüsse, die die Finanzen oder das Vermögen der Firma betreffen, benötigen seine Zustimmung. Beschlüsse von forschungspolitischer Bedeutung und Beschlüsse mit erheblichen finanziellen Auswirkungen können nicht gegen die Stimmen der Vertreter des Bundes und der beiden Sitzländer gefasst werden. Der Aufsichtsrat entscheidet zudem über die Erteilung einer Prokura.

■ Die Geschäftsführung

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH hat zwei Geschäftsführer (siehe Fig. 11.3), die gemeinsam die Gesellschaft vertreten. Die Geschäftsführung besteht aus dem *Wissenschaftlichen Direktor* und dem *Technisch-administrativen Geschäftsführer*.

Der Wissenschaftliche Direktor ist verantwortlich für die wissenschaftlich-fachliche Zielsetzung und die Programmgestaltung, und ist zudem Mitglied und Vorsitzender des Wissenschaftlichen Direktoriums. Seit Mai 2014

■ Shareholders' Meeting

All changes to the company, in particular the inclusion of new associates, the revision of the Shareholders' agreement, and the dissolution of the company, are decided by the shareholders. Shareholders also confirm new members forwarded by them to the Supervisory Board and the appointment or recall of the managing directors. In accordance with their shares, all shareholders currently have the same number of votes except the Gesellschaft für Informatik, which has three times the number of votes of the other shareholders in proportion to its larger number of shares. Decisions are made in shareholders' meetings which take place at least once a year, or via a written vote.

■ Supervisory Board

The Supervisory Board is responsible for ensuring that the management complies with the center's objectives in a legally and economically meaningful manner. The board is involved in all essential matters with regard to research and financial planning.

The 12-member board (see Fig. 11.2) is composed of four representatives of the Gesellschaft für Informatik, one representative from each of the three founding universities, two representatives of the four universities that subsequently joined, and one representative from each of the German federal government and the two host state governments of Saarland and Rhineland-Palatinate. The Supervisory Board members typically hold office for at least four full fiscal years. The term of office ends with the approval for the fourth fiscal year. In general, representatives of the universities in Darmstadt and Stuttgart and of the universities in Frankfurt and Trier rotate after each term of office.

The Supervisory Board formally appoints and recalls the managing directors and members of the Scientific Directorate, Scientific Advisory Board, and Industrial Curatory Board. Furthermore, all decisions regarding financial issues and company assets must be approved by the Supervisory Board. Consent cannot be given against the votes of the represented (federal) state governments if the matter affects political issues in the area of science or has considerable financial weight. The Supervisory Board also holds decision power with respect to the granting of power of attorney.

■ Management

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH has two managing directors (see Fig. 11.3) who jointly represent the company. These are the *Scientific Director* and the *Technical Administrative Director*.

The Scientific Director is in charge of drafting the company's scientific goals and program planning, and is also a member and the chairperson of the Scientific Directorate. Since May 2014, Prof. Raimund Seidel, Ph.D., is the Scientific Director of Schloss Dagstuhl.

ist Prof. Raimund Seidel, Ph.D., der wissenschaftliche Direktor von Schloss Dagstuhl.

Der Wissenschaftliche Direktor wird dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Dieser Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Seit Juli 2014 hat Frau Heike Meißner diese Position inne.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, Ms Heike Meißner holds this position.

Gremien der Gesellschaft

11.3

Dagstuhl Bodies

Die Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH werden durch drei Gremien unterstützt. Es sind die folgenden:

- Wissenschaftliches Direktorium
- Wissenschaftlicher Beirat
- Kuratorium

Details zu den Gremien werden in den folgenden Abschnitten ausgeführt.

The organs of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH are supported by the following bodies:

- Scientific Directorate
- Scientific Advisory Board
- Industrial Curatory Board

Detailed information about these boards can be found in the sections below.

■ Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 11.4) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Der Wissenschaftliche Direktor gehört dem Wissenschaftlichen Direktorium an. Er empfiehlt dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können das bestehende Wissenschaftliche Direktorium sowie der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt dem Beirat zusammen mit dem Wissenschaftlichen Direktor.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums – mit Ausnahme der des Wissenschaftlichen Direktors – beträgt drei Jahre. Sie beginnt am 1. November des Jahres ihrer Berufung und endet drei Jahre später am 31. Oktober. Eine Wiederberufung ist möglich.

■ Der Wissenschaftliche Beirat

Die Aufgaben des Wissenschaftlichen Beirats (siehe Fig. 11.5) werden nicht nur durch den Gesellschaftsvertrag festgelegt, sondern auch durch die Empfehlungen der Leibniz-Gemeinschaft. Im Sinne dieser wirkt der Wissen-

■ Scientific Directorate

The Scientific Directorate (see Fig. 11.4) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In its biannual directorate meetings, the Scientific Directorate discusses the proposals and decides which of them to accept or reject.

The Scientific Director is member of the Scientific Directorate. He recommends to the Supervisory Board the number of Scientific Directorate members. Candidates for the Scientific Directorate may be suggested not only by the shareholders, but also by the Scientific Directorate and the Scientific Advisory Board. The selection of candidates, which are recommended to the Supervisory Board for appointment, is carried out by the Scientific Advisory Board together with the Scientific Director.

The term of office of Scientific Directorate members – with the exception of the Scientific Director – is three years. It begins on November 1 of the year of appointment and ends three years later on October 31. Reelections are possible.

■ Scientific Advisory Board

The tasks of the Scientific Advisory Board (see Fig. 11.5) are not only defined by the Shareholders' Agreement, but also by the recommendations of the Leibniz Association. The latter stipulates two different ways in

schaftliche Beirat auf zwei Wegen bei der Qualitätssicherung mit. Zum einen berät er die Leitung in Fragen der Forschungs- und Entwicklungsplanung, nimmt Stellung zu den Programmbudgets und gibt Empfehlungen zum Ressourceneinsatz. Er unterstützt weiterhin den Aufsichtsrat bei wichtigen Entscheidungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung (SAE) der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat sollte aus sechs bis zwölf international angesehenen, im Berufsleben stehenden Wissenschaftlern aus dem In- und Ausland bestehen. Die Amtszeit der Mitglieder beträgt vier Jahre, eine einmalige Wiederberufung ist möglich. Der Beirat wählt aus seiner Mitte einen Vorsitzenden. Der Wissenschaftliche Beirat tagt einmal im Jahr. Mitglieder des Beirats werden vom Aufsichtsrat auf Vorschlag des Beirats ernannt.

■ Das Kuratorium

Das Kuratorium (siehe Fig. 11.6) erfüllt eine Transmissionsfunktion zwischen Schloss Dagstuhl und den Forschungsabteilungen und Entwicklungslaboren der Industrie. Es hat die Aufgabe, die Akzeptanz des Zentrums in Verwaltung, Industrie und Wirtschaft abzusichern und als Förderungsorganisation die wirtschaftliche Basis des Zentrums zu verbreitern. Mitglieder des Kuratoriums werden vom Aufsichtsrat ernannt.

Nach seiner Geschäftsordnung hat das Kuratorium mindestens fünf Mitglieder, deren Amtszeit vier Jahre beträgt. Eine einmalige Wiederberufung ist möglich. Die Mitglieder des Kuratoriums unterstützen das Zentrum dabei, aktuelle Themen zu identifizieren und dazu geeignete zugkräftige Organisatoren aus der Industrie zu gewinnen. Sie werden ebenso gebeten, geeignete Personen aus der Industrie als Teilnehmer von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops zu benennen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

which the Scientific Advisory Board is involved in quality assurance. On the one hand, the board offers advice to the management with regard to research as well as development planning and issues comments on the program budget draft, making recommendations on the use of resources. It also assists the Supervisory Board in important decisions with regard to future development of the institute as well as the acquisition of management staff. On the other hand, it carries out an audit of the entire institute between two evaluations by the Senatsausschuss Evaluierung (SAE, Senate Committee Evaluation) of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board, and the SAE.

The Scientific Advisory Board should consist of six to twelve internationally reputable, well established scientists and academics from Germany and abroad. The term of office for members is four years and can be prolonged once. The Scientific Advisory Board members elect a chairperson from their midst. The board convenes once a year. Members are appointed by the Supervisory Board in accordance with the suggestions of the Scientific Advisory Board.

■ Industrial Curatory Board

The Industrial Curatory Board (see Fig. 11.6) performs a transmissional function between the center and the industrial R&D departments and laboratories. Its role is to secure acceptance of Schloss Dagstuhl within the business, industry and administrative communities, and as a promotional organization to broaden the economic basis of the center. Board members are appointed by the Supervisory Board.

According to its rules of procedure, the Industrial Curatory Board consists of at least five members whose term of office is four years. A one-off reappointment for a second term is possible. The board members help the center to identify current R&D topics for seminars and locate attractive organizers in industry. The Industrial Curatory Board is regularly called upon to propose suitable participants for Dagstuhl Seminars and Dagstuhl Perspectives Workshops known to it from its activities. It convenes once a year together with the Scientific Advisory Board.

Gesellschafter Associates
Centrum Wiskunde & Informatica (CWI), The Netherlands
Gesellschaft für Informatik e. V., Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Karlsruher Institut für Technologie (KIT), Germany
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Berlin, Germany
Technische Universität Darmstadt, Germany
Technische Universität Kaiserslautern, Germany
Universität des Saarlandes, Germany
Universität Stuttgart, Germany
Universität Trier, Germany

Fig. 11.1

Associates.

Aufsichtsrat Supervisory Board
Dr. Marc Brüser Ministerium für Wissenschaft, Weiterbildung und Kultur, Mainz, Germany Representative of Rhineland-Palatinate state
Prof. Dr.-Ing. Hannes Federrath Universität Hamburg, Germany Representative of Gesellschaft für Informatik e. V. <i>tenure started in May 2018</i>
Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen Technische Universität Berlin, Germany Representative of Gesellschaft für Informatik e. V. Chairman of the Supervisory Board
Prof. Dr.-Ing. Peter Liggesmeyer Technische Universität Kaiserslautern und Fraunhofer IESE, Germany Representative of Gesellschaft für Informatik e. V. <i>tenure ended in May 2018</i>
Prof. Dr. Volker Lindenstruth Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany Representative of Johann Wolfgang Goethe-Universität Frankfurt am Main
Dr. Rainer Müssner Bundesministerium für Bildung und Forschung, Bonn, Germany Representative of the German federal government <i>tenure started in May 2018</i>
Prof. Dr. Arnd Poetzsch-Heffter Technische Universität Kaiserslautern, Germany Representative of Technische Universität Kaiserslautern
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany Representative of the Saarland
Prof. Dr. Ralph Schenkel Universität Trier Representative of Universität Trier
Prof. Dr. Manfred J. Schmitt Universität des Saarlandes, Saarbrücken, Germany Representative of Universität des Saarlandes
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.
Cornelia Winter Gesellschaft für Informatik e. V., Bonn, Germany Representative of Gesellschaft für Informatik e. V

Fig. 11.2

Supervisory Board members.

Geschäftsführung Management
Heike Meißner (Technisch-administrative Geschäftsführerin Technical Administrative Director) Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern, Germany
Prof. Raimund Seidel, Ph. D. (Wissenschaftlicher Direktor Scientific Director) Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern and Universität des Saarlandes, Saarbrücken, Germany

Fig. 11.3

Management.

Wissenschaftliches Direktorium Scientific Directorate
Prof. Gilles Barthe, Ph. D. IMDEA Software Institute, Madrid, Spain
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Daniel Cremers Technische Universität München, Germany
Prof. Dr. Stefan Diehl Universität Trier, Germany
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany
Prof. Dr. Lynda Hardman Centrum Wiskunde & Informatica (CWI), Amsterdam and University of Utrecht, The Netherlands
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany <i>tenure ended in October 2018</i>
Prof. Dr.-Ing. Oliver Kohlbacher Eberhard Karls Universität Tübingen, Germany
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany
Prof. Dr. Bernhard Nebel Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Bernt Schiele Max-Planck-Institut für Informatik, Saarbrücken, Germany <i>tenure ended in October 2018</i>
Prof. Dr. Albrecht Schmidt Ludwig-Maximilians Universität München, Germany
Prof. Dr.-Ing. Wolfgang Schröder-Preikschat Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany <i>tenure started in November 2018</i>
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Dr. Emmanuel Thomé Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France
Prof. Dr. Heike Wehrheim Universität Paderborn, Germany
Prof. Dr. Verena Wolf Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Martina Zitterbart Karlsruher Institut für Technologie, Germany <i>tenure started in November 2018</i>

Fig. 11.4

Scientific Directorate.

Wissenschaftlicher Beirat Scientific Advisory Board
Prof. Dr. Christel Baier Technische Universität Dresden, Germany
Prof. Dr. Anja Feldmann Max-Planck-Institut für Informatik, Saarbrücken, Germany
Prof. Dr. Manuel V. Hermenegildo IMDEA Software Institute, Madrid and Technical University of Madrid, Spain
Prof. Dr. Claude Kirchner Institut National de Recherche en Informatique et en Automatique (INRIA), Villers-lès-Nancy, France
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, Paderborn and Universität Paderborn, Germany Chairman of the Scientific Advisory Board
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter Heidelberg Laureate Forum Foundation, Heidelberg, Germany
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany

Fig. 11.5

Scientific Advisory Board.

Kuratorium Industrial Curatory Board
Dr. Udo Bub Eötvös Loránd University, Budapest, Hungary
Dr.-Ing. Elmar Dorner SAP SE, Karlsruhe, Germany <i>tenure ended in May 2018</i>
Dr.-Ing. Uwe Franke Daimler AG, Böblingen, Germany
Dr. Goetz Graefe Google, Madison, Wisconsin, United States
Dr. Tim Harris Amazon, Cambridge, United Kingdom <i>tenure started in January 2018</i>
Dr. Michael May Siemens AG, München, Germany
Dr.-Ing. Andreas Wierse SICOS BW GmbH, Stuttgart, Germany

Fig. 11.6

Industrial Curatory Board.

12

**Förderverein „Freunde von
Dagstuhl“**

Association “Friends of Dagstuhl”

■ Förderverein „Freunde von Dagstuhl“

Holger Hermanns (Universität des Saarlandes, Germany)

Erich Reindel (Universität des Saarlandes, Germany)

Seit Mitte 2014 gibt es den Verein zur Förderung von Schloss Dagstuhl — Leibniz-Zentrum für Informatik e.V.. Der sehr technische und holprig klingende Name spiegelt dabei exakt den Vereinszweck wider: die Förderung von Wissenschaft und Forschung im Leibniz-Zentrum für Informatik in Schloss Dagstuhl. Für die Webpräsenz wurde allerdings ein wesentlich geschmeidigerer Name gewählt: „Friends of Dagstuhl“ (<http://www.friends-of-dagstuhl.de>).

Der Verein ist darauf ausgerichtet, finanzielle Mittel zur erfolgreichen Umsetzung des Vereinszwecks zu beschaffen und bereitzustellen sowie die ihm zu diesem Zweck anvertrauten Mittel treuhänderisch zu verwalten. Die Stiftung Informatikzentrum Schloss Dagstuhl wurde daher auch als nicht rechtsfähige Stiftung in den Verein überführt. Seit Ende 2014 vertreten nun die Freunde von Dagstuhl die Stiftung im Rechts- und Geschäftsverkehr und verwalten das Stiftungsvermögen. Der Verein wird von einem Vorstand (siehe Fig. 12.1 und Fig. 12.3) geleitet.

Nach sorgfältiger Prüfung aller Möglichkeiten und unter Einbeziehung des Stiftungsrates (siehe Fig. 12.2) wurden im Jahr 2016 die entscheidenden Schritte zur Anlage des Stiftungsvermögens vorgenommen. Trotz der noch immer andauernden Niedrigzinsphase ist das Vermögen in einer sicheren, aber dennoch nicht ganz rendite-lose Anlageform investiert. Hierzu wurden Verträge mit einer professionellen und auf Stiftungskapital spezialisierten Vermögensverwaltung geschlossen. Bisher erfüllt die gewählte Anlageform auch und gerade im Hinblick auf die schwierige Lage an den Finanzmärkten noch immer die Erwartungen.

Dem Verein gehören aktuell 41 persönliche sowie bisher leider nur vier institutionelle Mitglieder an. Gerade im Hinblick auf die noch geringe Anzahl institutioneller Mitglieder wünschen sich die Freunde von Dagstuhl noch regen Zulauf.

Weitere Informationen zum Verein, aber auch Mitgliedschaftsanträge finden Sie unter <http://www.friends-of-dagstuhl.de>.

■ Association “Friends of Dagstuhl”

Since mid 2014, the registered association to support of Schloss Dagstuhl – Leibniz Center for Informatics (Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.) exists. This very technical and rather clumsy name nevertheless reflects the precise purpose of the association: the support of science and research at the Leibniz Center for Informatics at Schloss Dagstuhl. A significantly smoother name, i.e. “Friends of Dagstuhl”, was chosen for the website (<http://www.friends-of-dagstuhl.de>).

The association aims at acquiring and providing funds for the successful execution of its purpose, as well as holding these funds in trust. The Dagstuhl Foundation (Stiftung Informatikzentrum Schloss Dagstuhl) was therefore integrated into the association as a dependent foundation. Since late 2014, Friends of Dagstuhl represent the foundation in legal and business transactions and manage the foundation assets. The association is chaired by a board (see Fig. 12.1 and Fig. 12.3).

After a careful evaluation of all options and in cooperation with the foundation council (see Fig. 12.2), crucial steps were taken regarding the foundation assets in 2016. It was possible to invest the capital safely but not without return, in spite of the persistently low interest rates. Subsequently, contracts with an investment management company specialized in foundation assets were made. So far, the chosen form of investment has met the expectations, especially in view of the difficult situation on the financial markets.

Currently, the association has 41 individual members but unfortunately only 4 institutional members. Especially with regard to the small number of institutional members, the Friends of Dagstuhl desire to welcome new members.

Further information about the association as well as the membership application form can be found at <http://www.friends-of-dagstuhl.de>.

Vorstand des Vereins Chair of the association
Prof. Dr. Holger Hermanns (Vorstandsvorsitzender First deputy chairperson) Universität des Saarlandes, Saarbrücken, Germany
Angelika Müller-von Brochowski (Schriftführerin Secretary)
Erich Reindel (Schatzmeister Treasurer) Universität des Saarlandes, Saarbrücken, Germany

Fig. 12.1
 Der Vorstand des Vereins “Friends of Dagstuhl”
 The chair of the association “Friends of Dagstuhl”

Stiftungsrat Foundation council
Prof. Dr. Holger Hermanns (Vorstandsvorsitzender des Vereins “Friends of Dagstuhl” First deputy chairperson of the association “Friends of Dagstuhl”) Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Dr. h.c. mult. Kurt Mehlhorn Max Planck Institute for Informatics (MPII), Saarbrücken, Germany
Prof. Dr. Dorothea Wagner Karlsruher Institut für Technologie (KIT), Germany

Fig. 12.2
 Der Stiftungsrat der Stiftung “Informatik-Zentrum Schloss Dagstuhl”
 The council of the foundation “Informatik-Zentrum Schloss Dagstuhl”



Fig. 12.3
 Der Vorstand des Vereins “Friends of Dagstuhl”, v.l.n.r.: Prof. Dr. Holger Hermanns, Angelika Müller-von Brochowski, und Erich Reindel.
 The chair of the association “Friends of Dagstuhl”, f.l.t.r.: Prof. Holger Hermanns, Angelika Müller-von Brochowski, and Erich Reindel.

13 Statistiken

Statistics

Statistiken zu Seminaren und Workshops

13.1

Statistics on Seminars and Workshops

In diesem Abschnitt werden statistische Daten zum wissenschaftlichen Programm und der Zusammenstellung der Teilnehmer aufgeführt. Die Diagramme und Tabellen sind dabei wie nachfolgend beschrieben gegliedert.

Antrags-bezogene Daten: Die Anzahl eingereicherter Anträge von Dagstuhl Seminaren und Dagstuhl Perspektiven Workshops sowie deren Akzeptanzraten sind in Fig. 13.1 dargestellt. Fig. 13.2 zeigt, wie die akzeptierten Seminare und Workshops sich bezüglich Größe und Länge aufgliedern.

Veranstaltungs-bezogene Daten: Fig. 13.3 zeigt Anzahl und Anteil der eingeladenen Seminar Teilnehmer, welche die Einladung annehmen bzw. ablehnen. Die Verteilung dieser Annahmerate ist in Fig. 13.4 dargestellt. Fig. 13.5 zeigt dagegen, wie viel Prozent der zugesagten Größe (gemessen an der Personenanzahl) tatsächlich von einem Seminar belegt wurde. Daten zu Anzahl, Größe und Dauer der durchgeführten Seminare sind in Fig. 13.6 angegeben. Fig. 13.7 zeigt die Anzahl der verschiedenen Veranstaltungstypen.

Teilnehmer-bezogene Daten: Die Teilnehmerzahlen – abhängig vom Veranstaltungstyp – gibt Fig. 13.8 an. Fig. 13.9 zeigt die Verteilung der Herkunftsländer unserer Gäste.

Umfrage-bezogene Daten: Hier stellen wir ausgewählte Daten unserer fortlaufenden Befragung von Teilnehmern an Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops dar. Ein Überblick über die Ergebnisse der regelmäßigen Gästebefragungen kann Fig. 13.10 entnommen werden. Die Anzahl von früheren Seminarbesuchen kann man Fig. 13.11 entnehmen. Fig. 13.12 gibt Auskunft über die Altersstruktur der Teilnehmer. Während Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops sich primär an Forscher aus Universitäten und Forschungseinrichtungen richten, sind auch Anwender und Forscher aus der Industrie stets willkommen. Die Verteilung ihres Anteils ist in Fig. 13.13 gezeigt.

Auslastungs-bezogene Daten: Die Auslastung des Zentrums wird schließlich in Fig. 13.14 an Hand der Übernachtungen und ihrer Verteilung über die einzelnen Wochen getrennt nach Veranstaltungstypen aufgezeigt.

Geschlechter-bezogene Daten: Fig. 13.15 enthält Daten zur Geschlechter-Verteilung in der Seminarleitung. Dagegen zeigt Fig. 13.16 die Quote von Frauen bei der Beantragung von Seminaren sowohl bezüglich der Teams als auch bezüglich der gesamten Antragsteller. Die Abbildungen Fig. 13.17 und Fig. 13.18 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler. Die Verteilung der Rate der weiblichen Teilnehmer in den einzelnen Seminaren wird in Fig. 13.19 aufgezeigt.

This section provides statistical data about the scientific program and the composition of program participants. Charts and tables in this chapter may be outlined as follows.

Proposal-related data: Fig. 13.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, as well as acceptance rates for recent years. The size and duration of accepted seminars and workshops are displayed in Fig. 13.2.


Event-related data: Fig. 13.3 shows the number and the fraction of invited seminar participants who accepted or declined the invitation. The distribution of the rate is given in Fig. 13.4. In contrast, Fig. 13.5 visualizes how much of the reserved space was actually used by seminar participants. Data related to the number of seminars held in the last years together with their sizes and durations are given in Fig. 13.6. Fig. 13.7 shows the distribution of different types of events at Dagstuhl.

Participant-related data: Fig. 13.8 shows the number of participants according to event type. Fig. 13.9 shows the distribution of country affiliations.

Survey-related data: In this section we present data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results of the participants survey for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Fig. 13.10. Fig. 13.11 displays how often participants have attended seminars in the past. Fig. 13.12 gives data on the seniority of participants. While Dagstuhl Seminars and Dagstuhl Perspectives Workshops are mainly oriented towards academic researchers, also researchers and developers from industry are welcome. The distribution of their ratio compared to all participants of a seminar is shown in Fig. 13.13.

Utilization-related data: Finally, Fig. 13.14 states the number of overnight stays – separated by event type – hosted at Schloss Dagstuhl as well as their distribution about the weeks.

Gender-related data: Fig. 13.15 shows mixed-gender data with respective to organizer teams of Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In contrast Fig. 13.16 presents this data with respect to proposed seminar events. In Fig. 13.17 and Fig. 13.18 data is given with regard to female participants and invitees, respectively. The distribution of the rate of female participants by seminar and year is displayed in Fig. 13.19.




(a) Chart for 2012–2018

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
2012	89	68	76.4	21	23.6
2013	107	72	67.3	35	32.7
2014	98	65	66.3	33	33.7
2015	99	65	65.7	34	34.3
2016	125	79	63.2	46	36.8
2017	102	70	68.6	32	31.4
2018	136	81	59.6	55	40.4

(b) Detailed numbers for 2012–2018

Fig. 13.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.




(a) Chart for 2012–2018

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2012	12	20	0	36	68
2013	13	25	1	33	72
2014	12	19	1	33	65
2015	10	20	2	33	65
2016	11	20	2	46	79
2017	15	15	1	39	70
2018	10	22	0	49	81

(b) Detailed numbers for 2012–2018

Fig. 13.2
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2012–2018. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.




(a) Chart for 2012–2018

Year	Invitees	Acceptances		Declines	
	#	#	%	#	%
2012	5033	2346	46.6	2687	53.4
2013	5591	2639	47.2	2952	52.8
2014	5285	2590	49.0	2695	51.0
2015	5023	2473	49.2	2550	50.8
2016	5060	2393	47.3	2667	52.7
2017	5267	2572	48.8	2695	51.2
2018	4692	2320	49.4	2372	50.6

(b) Detailed numbers for 2012–2018

Fig. 13.3
Total number of invitees, acceptances, and declines for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.




(a) Chart for 2012–2018

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2012	21.4	80.5	47.2	11.0
2013	21.9	71.6	48.4	11.2
2014	26.7	80.0	50.2	11.2
2015	28.4	71.6	50.7	12.4
2016	26.9	80.4	48.6	11.2
2017	25.5	77.5	50.3	12.4
2018	21.8	77.0	51.2	12.0

(b) Detailed numbers for 2012–2018

Fig. 13.4

Distribution of the acceptance rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2012–2018. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.




(a) Chart for 2012–2018

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2012	48.9	137.8	92.4	17.6
2013	55.6	113.3	92.1	12.2
2014	60.0	113.3	90.6	10.3
2015	63.3	116.7	89.6	12.5
2016	55.6	113.3	86.7	11.8
2017	60.0	123.3	87.3	12.3
2018	60.0	113.3	90.3	10.2

(b) Detailed numbers for 2012–2018

Fig. 13.5

Distribution of the occupancy rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2012–2018. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.




(a) Chart for 2012–2018

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2012	12	10	1	41	64
2013	11	23	1	40	75
2014	11	24	1	39	75
2015	14	19	1	38	72
2016	5	26	2	39	72
2017	18	16	2	41	77
2018	9	14	1	41	65

(b) Detailed numbers for 2012–2018

Fig. 13.6

Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2012–2018. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.




(a) Chart for 2012–2018

Year	DS	PW	GI	EDU	RGM	Total
2012	59	5	2	4	52	122
2013	74	1	0	5	33	113
2014	70	5	3	4	30	112
2015	68	4	3	2	30	107
2016	68	4	3	6	35	116
2017	76	1	1	3	40	121
2018	63	2	2	2	40	109

(b) Detailed numbers for 2012–2018

Fig. 13.7
Number of all events held at Dagstuhl, by type. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.



(a) Chart for 2012–2018

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	
2012	2226	64.4	120	3.5	48	1.4	144	4.2	916	26.5	3454
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503
2014	2463	72.2	127	3.7	86	2.5	144	4.2	589	17.3	3409
2015	2385	72.3	88	2.7	90	2.7	111	3.4	624	18.9	3298
2016	2280	68.0	113	3.4	78	2.3	232	6.9	650	19.4	3353
2017	2551	77.1	21	0.6	21	0.6	131	4.0	584	17.7	3308
2018	2268	70.8	52	1.6	50	1.6	99	3.1	733	22.9	3202

(b) Detailed numbers for 2012–2018


Fig. 13.8
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

Country	A	B	Total	Country	A	B	Total
Germany	480	687	1167	Brazil	9	1	10
United States	562	36	598	Singapore	10	0	10
United Kingdom	213	21	234	China	8	1	9
France	174	21	195	Slovenia	9	0	9
Canada	92	7	99	Greece	8	0	8
Netherlands	87	9	96	New Zealand	8	0	8
Switzerland	81	8	89	Russian Federation	8	0	8
Austria	76	11	87	Hong Kong	7	0	7
Italy	64	5	69	Republic of Korea	7	0	7
Sweden	54	6	60	Estonia	4	1	5
Israel	48	0	48	Hungary	4	0	4
Australia	37	6	43	Mexico	4	0	4
Luxembourg	8	30	38	Turkey	3	0	3
Belgium	33	0	33	Argentina	1	1	2
Denmark	26	4	30	Iceland	2	0	2
Japan	30	0	30	Lebanon	2	0	2
India	26	0	26	Romania	1	1	2
Spain	21	1	22	Saudi Arabia	2	0	2
Czech Republic	21	0	21	Croatia	1	0	1
Finland	20	0	20	Qatar	1	0	1
Norway	13	5	18	Serbia	1	0	1
Poland	16	0	16	Slovak Republic	1	0	1
Portugal	14	1	15	South Africa	1	0	1
Ireland	11	3	14	Taiwan	1	0	1
Pakistan	0	14	14	United Arab Emirates	1	0	1
Chile	9	2	11	Total	2320	882	3202


(a) Details for 2018 by country

Region	A		B		Total	
	#	%	#	%	#	%
Germany	480	20.7	687	77.9	1167	36.4
Europe (w/o Germany)	974	42	127	14.4	1101	34.4
North America	654	28.2	43	4.9	697	21.8
Asia	143	6.2	15	1.7	158	4.9
Australia	45	1.9	6	0.7	51	1.6
South America	23	1	4	0.5	27	0.8
Africa	1	0	0	0	1	0
Total	2320	100	882	100	3202	100

(b) Details for 2018 by region




(c) Graphical distribution of seminar type A in 2012–2018 by year and region



(d) Graphical distribution of seminar type B in 2012–2018 by year and region

Fig. 13.9

Number of Dagstuhl guests by country of origin. A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).




(a) Graphical distribution for 2018

	2011	2012	2013	2014	2015	2016	2017	2017 – Detailed Numbers					
	Ø	Ø	Ø	Ø	Ø	Ø	Ø	1	2	3	4	5	total
would come again	4.9	4.9	4.9	4.9	4.9	4.9	4.9	0	3	6	99	1255	1363
found inspiration	4.4	4.5	4.4	4.5	4.5	4.5	4.6	4	9	36	473	854	1376
found collaboration	4.0	4.0	4.1	4.1	4.1	4.2	4.2	9	34	179	538	600	1360
found insight from neighboring fields	4.1	4.1	4.2	4.3	4.2	4.2	4.2	12	43	154	542	615	1366
found new research direction	4.0	4.0	4.0	4.1	4.1	4.1	4.2	6	38	143	650	522	1359
group composition	4.4	4.4	4.4	4.5	4.5	4.5	4.5	3	18	45	521	784	1371
integration of junior researchers	4.2	4.2	4.2	4.2	4.3	4.3	4.3	7	48	123	512	672	1362
new professional contacts	3.6	3.7	3.7	3.6	3.7	3.8	3.8	34	136	296	459	417	1342
exchange between academia and industry	4.3	4.2	4.2	4.3	4.3	4.4	4.4	5	29	93	318	499	944
advance information from Dagstuhl	4.4	4.4	4.4	4.4	4.4	4.4	4.4	5	8	111	494	713	1331
advance information from organizers	4.1	4.1	4.1	4.1	4.2	4.1	4.3	6	32	201	458	615	1312
number and length of talks	4.1	4.2	4.1	4.2	4.3	4.2	4.3	7	62	98	485	699	1351
opportunity for one on one talks	4.4	4.5	4.5	4.5	4.6	4.6	4.6	4	23	48	331	954	1360
flexibility of schedule	4.2	4.2	4.3	4.3	4.4	4.3	4.3	4	57	124	442	720	1347
open and honest discussion	4.6	4.7	4.7	4.7	4.7	4.7	4.7	1	2	23	307	1027	1360
outing	4.1	4.1	4.1	4.1	4.2	4.2	4.3	12	34	135	362	524	1067
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	2	7	32	286	1030	1357
conference facilities	4.7	4.6	4.7	4.6	4.7	4.7	4.7	0	3	37	359	958	1357
IT facilities	4.4	4.4	4.4	4.3	4.4	4.3	4.4	5	21	111	426	627	1190
staff support	4.7	4.7	4.7	4.7	4.7	4.7	4.8	0	1	27	235	1023	1286
meals	4.2	4.1	4.1	4.1	4.1	4.1	4.1	12	49	222	607	460	1350
rooms	4.4	4.4	4.4	4.4	4.4	4.4	4.4	3	25	147	481	693	1349
leisure facilities	4.6	4.6	4.6	4.6	4.5	4.5	4.6	1	4	46	367	813	1231
library services	4.5	4.5	4.5	4.5	4.5	4.5	4.6	0	1	37	166	419	623

(b) Averages for 2012–2018 and detailed numbers for 2018: 1 = very low, 2 = low, 3 = medium, 4 = high, 5 = very high

Fig. 13.10 Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.




(a) Graphical distribution for 2012–2018

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	
2012	483	44	193	17	135	12	295	27	1106
2013	630	44	237	17	145	10	422	29	1434
2014	561	40	239	17	144	10	443	32	1387
2015	573	40	234	17	158	11	451	32	1416
2016	654	46	217	15	137	10	410	29	1418
2017	607	43	222	16	148	10	446	31	1423
2018	557	41	219	16	148	11	425	32	1349

(b) Detailed numbers for 2012–2018

Fig. 13.11

Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants and their previous instances of attendance in Dagstuhl Seminars or Dagstuhl Perspectives Workshops, according to our guest survey.




(a) Chart for 2012–2018

Year	Junior		Senior		Total
	#	%	#	%	
2012	307	34.6	580	65.4	887
2013	413	35.4	754	64.6	1167
2014	382	33.3	765	66.7	1147
2015	410	34.9	764	65.1	1174
2016	404	33.9	787	66.1	1191
2017	422	35.2	778	64.8	1200
2018	401	35.7	722	64.3	1123

(b) Detailed numbers for 2012–2018

Fig. 13.12

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.



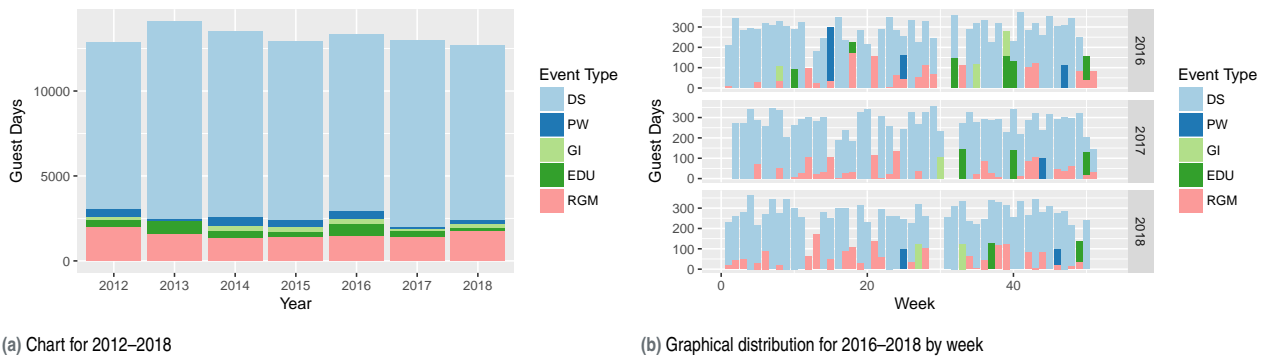
(a) Chart for 2012–2018

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2012	0.0	53.8	12.4	13.7
2013	0.0	66.7	11.6	12.8
2014	0.0	35.3	9.4	9.4
2015	0.0	58.8	9.8	10.5
2016	0.0	41.2	10.3	11.0
2017	0.0	60.0	10.9	11.6
2018	0.0	41.7	11.1	10.4

(b) Detailed numbers for 2012–2018

Fig. 13.13

Distribution of the rate of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2012–2018, according to our guest survey. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation. Occupation in business includes “industrial research”, “industrial development”, and “self employed”.



Year	DS	PW	GI	EDU	RGM	Total
2012	9798	458	190	393	2031	12870
2013	11612	130	0	753	1614	14109
2014	10939	475	348	390	1370	13522
2015	10491	380	344	261	1424	12900
2016	10362	495	315	703	1462	13337
2017	10989	102	105	401	1391	12988
2018	10270	182	250	231	1740	12673

(c) Detailed numbers for 2012–2018

Fig. 13.14

Number of overnight stays at Schloss Dagstuhl. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.





Fig. 13.15

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.




(a) Chart for 2012–2018

Year	Proposer Teams			Proposers		
	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2012	89	49	55.1	341	56	16.4
2013	107	53	49.5	431	66	15.3
2014	98	56	57.1	387	63	16.3
2015	99	62	62.6	391	80	20.5
2016	125	82	65.6	491	99	20.2
2017	102	62	60.8	394	81	20.6
2018	136	102	75.0	522	139	26.6

(b) Detailed numbers for 2012–2018

Fig. 13.16

Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.




(a) Chart for 2012–2018

Year	Invitees			Participants		
	Total	Female		Total	Female	
	#	#	%	#	#	%
2012	5033	826	16.4	2346	379	16.2
2013	5591	889	15.9	2639	401	15.2
2014	5285	943	17.8	2590	406	15.7
2015	5023	845	16.8	2473	369	14.9
2016	5060	978	19.3	2393	437	18.3
2017	5267	1110	21.1	2572	495	19.2
2018	4692	1086	23.1	2320	453	19.5

(b) Detailed numbers for 2012–2018

Fig. 13.17

Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.




(a) Chart for 2012–2018

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2012	826	379	45.9	447	54.1
2013	889	401	45.1	488	54.9
2014	943	406	43.1	537	56.9
2015	845	369	43.7	476	56.3
2016	978	437	44.7	541	55.3
2017	1110	495	44.6	615	55.4
2018	1086	453	41.7	633	58.3

(b) Detailed numbers for 2012–2018

Fig. 13.18

Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.



(a) Chart for 2012–2018

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2012	3.7	35.5	16.2	7.8
2013	0.0	36.7	15.1	7.3
2014	0.0	53.8	15.9	11.1
2015	0.0	31.8	14.8	7.7
2016	0.0	40.9	18.3	9.1
2017	4.2	50.0	19.7	9.8
2018	5.7	47.8	19.8	9.2

(b) Detailed numbers for 2012–2018

Fig. 13.19

Distribution of female participants rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2012–2018. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.


Statistiken zur Bibliographiedatenbank dblp

13.2

Statistics of the dblp computer science bibliography

Dieser Abschnitt enthält statistische Daten zur Bibliographiedatenbank dblp. Fig. 13.20 listet die durchschnittlichen Nutzungszahlen der letzten Jahre. Ein Überblick über die Entwicklung des dblp Datenbestandes kann Fig. 13.21 und Fig. 13.22 entnommen werden. Fig. 13.23–13.25 geben Auskunft über die kontinuierliche Datenkuration und -anreicherung des Bestandes.

This section provides statistical data about the dblp computer science bibliography. Fig. 13.20 shows the average usage statistics of the dblp servers in the past years. An overview of the development of the dblp database can be found in Fig. 13.21 and Fig. 13.22. Information about the continuous data curation and enrichment of existing records can be found in Fig. 13.23–13.25.




(a) Chart for 2014–2018

	Trier 1		Trier 2		Dagstuhl		Total	
	2017	2018	2017	2018	2017	2018	2017	2018
user sessions (visits) per day	27,931	31,530	2,836	3,233	5,366	11,483	36,133	46,247
page views per day	466,989	618,067	35,140	20,208	85,537	202,301	587,668	840,577
page views per user session	16.7	19.6	12.4	6.2	15.9	17.6	16.3	18.2
distinct users (IPs) per month	390,886	451,769	58,975	27,448	86,985	197,270	536,847	676,489
data served per month	1,082.3 GB	1,535.0 GB	82.8 GB	72.6 GB	235.0 GB	469.7 GB	1,400.1 GB	2,077.3 GB

(b) Detailed numbers for the past two years

Fig. 13.20

Average usage of the three dblp servers. Trier 1 = dblp.uni-trier.de, Trier 2 = dblp2.uni-trier.de, Dagstuhl = dblp.dagstuhl.de. All figures exclude traffic caused by recognized bots and web crawlers. Usage data has not been collected before 2014. In 2015, changes have been made in the server setup in order to shift traffic from development server Trier 2 to the more powerful server Trier 1. Since 2017, server Dagstuhl has been promoted to play a more prominent role under the domain dblp.org




(a) Chart for 1996–2018

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total #
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	
2012	16,037	0.7	850,603	39.3	1,199,892	55.5	9,631	0.4	20,154	0.9	13,125	0.6	0	0.0	53,810	2.5	2,163,252
2013	16,819	0.7	997,820	40.2	1,350,713	54.4	12,797	0.5	22,771	0.9	13,125	0.5	0	0.0	69,905	2.8	2,483,950
2014	17,533	0.6	1,129,231	39.8	1,545,065	54.5	14,470	0.5	26,137	0.9	14,690	0.5	0	0.0	88,217	3.1	2,835,343
2015	18,318	0.6	1,281,245	40.0	1,724,262	53.9	16,288	0.5	30,044	0.9	19,103	0.6	12	0.0	110,974	3.5	3,200,246
2016	51,070	1.4	1,429,427	39.7	1,912,895	53.1	19,774	0.5	33,782	0.9	20,174	0.6	26	0.0	134,354	3.7	3,601,502
2017	77,408	1.9	1,576,972	39.4	2,091,486	52.2	23,101	0.6	37,049	0.9	23,089	0.6	49	0.0	174,723	4.4	4,003,877
2018	83,249	1.9	1,725,704	39.0	2,306,585	52.2	24,708	0.6	40,795	0.9	23,150	0.5	514	0.0	216,984	4.9	4,421,689

(b) Detailed numbers for 2012–2018

Fig. 13.21

Development of the total size of the dbp database.




(a) Chart for 1996–2018

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total #
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	
2012	6,681	2.1	144,869	44.9	151,230	46.9	2,178	0.7	3,003	0.9	918	0.3	0	0.0	13,885	4.3	322,764
2013	782	0.2	147,217	45.9	150,821	47.0	3,166	1.0	2,617	0.8	0	0.0	0	0.0	16,095	5.0	320,698
2014	714	0.2	131,411	37.4	194,352	55.3	1,673	0.5	3,366	1.0	1,565	0.4	0	0.0	18,312	5.2	351,393
2015	785	0.2	152,014	41.7	179,197	49.1	1,818	0.5	3,907	1.1	4,413	1.2	12	0.0	22,757	6.2	364,903
2016	32,752	8.2	148,182	36.9	188,633	47.0	3,486	0.9	3,738	0.9	1,071	0.3	14	0.0	23,380	5.8	401,256
2017	26,338	6.5	147,545	36.7	178,591	44.4	3,327	0.8	3,267	0.8	2,915	0.7	23	0.0	40,369	10.0	402,375
2018	5,841	1.4	148,732	35.6	215,099	51.5	1,607	0.4	3,746	0.9	61	0.0	465	0.1	42,261	10.1	417,812

(b) Detailed numbers for 2012–2018

Fig. 13.22

Development of newly included publications in dbp. The negative number of new *Incollection* records in 2011 results from relabeling several thousand existing records with the newly introduced *Reference* type. Similarly, in the same year, several thousand *Articles* and *Inproceedings* records have been labeled as *Informal*.




(a) Chart for 2013–2018

Year	Merge		Split		Distribute		Total #
	#	%	#	%	#	%	
2013	6,927	68.6	781	7.7	2,395	23.7	10,103
2014	6,936	67.8	797	7.8	2,501	24.4	10,234
2015	7,515	69.5	788	7.3	2,514	23.2	10,817
2016	8,608	70.3	978	8.0	2,666	21.8	12,252
2017	9,092	66.1	1,464	10.6	3,192	23.2	13,748
2018	9,234	63.7	2,061	14.2	3,207	22.1	14,502

(b) Detailed numbers for 2013–2018

Fig. 13.23

Data curation of existing dblp person profiles. No curation data has been tracked before 2013. The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp person profiles. We distinguish between three curation cases: *Merge* = Two or more synonymous person profiles have been merged into a single profile. *Split* = A single, homonymous person profile has been split into two or more profiles. *Distribute* = A mixed case where records from two or more person profiles have been redistributed between two or more profiles.




(a) Chart for 2012–2018

Year	Homepages
2012	17,384
2013	19,511
2014	22,101
2015	25,814
2016	31,353
2017	35,973
2018	40,179

(b) Detailed numbers for 2012–2018

Fig. 13.24

Linked and verified academic homepages in dblp person profiles. A single person profile may be linked to multiple academic homepages. These figures exclude linked external IDs which are given in Figure 13.25.



(a) Chart for 2012–2018

Year	ORCID		Google Scholar		Wikidata		Wikipedia (en)		ACM DL		ResearcherID		VIAF	
	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2012	6	0.4	459	29.3	0	0.0	116	7.4	977	62.3	1	0.1	0	0.0
2013	14	0.6	947	39.9	0	0.0	207	8.7	1,169	49.2	2	0.1	2	0.1
2014	24	0.8	1,379	46.9	0	0.0	271	9.2	1,173	39.9	8	0.3	3	0.1
2015	89	1.2	2,510	34.6	0	0.0	1,002	13.8	1,225	16.9	34	0.5	573	7.9
2016	1,717	12.8	4,999	37.3	4	0.0	1,234	9.2	1,236	9.2	154	1.2	585	4.4
2017	11,591	31.8	7,326	20.1	4,046	11.1	1,550	4.3	1,425	3.9	1,297	3.6	945	2.6
2018	33,185	34.1	11,226	11.5	17,405	17.9	5,065	5.2	2,339	2.4	4,824	5.0	1,057	1.1

Year	Math Genealogy		Zentralblatt MATH		LOC		ISNI		GND		Twitter		Other		Total #
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	
2012	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	9	0.6	1,568
2013	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	1	0.0	32	1.3	2,374
2014	0	0.0	4	0.1	0	0.0	1	0.0	0	0.0	13	0.4	64	2.2	2,940
2015	341	4.7	622	8.6	357	4.9	1	0.0	240	3.3	67	0.9	200	2.8	7,261
2016	420	3.1	687	5.1	357	2.7	7	0.1	243	1.8	184	1.4	1,562	11.7	13,389
2017	879	2.4	701	1.9	677	1.9	671	1.8	481	1.3	453	1.2	4,419	12.1	36,461
2018	2,717	2.8	2,435	2.5	3,320	3.4	1,381	1.4	2,793	2.9	1,764	1.8	7,773	8.0	97,284

(b) Detailed numbers for 2012–2018

Fig. 13.25

Linked and verified external person IDs in dblp person profiles. A single person profile may be linked to multiple external IDs.

Statistiken zu Dagstuhl Publishing

13.3

Statistics of Dagstuhl Publishing

Dieser Abschnitt enthält statistische Daten zum Publikationswesen von Schloss Dagstuhl.

Ein Überblick über die Entwicklung der seminarbezogenen Veröffentlichungen kann den ersten drei Diagrammen und Tabellen entnommen werden. Fig. 13.26 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen, Fig. 13.27 die der Publikationen in der Reihe Dagstuhl Manifestos und schließlich Fig. 13.28 die der veröffentlichten Bände in der Reihe Dagstuhl Follow-Ups.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 13.29 fasst die Daten in der Reihe OASICs und Fig. 13.30 die der Reihe LIPICs zusammen.

Die Kennzahlen der Zeitschrift LITES können Fig. 13.31 entnommen werden.

Die verschiedenen Publikationsserien wurden in unterschiedlichen Jahren zwischen 2009 und 2015 gegründet. Wir stellen in den Statistiken dennoch stets den gesamten Zeitraum (2012–2018) dar.


In this section the statistical data of Dagstuhl Publishing are presented.

The first three figures present the development of the seminar-focused series: Fig. 13.26 summarizes the data of the periodical Dagstuhl Reports, Fig. 13.27 the data of the Dagstuhl Manifestos series, and, finally, Fig. 13.28 that of the volumes published in the Dagstuhl Follow-Ups series.

The statistical data of the service-focused series are presented afterwards. Fig. 13.29 presents numbers related to OASICs and Fig. 13.30 numbers related to LIPICs.

We summarize the publications of the journal LITES in Fig. 13.31.

Please note that the publication series were established in different years in the period between 2009 and 2015. However, we always consider this complete period (2012–2018).




(a) Graphical distribution for 2012–2018

Year	Articles	Pages
2012	42	913
2013	84	2059
2014	62	1464
2015	62	1636
2016	89	1910
2017	48	1138
2018	78	1938

(b) Detailed numbers for 2012–2018

Fig. 13.26
Statistics about Dagstuhl Reports published between 2012 to 2018.




(a) Graphical distribution for 2012–2018

Year	Articles	Pages
2012	3	62
2013	2	37
2014	2	52
2015	2	41
2016	2	39
2017	0	0
2018	1	22

(b) Detailed numbers for 2012–2018

Fig. 13.27
Statistics about Dagstuhl Manifestos published between 2012 to 2018.




(a) Graphical distribution for 2012–2018

Year	Volumes	Articles	Pages
2012	1	13	246
2013	3	25	641
2014	0	0	0
2015	0	0	0
2016	0	0	0
2017	1	12	346
2018	0	0	0

(b) Detailed numbers for 2012–2018

Fig. 13.28
Statistics about Dagstuhl Follow-Ups volumes published between 2012 to 2018.




(a) Graphical distribution for 2012–2018

Year	Volumes	Articles	Pages
2012	8	106	1192
2013	7	117	1265
2014	8	116	1264
2015	6	66	674
2016	6	85	1078
2017	3	50	684
2018	7	107	1312

(b) Detailed numbers for 2012–2018

Fig. 13.29
Statistics about OASlcs volumes published between 2012 to 2018.




(a) Graphical distribution for 2012–2018

(b) Detailed numbers for 2012–2018

Year	Volumes	Articles	Pages
2012	5	215	2591
2013	6	195	2607
2014	5	204	2752
2015	16	553	8565
2016	19	939	14222
2017	25	1127	17687
2018	32	1387	21871

Fig. 13.30
Statistics about LIPIcs volumes published between 2012 to 2018.




(a) Graphical distribution for 2012–2018

(b) Detailed numbers for 2012–2018

Year	Articles	Pages
2012	0	0
2013	0	0
2014	7	119
2015	3	58
2016	5	144
2017	8	218
2018	4	134

Fig. 13.31
Statistics about LITES articles published between 2012 to 2018.



(a) Graphical distribution for 2012–2018

(b) Detailed numbers for 2012–2018

Year	Articles	Pages
2012	0	0
2013	0	0
2014	0	0
2015	12	50
2016	14	30
2017	23	79
2018	19	90

Fig. 13.32
Statistics about DARTS artifacts published between 2012 to 2018.

14 **Veranstaltungen 2018** *Schedule of Events 2018*

Dagstuhl-Seminare

14.1

Dagstuhl Seminars**18021 – Symmetric Cryptography**

Joan Daemen (Radboud University Nijmegen, NL, and STMicroelectronics – Diegem, BE), Tetsu Iwata (Nagoya University, JP), Nils Gregor Leander (Ruhr-Universität Bochum, DE), Kaisa Nyberg (Aalto University, FI)

January 7–12, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18021>

18031 – Personalized Multiobjective Optimization: An Analytics Perspective

Kathrin Klamroth (Universität Wuppertal, DE), Joshua D. Knowles (University of Birmingham, GB), Günter Rudolph (TU Dortmund, DE), Margaret M. Wiecek (Clemson University, US)

January 14–19, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18031>

18041 – Foundations of Data Visualization

Helwig Hauser (University of Bergen, NO), Penny Rheingans (University of Maryland, Baltimore County, US), Gerik Scheuermann (Universität Leipzig, DE)

January 21–26, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18041>

18051 – Proof Complexity

Albert Atserias (UPC – Barcelona, ES), Jakob Nordström (KTH Royal Institute of Technology – Stockholm, SE), Pavel Pudlák (The Czech Academy of Sciences – Prague, CZ), Rahul Santhanam (University of Oxford, GB)

January 28 to February 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18051>

18052 – Genetic Improvement of Software

Stephanie Forrest (Arizona State University – Tempe, US), William B. Langdon (University College London, GB), Claire Le Goues (Carnegie Mellon University – Pittsburgh, US), Justyna Petke (University College London, GB)

January 28 to February 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18052>

18061 – Evidence About Programmers for Programming Language Design

Stefan Hanenberg (Universität Duisburg-Essen, DE), Brad A. Myers (Carnegie Mellon University – Pittsburgh, US), Bonita Sharif (Youngstown State University, US), Andreas Stefik (Univ. of Nevada – Las Vegas, US)

February 4–9, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18061>

18071 – Planning and Operations Research

J. Christopher Beck (University of Toronto, CA), Daniele Magazzeni (King’s College London, GB), Gabriele Röger (Universität Basel, CH), Willem-Jan Van Hoeve (Carnegie Mellon University – Pittsburgh, US)

February 11–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18071>

18081 – Designing and Implementing Algorithms for Mixed-Integer Nonlinear Optimization

Pierre Bonami (IBM Spain – Madrid, ES), Ambros M. Gleixner (Konrad-Zuse-Zentrum – Berlin, DE), Jeff Linderoth (University of Wisconsin – Madison, US), Ruth Misener (Imperial College London, GB)

February 18–23, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18081>

18082 – Formal Methods for the Synthesis of Biomolecular Circuits

Yaakov Benenson (ETH Zürich – Basel, CH), Neil Dalchau (Microsoft Research UK – Cambridge, GB), Heinz Koepl (TU Darmstadt, DE), Oded Maler (VERIMAG – Grenoble, FR)

February 18–23, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18082>

18091 – Data Consistency in Distributed Systems: Algorithms, Programs, and Databases

Annette Bieniusa (TU Kaiserslautern, DE), Alexey Gotsman (IMDEA Software – Madrid, ES), Bettina Kemme (McGill University – Montreal, CA), Marc Shapiro (Sorbonne University – Paris, FR)

February 25 to March 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18091>

18092 – The Logical Execution Time Paradigm: New Perspectives for Multicore Systems

Rolf Ernst (TU Braunschweig, DE), Stefan Kuntz (Continental Automotive GmbH – Regensburg, DE), Sophie Quinton (INRIA – Grenoble, FR), Martin Simons (Daimler AG – Böblingen, DE)

February 25–28, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18092>

18101 – Scheduling

Magnús M. Halldórsson (Reykjavik University, IS), Nicole Megow (Universität Bremen, DE), Clifford Stein (Columbia University, US)

March 4–9, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18101>

18102 – Dynamic Traffic Models in Transportation Science

Roberto Cominetti (Adolfo Ibáñez University – Santiago, CL), Tobias Harks (Universität Augsburg, DE), Carolina Osorio (MIT – Cambridge, US), Britta Peis (RWTH Aachen, DE)

March 4–9, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18102>

18111 – Loop Optimization

Sebastian Hack (Universität des Saarlandes, DE), Paul H. J. Kelly (Imperial College London, GB), Christian Lengauer (Universität Passau, DE)

March 11–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18111>

18112 – Coding Theory for Inference, Learning and Optimization

Po-Ling Loh (University of Wisconsin – Madison, US), Arya Mazumdar (University of Massachusetts – Amherst, US), Dimitris Papailiopoulos (University of Wisconsin – Madison, US), Rüdiger Urbanke (EPFL – Lausanne, CH)

March 11–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18112>

18121 – Machine Learning and Model Checking Join Forces

Nils Jansen (Radboud University Nijmegen, NL), Joost-Pieter Katoen (RWTH Aachen, DE), Pushmeet Kohli (Google DeepMind – London, GB), Jan Kretinsky (TU München, DE)

March 18–23, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18121>

18122 – Automatic Quality Assurance and Release

Bram Adams (Polytechnique Montreal, CA), Benoit Baudry (KTH Royal Institute of Technology – Stockholm, SE), Sigrid Eldh (Ericsson AB – Stockholm, SE), Andy Zaidman (TU Delft, NL)

March 18–21, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18122>

18151 – Program Equivalence

Shuvendu K. Lahiri (Microsoft Research – Redmond, US), Andrzej Murawski (University of Oxford, GB), Ofer Strichman (Technion – Haifa, IL), Matthias Ulbrich (KIT – Karlsruher Institut für Technologie, DE)

April 8–13, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18151>

18152 – Blockchains, Smart Contracts and Future Applications

Foteini Baldimtsi (George Mason University – Fairfax, US), Stefan Katzenbeisser (TU Darmstadt, DE), Volkmar Lotz (SAP Labs France – Mougins, FR), Edgar Weippl (Secure Business Austria Research, AT)

April 8–13, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18152>

18161 – Visualization of Biological Data – Crossroads

Jan Aerts (KU Leuven, BE), Nils Gehlenborg (Harvard University, US), Georgeta Elisabeta Marai (University of Illinois – Chicago, US), Kay Katja Nieselt (Universität Tübingen, DE)

April 15–20, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18161>

18171 – Normative Multi-Agent Systems

Mehdi Dastani (Utrecht University, NL), Jürgen Dix (TU Clausthal, DE), Harko Verhagen (Stockholm University, SE), Serena Villata (Laboratoire I3S – Sophia Antipolis, FR)

April 22–27, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18171>

18172 – Algebraic Effect Handlers go Mainstream

Sivaramakrishnan Krishnamoorthy Chandrasekaran (University of Cambridge, GB), Daan Leijen (Microsoft Research – Redmond, US), Matija Pretnar (University of Ljubljana, SI), Tom Schrijvers (KU Leuven, BE)

April 22–27, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18172>

18181 – Towards Accountable Systems

David Eysers (University of Otago, NZ), Christopher Millard (Queen Mary University of London, GB), Margo Seltzer (Harvard University – Cambridge, US), Jatinder Singh (University of Cambridge, GB)

April 29 to May 4, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18181>

18182 – Software Business, Platforms, and Ecosystems: Fundamentals of Software Production Research

Pekka Abrahamsson (University of Jyväskylä, FI), Jan Bosch (Chalmers University of Technology – Göteborg, SE), Sjaak Brinkkemper (Utrecht University, NL), Alexander Mädche (KIT – Karlsruher Institut für Technologie, DE)

April 29 to May 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18182>

18201 – Secure Compilation

Amal Ahmed (Northeastern University – Boston, US), Deepak Garg (MPI-SWS – Saarbrücken, DE), Catalin Hritcu (INRIA – Paris, FR), Frank Piessens (KU Leuven, BE)

May 13–18, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18201>

18202 – Inter-Vehicular Communication Towards Cooperative Driving

Onur Altintas (TOYOTA InfoTechnology Center USA – Mountain V, US), Suman Banerjee (University of Wisconsin – Madison, US), Falko Dressler (Universität Paderborn, DE), Geert Heijnen (University of Twente, NL), Katrin Sjöberg (Göteborg, SE)

May 13–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18202>

18211 – Formal Methods and Fault-Tolerant Distributed Computing: Forging an Alliance

Javier Esparza (TU München, DE), Pierre Fraigniaud (EPFL Lausanne, CH), Anca Muscholl (University of Bordeaux, FR), Sergio Rajsbaum (National Autonomous University of Mexico, MX)

May 21–25, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18211>

18212 – On-Body Interaction: Embodied Cognition Meets Sensor/Actuator Engineering to Design New Interfaces

Kasper Hornbaek (University of Copenhagen, DK), David Kirsh (University of California – San Diego, US), Joseph A. Paradiso (MIT – Cambridge, US), Jürgen Steimle (Universität des Saarlandes, DE)

May 21–24, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18212>

18231 – The Constraint Satisfaction Problem: Complexity and Approximability

Martin Grohe (RWTH Aachen, DE), Venkatesan Guruswami (Carnegie Mellon University – Pittsburgh, US), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU), Stanislav Živný (University of Oxford, GB)

June 3–8, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18231>

18241 – High-Performance Graph Algorithms

Henning Meyerhenke (HU Berlin, DE), Richard Peng (Georgia Institute of Technology – Atlanta, US), Ali Pinar (Sandia Nat. Labs – Livermore, US), Ilya Safro (Clemson University, US)

June 10–15, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18241>

18242 – Secure Routing for the Internet

Phillipa Gill (University of Massachusetts – Amherst, US), Amir Herzberg (University of Connecticut – Storrs, US), Adrian Perrig (ETH Zürich, CH), Matthias Wählisch (FU Berlin, DE)

June 10–13, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18242>

18251 – Database Architectures for Modern Hardware

Peter A. Boncz (CWI – Amsterdam, NL), Goetz Graefe (Google – Madison, US), Bingsheng He (National University of Singapore, SG), Kai-Uwe Sattler (TU Ilmenau, DE)

June 17–22, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18251>

18252 – Ubiquitous Gaze Sensing and Interaction

Lewis Chuang (LMU München, DE), Andrew Duchowski (Clemson University, US), Pernilla Qvarfordt (FX Palo Alto Laboratory, US), Daniel Weiskopf (Universität Stuttgart, DE)

June 18–21, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18252>

18261 – Discipline Convergence in Networked Systems

Yungang Bao (Chinese Academy of Sciences – Beijing, CN), Lars Eggert (NetApp Deutschland GmbH – Kirchheim, DE), Simon Peter (University of Texas – Austin, US), Noa Zilberman (University of Cambridge, GB)

June 24–29, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18261>

18271 – In Situ Visualization for Computational Science

Janine C. Bennett (Sandia National Labs – Albuquerque, US), Hank Childs (University of Oregon – Eugene, US), Christoph Garth (TU Kaiserslautern, DE), Bernd Hentschel (RWTH Aachen, DE)

July 1–6, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18271>

18281 – Synergies between Adaptive Analysis of Algorithms, Parameterized Complexity, Compressed Data Structures and Compressed Indices

Jérémy Barbay (University of Chile – Santiago de Chile, CL), Johannes Fischer (TU Dortmund, DE), Stefan Kratsch (HU Berlin, DE), Srinivasa Rao Satti (Seoul National University, KR)

July 8–13, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18281>

18291 – Extreme Classification

Samy Bengio (Google Inc. – Mountain View, US), Krzysztof Dembczynski (Poznan University of Technology, PL), Thorsten Joachims (Cornell University – Ithaca, US), Marius Kloft (TU Kaiserslautern, DE), Manik Varma (Microsoft Research India – Bangalore, IN)

July 15–20, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18291>

18321 – Web Application Security

Martin Johns (TU Braunschweig, DE), Nick Nikiforakis (Stony Brook University, US), Melanie Volkamer (KIT – Karlsruher Institut für Technologie, DE), John Wilander (Apple Computer Inc. – Cupertino, US)

August 5–8, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18321>

18322 – Human-Computer Integration

Jonathan Grudin (Microsoft Research – Redmond, US), Kristina Höök (KTH Royal Institute of Technology – Stockholm, SE), Pattie Maes (MIT – Cambridge, US), Florian Mueller (RMIT University – Melbourne, AU)

August 5–10, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18322>

18331 – Algorithmic Foundations of Programmable Matter

Spring Berman (Arizona State University – Tempe, US), Sándor Fekete (TU Braunschweig, DE), Matthew J. Patitz (University of Arkansas – Fayetteville, US), Christian Scheideler (Universität Paderborn, DE)

August 12–17, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18331>

18332 – Blockchain Technology for Collaborative Information Systems

Marlon Dumas (University of Tartu, EE), Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US), Jan Mendling (Wirtschaftsuniversität Wien, AT), Ingo Weber (Data61, CSIRO – Sydney, AU)

August 12–17, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18332>

18341 – Formalization of Mathematics in Type Theory

Andrej Bauer (University of Ljubljana, SI), Martín H. Escardó (University of Birmingham, GB), Peter L. Lumsdaine (University of Stockholm, SE), Assia Mahboubi (INRIA – Nantes, FR)

August 19–24, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18341>

18351 – Modeling for Sustainability

Gordon Blair (Lancaster University, GB), Betty H. C. Cheng (Michigan State University – East Lansing, US), Lorenz Hilty (Universität Zürich, CH), Richard F. Paige (University of York, GB)

August 26–31, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18351>

18361 – Measuring the Complexity of Computational Content: From Combinatorial Problems to Analysis

Vasco Brattka (Universität der Bundeswehr – München, DE), Damir D. Dzhafarov (University of Connecticut – Storrs, US), Alberto Marcone (University of Udine, IT), Arno Pauly (Swansea University, GB)

September 2–7, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18361>

18371 – Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web

Piero Andrea Bonatti (University of Naples, IT), Stefan Decker (RWTH Aachen, DE), Axel Polleres (Wirtschaftsuniversität Wien, AT), Valentina Presutti (CNR – Rome, IT)

September 9–14, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18371>

18381 – Quantum Programming Languages

Michele Mosca (University of Waterloo, CA), Martin Roetteler (Microsoft Corporation – Redmond, US), Peter Selinger (Dalhousie University – Halifax, CA)

September 16–21, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18381>

18391 – Algebraic Methods in Computational Complexity

Markus Bläser (Universität des Saarlandes, DE), Valentine Kabanets (Simon Fraser University – Burnaby, CA), Jacobo Torán (Universität Ulm, DE), Christopher Umans (Caltech – Pasadena, US)

September 23–28, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18391>

18401 – Automating Data Science

Tijl De Bie (Ghent University, BE), Luc De Raedt (KU Leuven, BE), Holger H. Hoos (Leiden University, NL), Padhraic Smyth (University of California – Irvine, US)

September 30 to October 5, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18401>

18411 – Progressive Data Analysis and Visualization

Jean-Daniel Fekete (INRIA Saclay – Orsay, FR), Danyel Fisher (Honeycomb – San Francisco, US), Arnab Nandi (Ohio State University – Columbus, US), Michael Sedlmair (Universität Stuttgart, DE)

October 7–12, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18411>

18412 – Encouraging Reproducibility in Scientific Research of the Internet

Vaibhav Bajpai (TU München, DE), Olivier Bonaventure (UC Louvain, BE), Kimberly Claffy (San Diego Supercomputer Center, US), Daniel Karrenberg (RIPE NCC – Amsterdam, NL)

October 7–10, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18412>

18421 – Algorithmic Enumeration: Output-sensitive, Input-Sensitive, Parameterized, Approximative

Henning Fernau (Universität Trier, DE), Petr A. Golovach (University of Bergen, NO), Dieter Kratsch (University of Lorraine – Metz, FR), Marie-France Sagot (University Claude Bernard – Lyon, FR)

October 14–19, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18421>

18422 – Shape Analysis: Euclidean, Discrete and Algebraic Geometric Methods

Michael Breuß (BTU Cottbus, DE), Alfred M. Bruckstein (Technion – Haifa, IL), Christer Oscar Kiselman (Uppsala University, SE), Petros Maragos (National Technical University of Athens, GR)
 October 14–19, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18422>

18431 – Computational Aspects of Fabrication

Marc Alexa (TU Berlin, DE), Bernd Bickel (IST Austria – Klosterneuburg, AT), Jessica K. Hodgins (Carnegie Mellon University – Pittsburgh, US), Kristina Shea (ETH Zürich, CH)
 October 21–26, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18431>

18441 – Data Physicalization

Jason Alexander (Lancaster University, GB), Petra Isenberg (INRIA Saclay – Orsay, FR), Yvonne Jansen (CNRS-Sorbonne University – Paris, FR), Bernice E. Rogowitz (Visual Perspectives – New York & Columbia University – New York, US), Andrew Vande Moere (KU Leuven, BE)
 October 28 to November 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18441>

18442 – Visualization and Processing of Anisotropy in Imaging, Geometry, and Astronomy

Andrea Fuster (TU Eindhoven, NL), Evren Özarslan (Linköping University, SE), Thomas Schultz (Universität Bonn, DE), Eugene Zhang (Oregon State University – Corvallis, US)
 October 28 to November 2, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18442>

18451 – Genomics, Pattern Avoidance, and Statistical Mechanics

Michael Albert (University of Otago, NZ), David Bevan (University of Strathclyde – Glasgow, GB), Miklós Bóna (University of Florida – Gainesville, US), István Miklós (Alfréd Rényi Institute of Mathematics – Budapest, HU)
 November 4–9, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18451>

18461 – Blockchain Security at Scale

Rainer Böhme (Universität Innsbruck, AT), Joseph Bonneau (New York University, US), Ittay Eyal (Technion – Haifa, IL)
 November 11–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18461>

18462 – Provenance and Logging for Sense Making

Jean-Daniel Fekete (INRIA Saclay – Orsay, FR), T. J. Jankun-Kelly (Mississippi State University, US), Melanie Tory (Tableau Software – Palo Alto, US), Kai Xu (Middlesex University – London, GB)
 November 11–16, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18462>

18471 – Next Generation Domain Specific Conceptual Modeling: Principles and Methods

Heinrich C. Mayr (Alpen-Adria-Universität Klagenfurt, AT), Sudha Ram (University of Arizona – Tucson, US), Wolfgang Reisig (HU Berlin, DE), Markus Stumptner (University of South Australia – Adelaide, AU)
 November 18–23, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18471>

18481 – High Throughput Connectomics

Moritz Helmstaedter (MPI for Brain Research – Frankfurt am Main, DE), Jeff Lichtman (Harvard University – Cambridge, US), Nir Shavit (MIT – Cambridge, US)
 November 25–30, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18481>

18482 – Network Visualization in the Humanities

Katy Börner (Indiana University – Bloomington, US), Oyvind Eide (Universität Köln, DE), Tamara Mchedlidze (KIT – Karlsruher Institut für Technologie, DE), Malte Rehbein (Universität Passau, DE), Gerek Scheuermann (Universität Leipzig, DE)
 November 25–30, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18482>

18491 – Multidirectional Transformations and Synchronisations

Anthony Cleve (University of Namur, BE), Ekkart Kindler (Technical University of Denmark – Lyngby, DK), Perdita Stevens (University of Edinburgh, GB), Vadim Zaytsev (RainCode – Brussels, BE)
 December 2–7, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18491>

18511 – Algebraic Coding Theory for Networks, Storage, and Security

Martin Bossert (Universität Ulm, DE), Eimear Byrne (University College Dublin, IE), Antonia Wachter-Zeh (TU München, DE)
 December 16–21, 2018 | Dagstuhl Seminar | <https://www.dagstuhl.de/18511>

Dagstuhl-Perspektiven-Workshops

14.2

Dagstuhl Perspectives Workshops

14

18262 – 10 Years of Web Science: Closing The Loop

Susan Halford (University of Southampton, GB), James A. Hendler (Rensselaer Polytechnic Institute – Troy, US), Eirini Ntoutsi (Leibniz Universität Hannover, DE), Steffen Staab (Universität Koblenz-Landau, DE)

June 24–29, 2018 | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/18262>

18472 – Implementing FAIR Data Infrastructures

Natalia Manola (University of Athens, GR), Peter Mutschke (GESIS – Köln, DE), Andrea Scharnhorst (Royal Netherlands Acad. of Arts & Sci. – Amsterdam, NL), Guido Scherp (ZBW-Leibniz-Informationszentrum Wirtschaft – Kiel, DE), Klaus Tochtermann (ZBW-Leibniz-Informationszentrum Wirtschaft – Kiel, DE), Peter Wittenburg (Max Planck Computing and Data Facility – Garching, DE)

November 18–21, 2018 | Dagstuhl Perspectives Workshop | <https://www.dagstuhl.de/18472>

GI-Dagstuhl-Seminare

14.3

GI-Dagstuhl Seminars

18283 – Visualizing Systems and Software Performance (VSSP)

Fabian Beck (Universität Duisburg – Essen, DE), Alexandre Bergel (University of Chile – Santiago de Chile, CL), Cor-Paul Bezemer (University of Alberta – Edmonton, CA), Katherine E. Isaacs (University of Arizona – Tucson, US)

July 8–13, 2018 | GI-Dagstuhl Seminar | <https://www.dagstuhl.de/18283>

18343 – Software Engineering for Intelligent and Autonomous Systems (SEFIAS)

Ada Diaconescu (Telecom ParisTech, FR), Simos Gerasimou (University of York, GB), Thomas Vogel (HU Berlin, DE)

August 19–24, 2018 | GI-Dagstuhl Seminar | <https://www.dagstuhl.de/18343>

Lehrveranstaltungen

14.4

Educational Events

18383 – Autumn School 2018 for Information Retrieval and Information Foraging

Ingo Frommholz (University of Bedfordshire – Luton, GB), Norbert Fuhr (Universität Duisburg-Essen, DE), Ralf Schenkel (Universität Trier, DE)

September 16–21, 2018 | Educational Event | <https://www.dagstuhl.de/18383>

18503 – Lehrerfortbildung in Informatik

Manuel García Mateos (Gymnasium am Steinwald – Neunkirchen, DE), Michael Gerke (Schloss Dagstuhl – Saarbrücken, DE), Martin Zimmol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)

December 12–14, 2018 | Educational Event | <https://www.dagstuhl.de/18503>

Forschungsgruppentreffen

14.5

Research Group Meetings

18029 – Forschungsaufenthalt

Carla Manni (University of Rome “Tor Vergata”, IT)

January 11–20, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18029>

18034 – Software Engineering Forschungsmethoden Training

Sven Apel (Universität Passau, DE), Eric Bodden (Universität Paderborn, DE), Lars Grunske (HU Berlin, DE)

January 14–17, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18034>

18043 – Digitalisierte Mobilität I

Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE)

January 22–23, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18043>

18044 – Retreat “SECUSO”

Melanie Volkamer (TU Darmstadt, DE)

January 21–24, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18044>**18045 – Digitalisierte Mobilität II**

Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE)

January 23–25, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18045>**18063 – Research Methods for Technology Enhanced Learning**

Marcus Specht (Open University – Heerlen, NL)

February 5–9, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18063>**18073 – Lehrstuhltreffen “Embedded Intelligence”**

Bernhard Sick (Universität Kassel, DE)

February 11–16, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18073>**18093 – SAW Projekttreffen**

Florian Reitz (Schloss Dagstuhl – Trier, DE)

February 28 to March 1, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18093>**18094 – Klausurtagung AG Robotersysteme TU KL**

Karsten Berns (TU Kaiserslautern, DE)

March 1–2, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18094>**18133 – GIBU 2018: GI-Beirat der Universitätsprofessoren**

Gregor Snelting (KIT – Karlsruher Institut für Technologie, DE)

March 25–27, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18133>**18143 – Retreat SFB 1102: Information Density and Linguistic Encoding**

Elke Teich (Universität des Saarlandes, DE)

April 3–6, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18143>**18144 – Klausurtagung Telematik Karlsruhe**

Robert Bauer (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)

April 4–6, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18144>**18163 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)**

Michaela Huhn (Ostfalia Hochschule – Wolfenbüttel, DE)

April 15–18, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18163>**18165 – iRRAM/MPFR/MPC Workshop**

Norbert T. Müller (Universität Trier, DE), Paul Zimmermann (INRIA Nancy – Grand Est, FR)

April 18–20, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18165>**18183 – Restructuring IEEE VIS for the Future**

Hans Hagen (TU Kaiserslautern, DE), Daniel A. Keim (Universität Konstanz, DE), Tamara Munzner (University of British Columbia – Vancouver, CA), Stephen North (Invisible – Oldwick, US), Hanspeter Pfister (Harvard University – Cambridge, US)

May 2–4, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18183>**18193 – Kolloquium zum GI Dissertationspreis 2017**

Steffen Hölldobler (TU Dresden, DE)

May 6–9, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18193>**18203 – Lehrstuhltreffen AG Zeller**

Andreas Zeller (Universität des Saarlandes, DE)

May 16–18, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18203>**18223 – Gemeinsamer Workshop der Graduiertenkollegs GRK 2167 und GRK 2193**

Ahmet Aker (Universität Duisburg-Essen, DE), Norbert Fuhr (Universität Duisburg-Essen, DE), Jakob Rehof (TU Dortmund, DE), David Scholz (TU Dortmund, DE), Sevda Tarkun (TU Dortmund, DE), Jan Winkels (TU Dortmund, DE)

May 27–30, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18223>**18233 – Workshop Buchprojekt “Applied Machine Intelligence”**

Thomas Hoppe (Datenlabor Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE), Anatol Reibold (OntoPort UG – Darmstadt, DE)

June 3–6, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18233>

18273 – GI Workshop Roboterarchitekturen

Karsten Berns (TU Kaiserslautern, DE)

July 3–4, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18273>**18293 – Klausurtagung**

Felix Freiling (Universität Erlangen-Nürnberg, DE)

July 15–18, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18293>**18294 – Deutsch-Pakistanischer Workshop**

Karsten Berns (TU Kaiserslautern, DE)

July 19–20, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18294>**18353 – Klausurtagung “LST Staab”**

Steffen Staab (Universität Koblenz-Landau, DE)

August 27–29, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18353>**18354 – Klausurtagung**

Marco Ragni (Universität Freiburg, DE)

August 26–31, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18354>**18363 – Lehrstuhltreffen Rechtsinformatik**

Christoph Sorge (Universität des Saarlandes, DE)

September 6–7, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18363>**18373 – TRR 195 Blockseminar über Spieltheorie**

Simon Schmidt (Universität des Saarlandes, DE), Andrea Thevis (Universität des Saarlandes, DE)

September 10–13, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18373>**18374 – Klausurtagung**

Armin Heinzl (Universität Mannheim, DE)

September 11–12, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18374>**18393 – Data Documentation Initiative (DDI) – Train the Trainers**

Jon Johnson (UK Data Service – Colchester, GB), Dan Smith (Colectica – Minneapolis, US), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)

September 23–28, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18393>**18403 – Interoperability of Metadata Standards in Cross-Domain Science, Health, and Social Science Applications**

Simon Cox (CSIRO – South Clayton, AU), Simon Hodson (CODATA – Paris, FR), Steven McEachern (Australian National University – Canberra, AU), Joachim Wackerow (GESIS – Mannheim, DE)

September 30 to October 5, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18403>**18433 – Arbeitstreffen “Challenges in Computer Aided Engineering”**

Bernd Becker (Universität Freiburg, DE)

October 21–24, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18433>**18434 – Klausurtagung Rossow**

Christian Rossow (Universität des Saarlandes, DE)

October 24–26, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18434>**18448 – Forschungsaufenthalt**

Brygg Ullmer (Clemson University, US)

November 3–4, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18448>**18449 – Forschungsaufenthalt**

Tobias Isenberg (INRIA Saclay – Orsay, FR)

October 28 to November 2, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18449>**18453 – Retreat for the Research Training Group “Adaptive Information Preparation from Heterogeneous Sources (AIPHES)”**

Johannes Fürnkranz (TU Darmstadt, DE), Stefan Roth (TU Darmstadt, DE)

November 4–7, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18453>**18473 – Klausurtagung des Forschungsprojekts Continuity: “Automated Performance Testing in Continuous Software Engineering”**

Stefan Siegl (Novatec Consulting GmbH – Leinfelden-Echterdingen, DE), André van Hoorn (Universität Stuttgart, DE)

November 21–23, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18473>

18474 – Klausurtagung

Jürgen Steimle (Universität des Saarlandes, DE)

November 22–23, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18474>**18493 – Digitalisierte Mobilität III**

Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE)

December 3–4, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18493>**18494 – VDV FhG Innovationsworkshop**

Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE), Christoph Terwelp (Fraunhofer FIT – Sankt Augustin, DE)

December 4–5, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18494>**18504 – Secan Lab Seminar**

Thomas Engel (University of Luxembourg, LU)

December 10–11, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18504>**18505 – “inf-schule.de: Weiterentwicklung des elektronischen und interaktiven Lehrbuches für den Informatikunterricht”**

Martin Zimmol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)

December 11–12, 2018 | Research Group Meeting | <https://www.dagstuhl.de/18505>





© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2018 | ISSN 2199-1995
<http://www.dagstuhl.de>