SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2016

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2016

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	Dr. Marcel R. Ackermann, Dr. Andreas Dolzmann, Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Dr. Roswitha Bardohl, Heike Clemens, Dr. Andreas Dolzmann, Jutka Gasiorowski, Dr. Marc Herbstritt, Heike Meißner, Petra Meyer, Thomas Schillo, Christina Schwarz, 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	Kern GmbH, Bexbach
Druckauflage	175
ISSN	2199-1995
Online-Version	http://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	April 2017
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	Creative Commons BY-ND 3.0 DE

Leibniz-Gemeinschaft

Zertifikat seit 2012 audit berufundfamilie

Vorwort
Foreword

2016 war ein besonderes Jahr für uns. Schloss Dagstuhl – Leibniz-Zentrum für Informatik wurde von der Leibniz-Gemeinschaft evaluiert. Das ist ein formaler Vorgang, dem sich jedes Institut der Leibniz-Gemeinschaft spätestens alle sieben Jahre unterziehen muss, um sicherzustellen, dass es nach wie vor die Bedingungen für eine gemeinsame Förderung durch Bund und Länder erfüllt.

Wir können jetzt sagen, dass wir diese Überprüfung mit wehenden Fahnen bestanden haben. Der Bericht unserer Evaluierungskommission, den sich auch der Leibniz-Senat zu eigen gemacht hat, ist für uns äußerst günstig, zum Teil sogar voll Begeisterung. Er beurteilt unsere drei Bereiche und benotet *Seminare und Workshops* und auch *Bibliographische Dienste* mit "exzellent" und benotet *Open-Access-Publikationswesen* mit "sehr gut". Dieser Bericht ist öffentlich und kann bei der Seite der Leibniz-Gemeinschaft eingesehen werden.

Ich möchte allen unseren Mitarbeitern herzlichst danken, wie auch den Mitgliedern unserer Kommissionen und allen anderen, die sich so sehr engagiert haben, Dagstuhl zu dem zu machen, was es heute ist, und unserem Insitut zu seinem Ruf verholfen haben. Insbesondere gilt großer Dank meinem Vorgänger, dem langjährigen Direktor, Reinhard Wilhelm.

Was unsere Gesamtaufgabe betrifft, war 2016 wieder ein sehr erfolgreiches Jahr. Wir hatten eine interessante Mischung an Seminarthemen, von reiner Informatik wie z.B. *Symmetric Cryptography* hin zu Anwendungen in den Ingenieurswissenschaften, z.B. *Symbolic-Numeric Methods for Reliable and Trustworthy Problem Solving in Cyber-Physical Domains*, oder auch in der Kunst wie z.B. *Computational Music Structure Analysis*. Aber dazu gab es noch Themen, die die Informatik im Kontext der Gesellschaft betrachten, wie z.B. *Engineering Moral*

2016 was an extraordinary year for us. Schloss Dagstuhl – Leibniz Zentrum für Informatik was evaluated by the Leibniz Association. This is a formal process that is applied to every institution of the Leibniz Association every seven years, at the latest, to ensure that it still fulfills the prerequisites for joint funding by the German federal and state governments.

Now we can proudly say that we passed this evaluation with flying colors. The report issued by our evaluation committee that was officially adopted by the Leibniz Senate is extremely favorable, at times enthusiastic. It gives "grades" to our three divisions and rates *Seminars and Workshops* and also *Bibliographic Services* as *excellent*, and rates *Open Access Publishing* as *very good*. The report is publicly available and can be found at the website of the Leibniz Association.

I would like to thank our dedicated staff, the external members of our committees, and all others who have helped to make Dagstuhl to the place that it is and that is perceived so positively. In particular I want to thank my predecessor, the long-time director Reinhard Wilhelm.

In terms of our mission, 2016 was again a very successful year. We had an interesting mix of seminar topics ranging from pure computer science topics such as *Symmetric Cryptography* to applications in engineering, e.g. *Symbolic-Numeric Methods for Reliable and Trustworthy Problem Solving in Cyber-Physical Domains* and in the arts, e.g. *Computational Music Structure Analysis*, but also including topics placing computing in the context of society, e.g. *Engineering Moral Agents – from Human Morality to Artifcial Morality* or *Data, Responsibly*. The dblp database and services continued to grow and our open-access publishing efforts have seen great acceptance, in particular with its *LIPIcs* series.

Agents – from Human Morality to Artifcial Morality oder *Data, Responsibly*. Die dblp Datenbank und ihre Dienste sahen weiteres Wachstum, und unsere Anstrengungen im Open-Access-Publikationswesen werden immer besser angenommen, insbesondere die *LIPIcs* Serie.

Natürlich ist noch einiges mehr passiert. Sie können es in diesem Bericht finden. Ich hoffe, Sie haben Freude daran.

Raimund Seidel

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

Of course many more things have happened. You can find them in this report. I hope you will find some enjoyment in them.

On behalf of the Managing Directors

Inhaltsverzeichnis
Contents

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.

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.

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 Ingenieurswissenschaften 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 jährlich mehr als 3 000 internationale Gäste.

Seit 2005 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 91 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.

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 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 91 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- 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

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 Januar 2017.
As of January 2017.

intensiv über ihre aktuelle Forschung zu diskutieren. Darüber hinaus trifft sich in Dagstuhl Perspektiven Workshops ein 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 im 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

Seit 2011 betreibt Schloss Dagstuhl in enger Zusammenarbeit mit der Universität Trier die Bibliographiedatenbank dblp, welche mit mittlerweile mehr als drei Millionen Publikationseinträgen die weltweit größte, offene Sammlung bibliographischer Daten in der Informatik ist. 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.

dblp computer science bibliography

Since 2011, Schloss Dagstuhl operates the dblp computer science bibliography in close cooperation with the University of Trier. Listing about three 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

Die Förderung der Kommunikation zwischen den Wissenschaftlern in der Informatik gehört zu der zentralen Aufgabe von Schloss Dagstuhl. Wissenschaftliche Veröf-

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

fentlichungen 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.

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 2016 — 1.2 — *News from 2016*

Evaluierung

Die Leibniz-Gemeinschaft evaluiert spätestens alle sieben Jahre ihre Einrichtungen und überprüft, ob sie den Kriterien für eine weitere Förderung im Rahmen der Bund-Länder-Finanzierung genügen. In 2016 stand die Evaluierung für Schloss Dagstuhl an. Eine diesbezügliche Begehung durch das internationale Gutachterkomitee fand im Juli statt. Die Veröffentlichung des Berichts der Evaluierungskommision, die Stellungnahme des Leibniz Senats sowie eine Überprüfung der Fördervoraussetzung durch die Gemeinsame Wissenschaftskonferenz (GWK) werden im ersten Halbjahr 2017 erwartet.

Evaluation

The Leibniz Association evaluates its institutions every seven years at the latest to ensure that they still fullfil the prerequisites for a joint funding by the German federal and state governments. In 2016, Schloss Dagstuhl has been evaluated. As part of this evaluation, an international evaluation committee visited Schloss Dagstuhl in July 2016. The report of the evaluation committee, a statement by the Leibniz Senate, and the assessment of the funding eligibility by the Joint Science Conference GWK are expected in the first half of 2017.

Das Team

Nahezu alle Mitarbeiter von Schloss Dagstuhl wurden 2016 über den Kernhaushalt des Zentrums bezahlt. Eine Ausnahme bildet das dblp-Team. Die Klaus Tschira Stiftung unterstützte Schloss Dagstuhl und dblp wie in den vergangenen Jahren auch 2016 mit einer großzügigen Spende von inzwischen 71 500 €. Zwei Mitarbeiter wurden im Rahmen des seit Juli 2015 laufenden Projektes „Skalierbare Autoren-Disambiguierung in Literaturdatenbanken" aus Fördermitteln des Leibniz-Wettbewerbes finanziert.

Schloss Dagstuhl beschäftigte 2016 in der Küche insgesamt drei Auszubildende. Eine schloss ihre Prüfung im August 2016 erfolgreich ab und wurde für zunächst ein Jahr übernommen. Ein im Februar 2016 eingestellter Praktikant begann im August 2016 seine Ausbildung in der Küche. Daneben arbeitete im Sommer 2016 ein Praktikant für 2 Wochen in der IT-Abteilung. Er fing an, die Gebäude des Tagungszentrum im Open-World-Computerspiel Minecraft nachzubilden.

Ende 2016 beschäftigte Schloss Dagstuhl insgesamt 33,86 Vollzeitäquivalente bzw. 49 Angestellte.

The Team

Nearly all staff at Schloss Dagstuhl were funded from the center's core budget in 2016. An exception is the dblp team, where – as in the previous years – one position was supported by a generous donation in the amount of 71,500 € from the Klaus Tschira Foundation. Additionally, two dblp staff members were financed by a project grant of the Leibniz Competition (project "Scalable Author Disambiguation for Bibliographic Databases").

Schloss Dagstuhl's kitchen employs three trainees in total. One who had started in 2016 finished her training in August 2016 and entered a one-year contract. An intern who had started in February began his training in August 2016. Additionally, the IT department hosted an intern for 2 weeks during the summer. He started to clone the center's buildings in the sandbox video game Minecraft.

At the end of 2016, Schloss Dagstuhl had a total of 49 staff members corresponding to 33.86 full-time positions.

Seminare und Workshops

Mit 126 Anträgen auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wurden mehr Anträge denn je eingereicht. Dies wird 2017 wieder zu einem leichten Anstieg der Veranstaltungszahl in diesen Kategorien führen. Im Berichtsjahr gab es wie auch 2015 mit 72 Seminaren und Workshops etwas weniger Veranstaltungen als das bisher zweimal erreichte Maximum von 75 pro Jahr.

Von den mehr als 3300 Gästen, die sich in Dagstuhl trafen, nahmen etwa 2400 an Seminaren teil. Etwas weniger als die Hälfte 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. Diese positiven Kennzahlen liegen bis auf vernachlässigbare Schwankungen im Bereich der Ergebnisse der letzen Jahre.

Bezüglich der Geschlechterverteilung gab es 2016 leicht bessere Ergebnisse als in den vorherigen Jahren. Mehr als 69 % aller Seminare hatten mindestens eine Frau im drei- bis fünfköpfigen Organisatorenteam, und fast ein Viertel aller Organisatoren waren Frauen. Der Frauenanteil unter allen Seminarteilnehmern war mit 18 % ebenfalls höher als in den vergangenen Jahren. Mehr Details und Zahlen zum Seminarprogramm finden sich in Kapitel 2.

Seminars and Workshops

126 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops were submitted, more than ever before. This will lead to a slight increase in these event categories in 2017. Just as 2015, the year under review saw 72 seminars and workshops, slightly fewer than the maximum of 75 that has been reached twice before.

2400 of the 3300 guests hosted at Dagstuhl participated in seminars. A little less than half of all seminar participants came to Dagstuhl for the first time and more than a third of all participants in our guest survey classified themselves as junior researchers. More than 75 % of all seminar participants were affiliated outside Germany. Apart from negligible deviations, these positive key figures are in the same range as the results of previous years.

Regarding gender distribution, the year under review saw slightly better results than previous years. More than 69 % of all seminars had at least one woman organizer in organizer teams of three to five scientists, whereas almost 25 % of all organizers were women. The ratio of women to all seminar participants was 18 %, i.e. also higher than in previous years. See Chapter 2 fore more details.

Bibliographiedatenbank dblp

Im Jahr 2016 konnte die *dblp computer science bibliographie* das erste mal die Marke von 400 000 neu aufgenommenen Publikationen überschreiten. Dies entspricht mehr als 1 500 neuen Publikationen pro Arbeitstag. Ende 2016 indexierte dblp somit bereits über 3,6 Millionen Fachartikel aus den verschiedenen Teilgebieten der Informatik.

Die Nutzung des dblp-Dienstes blieb dabei auf konstant hohem Niveau. Jeden Monat verzeichnet die dblp-Webseite bis zu 17 Millionen Seitenzugriffe von etwa 450 000 verschiedenen Nutzern aus aller Welt. Dies entspricht mehr als sechs Seitenzugriffen pro Sekunde; zu Spitzenzeiten erfolgen über 80 Zugriffe gleichzeitig. Im Durchschnitt beginnt etwa alle drei Sekunden ein neuer Nutzer, mit dblp zu arbeiten.

2016 konnten Metadaten von mehr als 9 000 Dissertationen der französischen Datenbank *Hyper Articles en Ligne (HAL)* sowie 10 000 Dissertationen des *EThOS*-Dienstes der britischen Nationalbibliothek in dblp integriert werden. Diese Datensätze umfassen ein Großteil der französischen und britischen Dissertationen seit den 1950er Jahren und ergänzen die über 20 000 bereits vorhandenen Datensätzen der *Deutschen Nationalbibliothek (DNB)*.

Mehr Informationen zu dblp finden sich in Kapitel 3.

dblp computer science bibliography

Between January 1, 2016, and December 31, 2016, the dblp database grew by more than 400,000 publication records to reach a total of more than 3.6 million records. This corresponds to more than 1,500 new records for each working day of the year.

Up to 17 million web pages are visited each month by about 450,000 researchers and computer science enthusiasts all over the world. On average, about six web pages are requested from the dblp web servers in every second; at peak times, as many as 80 request are made concurrently. About every three seconds, a new user session is started.

In 2016, metadata of more than 9,000 PhD theses from the French archive *Hyper Articles en Ligne (HAL)* and more than 10,000 PhD theses from the *EThOS Repository* of the British Library have been imported into dblp. These theses cover a significant part of the French and British computer science community, with some theses reaching back as far as the 1950s. These theses complement the more than 20,000 theses already added from the *Deutsche Nationalbibliothek (DNB)*, the German National Library.

More information about dblp can be found in Chapter 3.

Dagstuhl Publishing

Auch in 2016 haben die Open-Access-Publikationsaktivitäten starken Zuspruch bekommen. So wurden in den beiden Konferenzbandreihen LIPIcs und OASIcs zusammen erstmals über 1 000 Publikationen innerhalb eines Jahres veröffentlicht. Zudem gab es auch in 2016 wieder

Dagstuhl Publishing

Schloss Dagstuhl's open-access publishing services experienced an on-going strong increase in demand from the community in 2016. For the first time, more than 1,000 articles have been published within one year in the two conference proceedings series LIPIcs and OASIcs.

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.

Öffentlichkeitsarbeit und Weiterbildung

Am 28. Oktober 2016 fanden die ersten „Dagstuhler Gespräche" statt, eine gemeinsame Veranstaltung von Schloss Dagstuhl und der Stadt Wadern. 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. Der Dialog zwischen Entscheidern und Gestaltern aus Wirtschaft und Politik sowie der interessierten Öffentlichkeit wurde durch den Impulsvortrag „Wenn Schweine schwitzen" von Prof. Holger Hermanns angeregt und fand starken Anklang.

Unter dem Motto „Schreiben über Informatik" nahmen unter der Leitung der Wissenschaftsjournalisten Tim Schröder und Gordon Bolduan 12 Volontäre und mit Öffentlichkeitsarbeit betraute Mitarbeiter aus Organisationen teil.

In Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisierte Schloss Dagstuhl 2016 zum 26. mal eine Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet.

Mehr Informationen zur Öffentlichkeitsarbeit und zu den Weiterbildungsaktivitäten finden sich in Kapitel 7.

Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch im Jahr 2016 gab es wieder eine Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum[2] (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 2016 erhielten in der Woche vor der vierten Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Public-Key Cryptography" (16371). Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2017 zugestimmt.

[2] http://www.heidelberg-laureate-forum.org

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.

Public Relations and Professional Training

On October 28, 2016, the first session of "Dagstuhler Gespräche" (Dagstuhl Talks) took place, an event organized jointly by Schloss Dagstuhl and the town 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. Prof. Holger Hermann's kick-off talk "Wenn Schweine schwitzen" (When Pigs Perspire) achieved this goal and was very well received.

There were also 12 trainee journalists and PR staff members from different organizations who, led by science journalists Tim Schröder and Gordon Bolduan, attended under the motto "Writing about Computer Science".

In 2016, Schloss Dagstuhl hosted its teacher training course for the 26th 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).

Further details about public relations and professional training at Schloss Dagstuhl can be found in Chapter 7.

Joint Outreach with the Heidelberg Laureate Forum

2016 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum[2] (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 "Public-Key Cryptography" (16371), taking place during the week before the fourth edition of the forum. Satisfied with the outstanding success of the initiative, both partners agreed to continue the cooperation in 2017.

Spender und Förderer

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

2015 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 684 Bände als Spende, davon 618 Monographien des Springer-Verlags. 23 Bände wurden der Bibliothek von Gästen und Forschern überlassen.

Wie in den vergangenen fünf Jahren förderte die Klaus Tschira Stiftung auch in diesem Jahr die Bibliographiedatenbank dblp mit einer Spende von 71 500 €.

Das Heidelberger Institut für Theoretische Studien (HITS) hat 2016 Dagstuhl Publishing mit 45 000 € unterstützt. Für 2017 und 2018 sind weitere Spenden in einer Gesamthöhe von 66 000 € zugesagt worden.

Sponsors and Donors

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generous donations for the 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 684, including 618 monographs from the Springer publishing house. 23 books were donated by guests and researchers.

As in the previous five years, Schloss Dagstuhl was grateful to receive a grant of now 71,500 € from the Klaus Tschira Foundation in support of the dblp computer science bibliography in 2016.

The Heidelberg Institute for Theoretical Studies (HITS) supported Dagstuhl Publishing with 45,000 €. Further support with a total amount of 66,000 € has been confirmed for the years 2017 and 2018.

NSF Förderung von Nachwuchswissenschaftlern

Schloss Dagstuhl hatte im August 2012 bei der National Science Foundation (NSF) in den USA einen Antrag auf Förderung junger Wissenschaftler aus den USA gestellt. Unter dem Titel „Schloss Dagstuhl – NSF Support Grant for Junior Researchers“ [3] wurde der Antrag mit Förderungsbeginn am 1. Oktober 2013 für drei Jahre genehmigt. Da nicht alle bereitgestellten Mittel im Förderzeitraum bis September 2016 abgerufen wurden, wurde die Förderung bis September 2017 verlängert. Im Berichtsjahr konnte durch die Förderung 45 Forschern aus den USA eine Teilnahme an insgesamt 36 Seminaren ermöglicht werden. Insgesamt wurden dafür 48 950 $ Fördermittel ausgegeben.

NSF Grant for Junior Researchers

In August 2012, Schloss Dagstuhl applied to the National Science Foundation (NSF), USA, for support for junior researchers working in the United States. The application was approved, and the NSF Grant for Junior Researchers was established for a duration of three years, effective from October 1, 2013. In 2016, 45 US-based scientists were supported with a total amount of 48,950 $ and hence able to participate in overall 36 Dagstuhl Seminars.

Baumaßnahmen und Renovierung

Schloss Dagstuhl hält die verschiedenen Gebäude des Zentrums laufend instand und modernisiert sie. Als größte Maßnahme wurde 2016 damit begonnen, einen zweiten großen Hörsaal zu schaffen. Im Rahmen eines im August begonnenen Umbaus werden im 1993 eröffneten Teil des Seminarzentrums der kleinste Hörsaal und der Computerraum zu einem neuen großen Hörsaal zusammengelegt. Dieser wird den zweitgrößten und fast wöchentlich genutzten Hörsaal „Kaiserslautern“ ersetzen und bezüglich der Größe als auch der technischen Ausstattung den aktuellen Anforderungen genügen.

Ende Oktober wurde südlich des derzeitigen provisorischen Parkplatzes mit dem Bau eines neuen Parkplatzes für Gäste und Mitarbeiter begonnen. Bis auf den Bau einer Umgrenzung und die Installation der Beleuchtung wurde der Parkplatz noch Ende 2016 fertig gestellt.

In der Bibliothek wurde ein Treppenlift installiert, um auch das oberste Geschoss des Gebäudes Menschen mit Einschränkungen zugänglich zu machen. Weiterhin wurden im Rahmen der Sanierung aller Bäder des Altbaus

Construction Work and Renovation

Schloss Dagstuhl continually maintains and modernizes all of the center's buildings. The biggest project in 2016 was the start on construction works related to the creation of a second large lecture hall. Renovation work started in August to merge our smallest meeting room, "Karlsruhe," and the adjacent computer room to a new, large lecture hall. It will replace the second largest, very frequently used lecture hall "Kaiserslautern" and will meet current requirements, both in terms of size and technical equipment.

In late October, construction on a new parking lot for guests and employees started south of the current, provisional one. Apart from boundary and lighting, the lot was completed by the end of 2016.

A stair lift was installed in the library so as to make the top floor accessible for people with disabilities. Furthermore, as part of the refurbishment of all bathrooms in the old building, two further bathrooms were completely refurbished. Other, smaller construction projects contributed to the conservation of the building stock and fire safety.

[3] dt.: „Schloss Dagstuhl – NSF Unterstützung für Nachwuchswissenschaftler“.

zwei weitere Bäder von Grund auf saniert. Eine Reihe kleinerer Baumaßnahmen dienten dem Brandschutz und der Sicherung des Baubestandes.

Im Speisesaal wurde die Theke für das kalte und warme Buffet erneuert. Die neue Theke entspricht nun den Vorgaben des HACCP-Konzepts. Ebenso wurde der gestiegenen durchschnittlichen Anzahl der Gäste, die mittags bewirtet werden, durch eine zusätzliche lange Theke Rechnung getragen. Im Rahmen des Umbaus wurde außerdem das Kühlaggregat in den Außenbereich verlagert.

In Zusammenarbeit mit der Universität des Saarlandes wurde der Kopierraum am Standort Saarbrücken grundlegend saniert und anschließend mit einem höhenverstellbaren Schreibtisch ausgestattet. Dies erleichtert den Mitarbeitern Arbeiten wie Frankieren, Schneiden und Binden.

Ausstattung

Neben diesen größeren Maßnahmen und Neuerungen hat Schloss Dagstuhl auch darauf geachtet, weitere Maßnahmen zur Verbesserung des Komforts und Ambiente umzusetzen. In weiteren Gästezimmern wurden die bisherigen Stühle mit Korbgeflecht durch aktuelle Schwingstühle mit Lederbezug ersetzt.

Ein Seminarraum wurde mit 12 höhenverstellbaren und rollbaren Tischen ausgestattet, die zudem platzsparend ineinandergeschoben werden können. Dies ermöglicht es Arbeitsgruppen, die Möblierung rasch ihren eigenen Bedürfnisse anzupassen.

In der Geschäftsstelle wurden vier Arbeitsplätze mit ergonomischen Bürostühlen ausgestattet. Weiterhin wurde nach Kündigung von Seiten der Deutschen Post der Freistempler durch eine aktuelle Frankiermaschine ersetzt.

IT Services

Das Jahr 2016 war für die IT geprägt durch größere Umbauarbeiten bedingt durch die Bauarbeiten an dem neuen Hörsaal. Um den Gästen einen durchgehenden Betrieb der IT-Infrastrukturen zu gewährleisten, musste der zentrale Serverraum erst provisorisch in einen anderen Raum umziehen, ehe er am aktuellen Standort neu aufgebaut werden konnte. Weiterhin setzt die IT, wie bereits im vorherigen Jahr begonnen, vermehrt den Raspberry Pi Kleinstcomputer ein. Im Hörsaal „Kaiserlsautern“ kommt ein solcher jetzt zusammen mit einer von der IT selbst entwickelten Software zur Steuerung des Beamers zum Einsatz. Das Infoterminal vor diesem Hörsaal wurde ersetzt und wird ebenfalls von einem Raspberry Pi gesteuert. Schließlich wurden die Thin-Clients an den schon vorhandenen Recherche-Arbeitsplätzen in der Bibliothek und die beiden Arbeitsplätze am Check-In durch Raspberry Pi-basierte Lösungen ersetzt.

Bibliothek

Zu den meisten Dagstuhl-Seminaren stellt Schloss Dagstuhl alle in der Bibliothek vorhandenen Bücher der anwesenden Gäste zusammen und präsentiert sie in einer separaten Auslage. So wurden 2016 zu 59 Dagstuhl-Seminaren ingesamt 2667 Bücher präsentiert, das entspricht

In the dining hall, the counter used for the cold and hot buffet was replaced, and now complies with HACCP regulations. Moreover, an additional counter was installed in order to accommodate the increasing average number of guests hosted at lunchtime. As part of the renovation, the cooling unit was relocated outside the building.

In cooperation with Saarland University, the copy room in the Dagstuhl Office, Saarbrücken, was refurbished extensively and equipped with a height-adjustable desk. This simplifies tasks like franking mail, cutting, and binding.

Facilities

Aside from these large-scale construction projects and renovations, Schloss Dagstuhl took further steps to modernize its facilities in order to enhance comfort and atmosphere. More guest rooms were equipped with up-to-date, leather-upholstered swinger chairs, replacing the old wicker chairs.

One seminar room was equipped with 12 height-adjustable, rollable, and nestable desks. This enables work groups to individually adjust the room's layout according to their needs.

In the Dagstuhl Office, four work stations were equipped with ergonomic office chairs. Besides, due to the termination of contract on the part of Deutsche Post AG, the old franking machine was replaced with an up-to-date model.

IT Services

For the IT department, 2016 was defined by large-scale renovations due to the construction works related to the new lecture hall. In order to ensure the guests' continued access to Dagstuhl's IT infrastructure, the main server room had to be moved provisionally, only to be rebuilt at its current, permanent location. In continuation of the development in 2015, the IT department increasingly utilizes Raspberry Pi minicomputers. The lecture hall "Kaiserslautern" was equipped with one, along with software developed by our IT department for controlling the computer projector. The information terminal in front of the hall was replaced and is now also operated by a Raspberry Pi. Finally, the thin clients at the research work stations in the library and at the check-in were also replaced with solutions based on the Raspberry Pi.

Library

For most Dagstuhl Seminars, books available in our library that were written by seminar participants are displayed separately. In 2016, 2667 books were displayed for 59 seminars, on average that corresponds to more than 45 books per seminar.

durchschnittlich mehr als 45 Büchern pro Dagstuhl-Seminar.

Die umfangreiche Sammlung „Dagstuhl's Impact"[4], die publizierte Ergebnisse aus Dagstuhl-Seminaren dokumentiert, wurde 2016 um zahlreiche weitere Publikationen ergänzt: 89 Artikel, 12 Zeitschriftensonderhefte und 8 Bücher konnten hinzugefügt werden.

Im Rahmen des Projekts „DEAL – Bundesweite Lizenzierung von Angeboten großer Wissenschaftsverlage", das durch die Allianz der deutschen Wissenschaftsorganisationen auf Anregung der HRK ins Leben gerufen wurde, hat die Bibliothek gemeinsam mit 60 weiteren bundesweiten Wissenschaftseinrichtungen zum Ende 2016 die laufenden Elsevier-Abonnements gekündigt. Der Ausgang der DEAL-Verhandlungen mit Elsevier ist noch offen, sodass davon auszugehen ist, dass 2017 zunächst kein Zugriff auf laufende Elsevier-Zeitschriften zur Verfügung steht.

Kunst

Seit 1995 finden in Schloss Dagstuhl regelmäßig Kunstausstellungen statt. Bisher wurden die Ausstellungen meist von einem einzelnen Künstler oder manchmal auch von einer kleinen Gruppe von Künstlern gestaltet. Im Berichtsjahr wurde nun eine Kooperation zwischen Saartoto, der HBKsaar und Schloss Dagstuhl begründet, in deren Rahmen die sonst unzugänglichen Werke des Kunstförderers Saartoto zeitlich beschränkt in Dagstuhl ausgestellt werden. Ungewöhnlich war 2016 ebenso die Ausstellung von Werken eines privaten Kunstsammlers. Weitere Informationen über das generelle Konzept, die Ausstellungen 2016 und die neugestartete Zusammenarbeit finden sich in Kapitel 10.

In 2016, the comprehensive collection "Dagstuhl's Impact", which documents published results of Dagstuhl Seminars, grew substantially: 89 articles, 12 special journal issues and 8 books were added.

As part of the project "DEAL – Bundesweite Lizenzierung von Angeboten großer Wissenschaftsverlage" (nationwide licensing of offers from major publishers of scientific books), created by the Alliance of Science Organisations in Germany at the suggestion of the German Rectors' Conference, Dagstuhl's library and 60 other nationwide scientific institutions discontinued the current Elsevier subscriptions as of the end of 2016. The DEAL negotiations with Elsevier have not been concluded, so it must be assumed for the time being that in 2017, current Elsevier journals will not be available.

Art

Since 1995, Schloss Dagstuhl has hosted art exhibitions on a regular basis. Until now, the exhibitions were organized by individual artists or sometimes small groups of artists. The year under review saw the establishment of a cooperation between Saartoto, HBKsaar, and Schloss Dagstuhl. Due to this cooperation, Schloss Dagstuhl was able to exhibit otherwise inaccessible works of art in the possession of art sponsor Saartoto for a limited period of time. Another unusual event was the exhibition of works owned by a private art collector. Further information about the art program in general, exhibitions in 2016 as well as the newly established cooperation can be found in Chapter 10.

[4] http://www.dagstuhl.de/bibliothek/dagstuhls-impact/

Fig. 1.1
Aerial photographies of Schloss Dagstuhl.

2 Seminare und Workshops
Seminars and Workshops

2.1 Dagstuhl-Seminare

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 2016 stattfanden, aufgelistet. Auf der Dagstuhl-Website ist das Programm der kommenden 24 Monate verfügbar.

2.1 *Dagstuhl Seminars*

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 2016 Seminars and Perspectives Workshops. Chapter 14 provide 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] http://www.dagstuhl.de/dagrep/

Dagstuhl-Perspektiven-Workshops

2.2

Dagstuhl Perspectives Workshops

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 2016 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]

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 2016 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].

Past and upcoming Dagstuhl Perspectives Workshop can be found on our web site.[7]

[6] http://www.dagstuhl.de/dagman
[7] http://www.dagstuhl.de/pw-list

Foundations of Data Management http://www.dagstuhl.de/16151
Engineering Academic Software http://www.dagstuhl.de/16252
Tensor Computing for Internet of Things http://www.dagstuhl.de/16152
QoE Vadis? http://www.dagstuhl.de/16472

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2016.

2.3 Einreichung der Anträge und Begutachtungsverfahren

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch 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 ausgeglichenen Repräsentation wissenschaftlicher Gemeinden, geographischer Regionen und besonders auf das Miteinbeziehen junger und weiblicher Wissenschaftler geachtet.

Die Informatikforscher zeigten 2016 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 125 Anträgen in den Antragsrunden im Januar und Juni 2016. In keinem Jahr zuvor gingen mehr Anträge ein. Der hohen Qualität der Anträge entsprechend, wurden etwa 63 % der eingereichten Anträge genehmigt. In den vergangenen 7 Jahren variierte die Rate der angenommen Anträge zwischen 63 % und 76 % (siehe Fig. 2.2).

Unter den 79 in 2016 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 13 Seminare bereits 2016 ausgerichtet werden, der Großteil wurde jedoch für das Seminar-Programm in 2017 eingeplant. Nur fünf der 2016 genehmigten Seminare werden – oft auf ausdrücklichen Wunsch der Organisatoren – in 2018 stattfinden.

2.3 *Proposal Submission and Review Process*

Schloss Dagstuhl maintains the high quality of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop series by identifying those proposals that promise a high potential to engage researchers – often from different disciplines – in scientific discussion 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, when the selection is made.

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 2016, submitting 125 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2016 and June 2016 submission rounds. There were never before so many proposals submissions. The quality of the proposals was excellent, resulting in a 63 % acceptance rate by Dagstuhl's Scientific Directorate. In the previous seven years, proposal acceptance rates have tended to range between 63 % and 76 % (see Fig. 2.2).

Among the 79 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2015 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 2017 seminar program, although it was possible to schedule 13 of them already in 2016. Only five in 2016 approved seminars will be held in 2018, often due to an explicit request of the seminars organizers.

2.4 Seminar-Programm 2016

In 42 von 49 Wochen, in denen das Tagungszentrum 2016 geöffnet war, fand mindestens ein Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop statt. In 30 Wochen waren es sogar zwei. An drei Wochen hatte das Zentrum ganz geschlossen während in den verbleibenden sieben

2.4 *The Seminar Program in 2016*

At least one Dagstuhl Seminar or Dagstuhl Perspectives Workshop was held in 42 of 49 week in 2016. For three weeks the center were completely closed while in seven weeks there were exclusively other events scheduled. In 30 weeks there were actually two seminars in parallel.

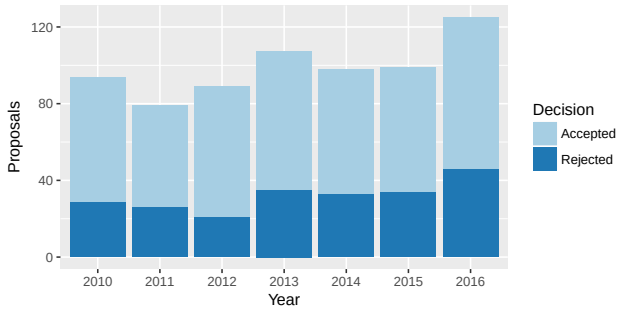

Fig. 2.2
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2010–2016.

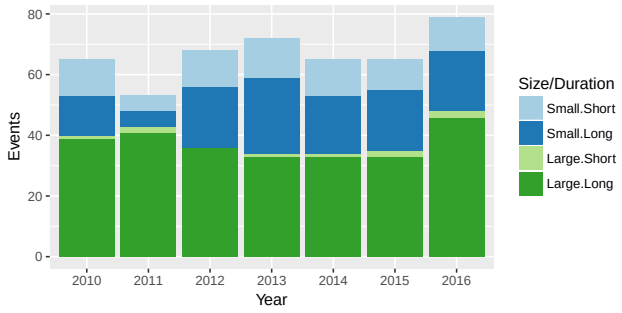

Fig. 2.3
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2010–2016. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

Wochen das Zentrum nur durch andere Veranstaltungen belegt war.

Seit 2012 ist es aufgrund des damals fertiggestellten Gästehauses möglich, zwei Seminare parallel in einer Woche zu veranstalten. Dadurch ist, verglichen mit den Jahren zuvor, seit 2012 die Gesamtanzahl an Seminaren pro Jahr gestiegen. In 2016 fanden mit 72 Seminaren genau so viele statt wie 2015. In Fig. 2.4 ist die Entwicklung der vergangenen Jahre dargestellt.

Since the guest house opened in 2012, it has been possible for the center to schedule two parallel seminars in any given week. Due to this, there were an increase of seminars held since 2012 compared with the years before. With 72 seminars there were in 2016 exactly the same number of seminars as 2015. 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 2016 knapp die Hälfte (49 %, 1 115 von 2 267) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 16 % der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten. Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass etwa 46 % der Antwortenden 2016 das erste Mal und weitere 15 % zum zweiten 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

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 year, also in 2016 a bit less than the half (1,115 of 2,267, or 49 %) researchers were first-time visitors to Dagstuhl. About 16 % additional researchers had already attended one previous seminar in the years before. Slightly different numbers are obtained from our guest survey: About 46 % of the reponders were first-time visitors and additional 15 % states their second visit (see Figure 2.5a).

A healthy number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 34 % of 2016 seminar and workshop survey respon-

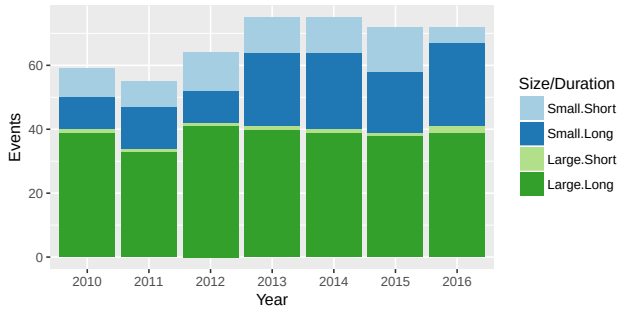

Fig. 2.4
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2010–2016. Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

ist für den weiteren Verlauf ihres Lebenswegs. Etwa 34 % der Gäste der Seminare und Workshops in 2016, die an unserer Umfrage 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 76 % war der Anteil von Gästen aus dem Ausland 2016 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2016 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details können Kapitel 13 entnommen werden.

In 2016 waren etwa 69 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 18 % wieder erfreulich hoch (siehe Fig. 2.6b).

dents self-classified as junior (see Fig. 2.5b). This proportion of junior to senior researchers has remained relatively constant over the years, reflecting the center's determined effort to maintain the "Dagstuhl connection" between brilliant junior scientists and their senior colleagues.

At over 76 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2016. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2016. For a detailed breakdown please refer to Chapter 13.

In 2016 were 69 % 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 also high both in total and relative terms, at 18 % (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 2016.

In den Seminaren, die sich Themen aus der theoretischen Informatik gewidmet haben, wurden sowohl klas-

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 seminars from 2016. 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 2016. The seminar summaries in Chapter 6 provide a full overview of the 2016 scientific seminar program.

Among the seminars which addressed topics from theoretical computer science, there were classical topics like *Scheduling* (16081) and *Structure and Hardness in P* (16451) but also topics, which were initially started in the mathematics and were applied in the computer science like *Algorithms and Effectivity in Tropical Mathematics and*

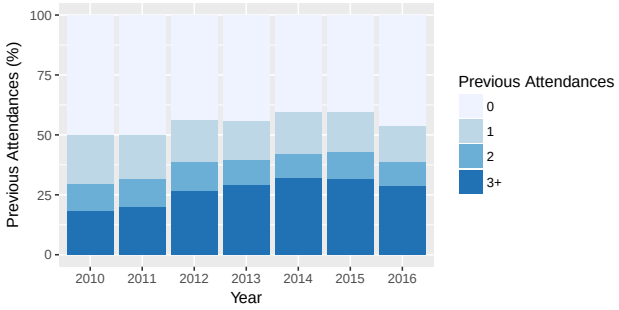

(a) Distribution of the number of previous attendances of participants, according to survey data.

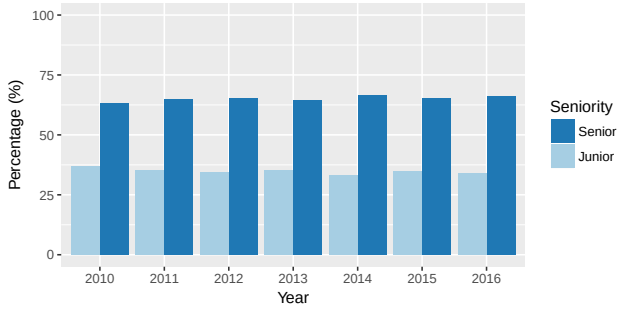

(b) Percentage of junior researchers, according to survey data.

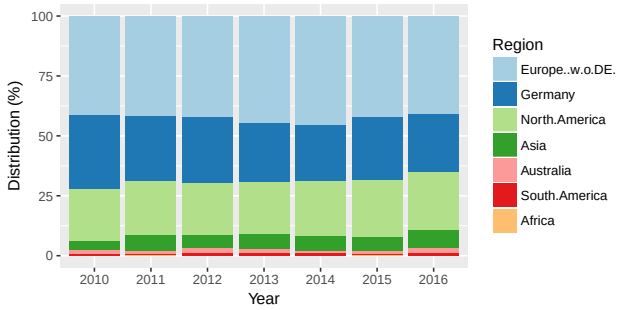

(c) Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2010–2016.

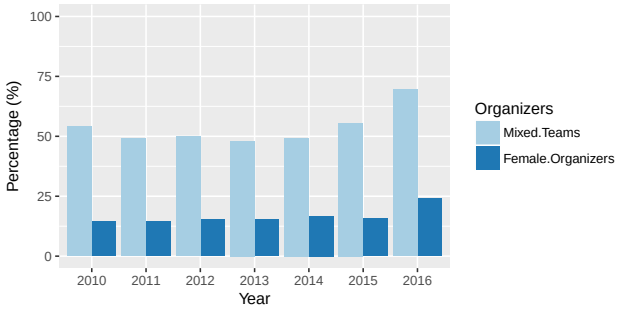

(a) Female organizers and mixed-gender organizer teams.

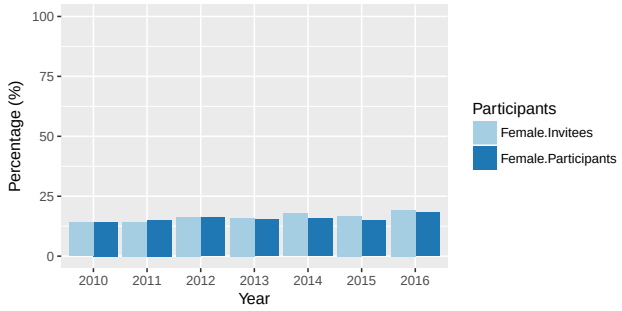

(b) Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2010–2016.

sische Themen wie *Scheduling* (16081) und *Structure and Hardness in P* (16451) als auch Themen, die aus der Mathematik kommend, erste Anwendungen in der Informatik haben wie z. B. *Algorithms and Effectivity in Tropical Mathematics and Beyond* (16482) diskutiert. Ergänzt wurde dies durch Themen, die Anwendungen in den aktuellen Gebieten der Informatik erörterten – wie *Coding Theory in the Time of Big Data* (16321) oder *Data Structures and Advanced Models of Computation on Big Data* (16101).

Die fortschreitende digitale Vernetzung der Welt gab Anlass für Seminare im Bereich *Netzwerke*. Beispiele sind *Network Latency Control in Data Centres* (16281) oder *Global Measurements: Practice and Experience* (16012). Mit Seminaren wie *Information-centric Networking and Security* (16251) wurde eine Brücke zu den Themen Security, Cryptography und Privacy geschlagen, die bei *Beyond* (16482). This has been supplemented by topics, that discuss current applications in the area of computer science like *Coding Theory in the Time of Big Data* (16321) or *Data Structures and Advanced Models of Computation on Big Data* (16101).

The ever-increasing digital interconnectedness gave motivation for several seminars about *Networking* topics, e.g. *Network Latency Control in Data Centres* (16281) or *Global Measurements: Practice and Experience* (16012). Seminars like *Information-centric Networking and Security* (16251) bridged the gap between *Networking* and topics like *Security*, *Cryptography* and *Privacy*. These topics were well represented in 2016 seminars and were discussed in a wide range: Starting from the basics like *Symmetric Cryptography* (16021) and *Public-Key Cryptography* (16371), over the interface of Cryptography und Security (*Modern Cryptography and Security: An Inter-Community*

den Seminaren in 2016 weiterhin gut vertreten waren. Diese Themen wurden von den Grundlagen wie *Symmetric Cryptography* (16021) und *Public-Key Cryptography* (16371) über die Schnittstelle von Cryptography und Security (*Modern Cryptography and Security: An Inter-Community Dialogue* (16051)) bis hin zu der ganzen Breite der Anwendungen in Seminaren wie *Hardware Security* (16202) oder *Privacy and Security in Smart Energy Grids* (16032) diskutiert.

Auch der Bereich Machine Learning, der seit einiger Zeit in der weltweiten Forschung einen Aufschwung erlebt, war in Dagstuhl gut vertreten und wurde in Seminaren wie *New Directions for Learning with Kernels and Gaussian Processes* (16481) oder *Foundations of Unsupervised Learning* (16382) diskutiert.

Erfreulicherweise waren auch Seminare über die verschiedenen Anwendungen der Informatik im täglichen Leben vertreten, wie z. B. *Automotive User Interfaces in the Age of Automation* (16262) oder *Computational Challenges in Cooperative Intelligent Urban Transport* (16091).

Herausforderungen, die durch den Einsatz von autonomen Systemen entstehen, wurden in Seminaren wie *Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR)* (16442) oder *Engineering Moral Agents – from Human Morality to Artificial Morality* (16222) diskutiert.

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2016 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.

Dialogue (16051)), up to the full width of application in seminars like *Hardware Security* (16202) or *Privacy and Security in Smart Energy Grids* (16032).

Also the area *Machine Learning*, which has attracted a lot of attention in the recent past in research worldwide, was well represented in Dagstuhl and has been discussed in seminars like *New Directions for Learning with Kernels and Gaussian Processes* (16481) or *Foundations of Unsupervised Learning* (16382).

Fortunately, also several seminars on applications of computer-science in daily life took place in Dagstuhl in 2016, e.g. *Automotive User Interfaces in the Age of Automation* (16262) oder *Computational Challenges in Cooperative Intelligent Urban Transport* (16091).

This brief selection of seminars should not hide the fact that each of the 2016 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, Lehrerfortbildungen und Fortbildung von jungen Journalisten und Volontären
- 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.

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, vocational training for teachers and instructors, and educational and training workshops for young journalists and trainees
- research group meeting 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.

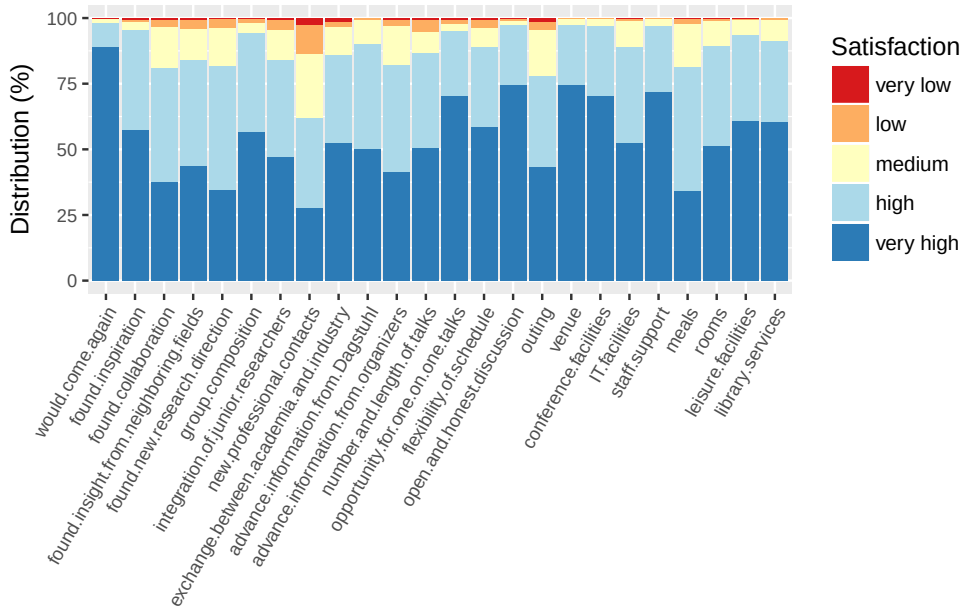

Fig. 2.7
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2016. According to survey results.

Qualitätssicherung 2.8 *Quality Assurance*

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

Fig. 2.7 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2016 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 471 Fragebögen, welche die Meinung von etwa 61 % der 2 393 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.

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

Fig. 2.7 shows the satisfaction of responding participants in 2016 with regard to selected aspects of their stay. The results were compiled from 1,471 questionnaires, representing the responses of about 61 % of all 2,393 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 Dagstuhl Seminar organizers a more transparent invitation process by giving them direct access to the status of invitee replies via a dedicated webpage. The page is available 24/7 and has met with very positive feedback from organizers.

Auslastung des Zentrums 2.9 *Utilization of the Center*

Auch 2016 konnte Schloss Dagstuhl die durch das neue Gästehauses ermöglichte hohe Auslastung weitgehend halten. Es gab 2016 insgesamt 13 337 Gasttage, wobei 10 857 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Bezogen auf die Seminar- und Workshopgäste bedeutet dies einen minimalen Rückgang verglichen mit 2015. Allerdings gab es insgesamt mehr Gasttage als in 2015. Es fanden im Berichtsjahr 116 Veranstaltungen mit insgesamt 3 353 Gästen statt. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2016 ebenso unbelegt wie jeweils zwei Wochen im Juli/August und am Jahresende. Diese wurden zu Instandhaltungs- und Verwaltungsarbeiten benötigt.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2016 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite ist ein Kalender[8] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

Thanks to the new guest house, Schloss Dagstuhl was able to uphold the high capacity utilization again in 2016. There were 13,337 overnight stays in total, with 10,857 overnight stays in seminars and perspective workshops. The latter was slightly less than in 2015. However, there were more overnight stays in total in 2016 compared to stays in 2015. The center hosted a total of 116 events with 3,353 guests in 2016. See Chapter 13 for further details.

Weekends were kept free in 2016, as well as two weeks in July/August and 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 2016, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar[8] of upcoming events and further information and material on all events past, present and future, e.g. aims and scope, participant list, and concluding report.

[8] http://www.dagstuhl.de/no_cache/programm/kalender/

Fig. 2.8
„Schloss Dagstuhl." Twitter post by 16241 Dagstuhl Seminar participant Krystal Guo. https://twitter.com/guo_krystal/status/741917085899575296. Photo courtesy of Krystal Guo.

3 Bibliographiedatenbank dblp
dblp computer science bibliography

3.1 Offene Bibliographiedaten für die Informatik

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 immensem 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 20 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.

3.1 *Open Bibliographic Data in Computer Science*

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 20 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 24 200 Treffer; im Einzelnen weisen SpringerLink ca. 2 400 Artikel, Elsevier ScienceDirect über 550 Artikel, die ACM Digital Library ca. 475 Artikel und IEEE Xplore über 200 Artikel nach.

[10] The search term “dblp” results in 24,200 hits at Google Scholar; in particular, SpringerLink lists about 2,400 articles, Elsevier ScienceDirect lists more than 550 articles, the ACM Digital Library lists 475 articles, and IEEE Xplore lists more than 200 articles.

Schloss Dagstuhl und dblp 3.2

Die Zusammenarbeit zwischen Schloss Dagstuhl und der ursprünglich an der Universität Trier entwickelten Bibliographiedatenbank dblp besteht bereits seit Ende 2010. Zunächst durch ein Projekt im Leibniz-Wettbewerb gefördert, wird das Engagement seit Juni 2013 von Schloss Dagstuhl direkt mitfinanziert. Die Finanzierung wird zudem seit November 2010 durch eine großzügige Spende der Klaus-Tschira-Stiftung unterstützt. Bereits seit 2012 steht nun auch unter dblp.dagstuhl.de ein eigener dblp-Webservice unter der Domain von Schloss Dagstuhl bereit und ergänzt damit das dblp-Angebot der Universität Trier unter dblp.uni-trier.de. Das Kooperationsabkommen zwischen Schloss Dagstuhl und der Universität Trier wurde Ende 2016 um zunächst weitere zwei Jahre verlängert.

Im Zuge der Konsolidierung der Zusammenarbeit wurden unter dem Dach von Schloss Dagstuhl zweieinhalb Mitarbeiterstellen im wissenschaftlichen Stab geschaffen, die hauptamtlich für die Betreuung und Weiterentwicklung von dblp abgestellt sind. Der dblp-Beirat (siehe Fig. 3.1) leistet seit November 2011 unter dem Dach von Schloss Dagstuhl die wissenschaftliche Aufsicht und unterstützt das dblp-Team mit seiner Expertise.

Schloss Dagstuhl and dblp

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. Since November 2010, Schloss Dagstuhl's dblp team has also been supported by a generous donation from the Klaus Tschira Foundation. Schloss Dagstuhl's own dblp web service at dblp.dagstuhl.de was established in 2012 and complements the dblp service available at the University of Trier at dblp.uni-trier.de. In late 2016, the cooperation agreement between Schloss Dagstuhl and the University of Trier was renewed for another two years.

As part of the consolidation of this cooperation, two and a half Schloss Dagstuhl 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.

dblp-Beirat | dblp Advisory Board
Prof. Dr. Hannah Bast University of Freiburg, Germany | *Chair*
Prof. Dr. Andreas Butz Ludwig Maximilians University Munich, Germany
Prof. Dr.-Ing. Rüdiger Dillmann Karlsruhe Institute of Technology, Germany
Prof. 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.

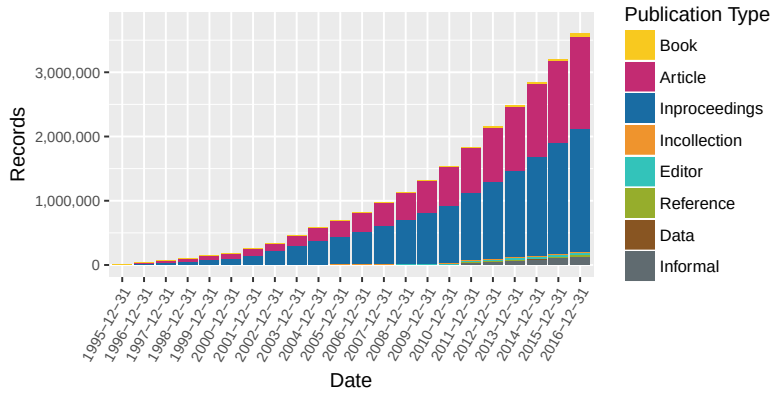

(a) Total number of records by year and type

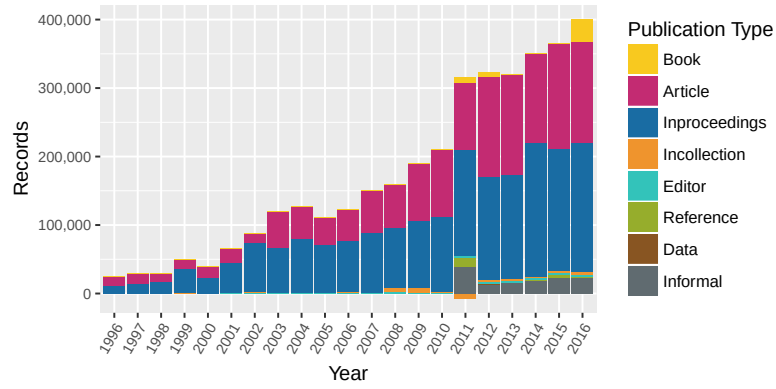

(b) New records by year and type

Fig. 3.2
Development of the dblp data stock.

Statistiken der Datenakquise 3.3 *Data Acquisition Statistics*

Die Bibliographiedatenbank 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.

Zum 31. Dezember 2016 indexierte dblp mehr als 3,6 Millionen Publikationen. In dem Zeitraum von Anfang Januar 2016 bis Ende Dezember 2016 wurden dabei mehr als 400 000 neue Publikationseinträge in dblp aufgenommen. Dies entspricht mehr als 1 500 neuen Publikationen pro Arbeitstag. Somit konnte nun bereits zum dritten mal in Folge die Rekordaufnahmequote des Vorjahres übertroffen werden. Die neu aufgenommenen Einträge verteilen sich zu 46,9% auf Konferenzbeiträge, zu 36,8% auf Journalartikel, zu 8,1% auf Monographien und Dissertationen, sowie zu 8,4% 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.

Between January 1, 2016, and December 31, 2016, the dblp database grew by more than 400,000 publication records to reach a total of more than 3.6 million records. This corresponds to more than 1,500 new records for each working day of the year. Hence, for three successive years, dblp has been able to further increse the record number of new records included in the previous year. Of these new records, 46.9% have been conference papers, 36.8% have been journal articles, 8.1% have been monographs and PhD theses, and 8.4% have been other publications.

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

	Trier 1	Trier 2	Dagstuhl	Total
user sessions (visits) per day	26 911	1 428	1 254	29 593
page views per day	501 208	26 355	35 406	562 969
page views per user session	18,6	18,4	28,2	
distinct users (IPs) per month	393 273	25 249	20 416	
data served per month	1 187,6 GB	72,7 GB	120,7 GB	1 381 GB

Fig. 3.3
Average usage of the three dblp servers in 2016. Trier 1 = http://dblp.uni-trier.de, Trier 2 = http://dblp2.uni-trier.de, Dagstuhl = http://dblp.dagstuhl.de

Nutzungsstatistiken 3.4 *Usage Statistics*

2016 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

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

Die 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 ist zu beachten, dass Server Trier 1 aufgrund seiner prominenten Sichtbarkeit in den Google-Suchergebnissen die mit Abstand bekannteste Adresse besitzt. Der Server in Dagstuhl hat dabei jedoch zunehmend an Sichtbarkeit gewonnen und konnte im Laufe des Jahres die Zahl seiner Nutzer vervierfachen.

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 2016, three official dblp web servers were updated and synchronized on a daily basis:

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

The 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. The three servers do show a very different rate of usage, with Trier 1 being the by far most widely known server. This is of course due to the fact that server Trier 1 is ranked so highly by the Google search engine. However, server Dagstuhl has become increasingly more visibile, and the number of its users quadrupled during the course of 2016.

Figure 3.3 shows the average usage of all three servers in 2016. These figures ignore the traffic caused by known bots and crawlers.

Dissertationen in dblp 3.5 *PhD theses added to dblp*

2016 konnten Metadaten von mehr als 9 000 Informatik-Dissertationen der französischen Datenbank Hyper Articles en Ligne (HAL)[11] sowie 10 000 Dissertationen des EThOS[12]-Dienstes der britischen Nationalbibliothek in dblp integriert werden. Diese Datensätze umfassen einen Großteil der französischen und britischen Dissertationen seit den 1950er Jahren. Mit den beiden Diensten wurden somit neben den über 20 000 Datensätzen der *Deutschen Nationalbibliothek (DNB)* zwei weitere, kontinuierliche Metadatenquellen für Dissertation in dblp integriert.

Da dblp Dissertationen als eine zentrale wissenschaftliche Veröffentlichung in der Laufbahn eines jeden Forschers versteht, ist es ein erklärtes Ziel von dblp, die allgemeine Abdeckung der Dissertationen zu verbessern und auf weitere, internationale Datenquellen auszuweiten.

In 2016, metadata of more than 9,000 computer science PhD theses from the French archive Hyper Articles en Ligne (HAL)[11] and more than 10,000 PhD theses from the EThOS[12] repository of the British Library have been imported into dblp. These theses cover a significant part of the French and British computer science community, with some theses reaching back as far as the 1950s. The HAL and EThOS archives are the second and third source of PhD metadata to be continuously imported into dblp in addition to the more than 20,000 data records already added from the *Deutsche Nationalbibliothek (DNB)*, the German National Library.

Since dblp understands the PhD thesis as a very important and central publication in the professional life of a computer scientist, we aim to further increase the coverage of PhD theses in dblp.

[11] https://hal.archives-ouvertes.fr/
[12] http://ethos.bl.uk

Gemeinsames Projekt von dblp, Zentralblatt MATH und HITS

3.6

Joint Project of dblp, Zentralblatt MATH, and HITS

Die Urheberschaft wissenschaftlicher Publikationen eindeutig zu erkennen und zuzuordnen ist eine der großen Herausforderungen bibliographischer Datendienste. Die Forschung kennt dieses Problem in seiner allgemeinen Form als das Problem der „Entity-Resolution“ oder der „Autorennamen-Disambiguierung“, welches ein wichtiges Forschungsthema im Bereich der linguistischen Datenverarbeitung darstellt. In einem gemeinsamen Projekt nehmen sich die Bibliographiedatenbank dblp, die Datenbank zbMATH des FIZ Karlsruhe und das Heidelberger Institut für Theoretische Studien (HITS) diesem Problem an und entwickeln mit Hilfe des aktuellen Forschungsstandes gemeinsame Lösungsstrategien. Die Datensätze von zbMATH und dblp teilen dabei die Probleme bei der Identifikation von Autorennamen. Die Kombination beider Datensätze, bestehend aus teils überlappenden und teils disjunkten Einträgen, stellt dabei eine einzigartige Möglichkeit dar, Fehler in den Datensätzen aufzudecken und voneinander zu lernen. Die Natural-Language-Processing (NLP) Forschungsgruppe des HITS um Prof. Michael Strube bringt dabei ihre Erfahrung mit graph- und netzwerkbasierten NLP-Methoden bei der Co-Referenz-Resolution und der Konzept- bzw. Entitäts-Disambiguierung ein.

Seit 2015 wird das Projekt im Leibniz Wettbewerb in der Förderline „Nationale und internationale Vernetzung“ gefördert, wodurch ein weiterer wissenschaftlicher Mitarbeiter zur Verstärkung des dblp-Teams gewonnen werden konnte. Seitdem wurde ein steter Datenaustausch zwischen den Projektpartnern initiiert. Zahlreiche in dblp und zbMATH gemeinsam vertretene Autoren konnten identifiziert und deren Autorenprofile in den beiden Datenbanken verlinkt werden. Zudem wurde eine neuer, hochqualitativer Test- und Trainingsdatensatz (sogenannte „Goldstandard-Daten“) für die Autorennamen-Disambiguierung erstellt und dessen Veröffentlichung im Fachjournal *Scientometrics* angenommen. Derzeit werden erste Entwürfe eines auf neuronalen Netzwerken basierenden Disambiguierungsverfahrens in den Evaluations- und Testumgebungen von dblp und zbMATH erprobt. Das Projekt läuft noch bis Juni 2018.

The correct attribution of scholarly material to their unambiguous authors ranks among the most critical challenges for digital libraries. More generally, the problem of determining which records in a database refer to the same entities is known as “entity resolution” or “author name disambiguation” and constitutes an important field of research within the discipline of natural language processing. In a joint project, the dblp computer science bibliography and the zbMATH database (operated by FIZ Karlsruhe) are partnering with the Heidelberg Institute for Theoretical Studies (HITS) to find and implement new and state-of-the-art strategies to overcome the challenges of author identification and disambiguation. zbMATH and dblp share the challenges associated with author name disambiguation. Their partially overlapping, but also partially disjointed data, allow for a joint effort to identify authors based on the combination of the two data sets. The Natural Language Processing (NLP) Group at the HITS, lead by Prof. Michael Strube, provides its extensive experience with graph-based and network methods for NLP tasks such as co-reference resolution, concept disambiguation, and entity disambiguation.

Since 2015, the project is funded by a grant from the “National and International Networking” funding line of the Leibniz Competition. Since then, one further scientific project staff member reinforced the dblp team. zbMATH and dblp established an ongoing exchange of metadata which allowed for numerous common author profiles to be identified in and linked between the dblp and zbMATH data sets. Additionally, a new, high quality training and test data set (“gold standard data”) for the author disambiguation task has been developed and accepted for publication in the journal *Scientometrics*. Currently, a first prototype of a new disambiguation approach based on neural networks is evaluated in the testing environments of dblp and zbMATH. The project will conclude in June 2018.

4 Dagstuhl Publishing
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.

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 artefacts since 2015.

Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*[13] dokumentiert, was eine Zitation der Seminare im wissenschaftlichen Kontext ermöglicht. Zudem erlaubt es auch denjenigen Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Seminaren und Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 11.4) agiert als Herausgebergremium 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 2016 wurde für 89 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 Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*[13] which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

The periodical started with the first seminars of January 2011 and publishes in monthly issues reports on seminars and workshops that took place on a given month. The content is not peer-reviewed. The Scientific Directorate (see Fig. 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).

89 reports of Dagstuhl Seminars and Dagstuhl Perspectives Workshops have been published in 2016. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*[14] die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 11.4) fungiert hier ebenfalls als Herausgebergremium. In 2016 wurde keine Ausgabe veröffentlicht. Allerdings wurden mehrere Manifestos zur Begutachtung eingereicht, die nach Annahme durch das Herausgebergremium vorrausichtlich 2017 veröffentlicht werden.

Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*[14] in Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. In 2016 no volume was published. However, several Dagstuhl Manifestos has been submitted for review and – if accepted by the editorial board – which are scheduled for publishing in 2017.

Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*[15] ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspekti-

Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*[15] book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Per-

[13] http://drops.dagstuhl.de/dagrep
[14] http://drops.dagstuhl.de/dagman

Dagstuhl Publishing

ven-Workshop basiert. Für jedes Buch ist ein Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausgebergremium verantwortlich ist) begutachtet und freigegeben werden muss. In 2016 wurde kein Buch in der Reihe veröffentlicht, jedoch gab es einen Antrag ausgehend vom Dagstuhl Seminar 15301 „The Constraint Satisfaction Problem: Complexity and Approximability“, welcher angenommen wurde. Dieser Band wird voraussichtlich im ersten Quartal 2017 veröffentlicht.

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[16] veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausgebergremium (Fig. 4.1), 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 2016 wurden 6 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.2.

spectives Workshop. Each book needs a proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). In 2016, no volume was published in the series. However, a proposal based on Dagstuhl Seminar 15301 “The Constraint Satisfaction Problem: Complexity and Approximability” was submitted and accepted. This volume will be published in the first quarter of 2017.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[16] aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 4.1, 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 2016, Dagstuhl published 6 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.2.

[15] http://drops.dagstuhl.de/dfu
[16] 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. 4.1
OASIcs Editorial Board.

Vol. 50 | 5th Student Conference on Operational Research (SCOR'16) http://www.dagstuhl.de/dagpub/978-3-95977-004-0
Vol. 51 | 5th Symposium on Languages, Applications and Technologies (SLATE'16) http://www.dagstuhl.de/dagpub/978-3-95977-006-4
Vol. 52 | Technical Communications of the 32nd International Conference on Logic Programming (ICLP 2016) http://www.dagstuhl.de/dagpub/978-3-95977-007-1
Vol. 53 | 7th Workshop on Computational Models of Narrative (CMN 2016) http://www.dagstuhl.de/dagpub/978-3-95977-020-0
Vol. 54 | 16th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2016) http://www.dagstuhl.de/dagpub/978-3-95977-021-7
Vol. 55 | 16th International Workshop on Worst-Case Execution Time Analysis (WCET 2016) http://www.dagstuhl.de/dagpub/978-3-95977-025-5

Fig. 4.2
OASIcs volumes published in 2016.

Dagstuhl Publishing

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[17] veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausgebergremium (siehe Fig. 4.3) besteht aus einschlägig bekannten Wissenschaftlern und wird seit Mai 2015 von Wolfgang Thomas als Hauptherausgeber geleitet.

In 2016 wurden Tagungsbände von 19 Konferenzen veröffentlicht, so viel wie noch nie zuvor; siehe Fig. 4.4.

Die Konferenz *Computer Science Logic (CSL)* wurde vom Herausgebergremium im Rahmen der bestehenden Kooperation erneut evaluiert und für weitere fünf Jahre (2016–2020) aufgenommen.

In 2016 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. 4.5 sind alle Konferenzen aufgelistet, deren Anträge bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischweise 5 Jahre) eingegangen wurde.

LIPIcs aber auch die Reihe OASIcs erheben eine Veröffentlichungsgebühr (article-processing charge, APC). Nach einer Prüfung der Kostendeckung in 2016, haben die Geschäftsführung und der Aufsichtsrat von Schloss Dagstuhl eine Erhöhung dieser Gebühr von 15 € auf 60 € beschlossen. Allerdings erfolgt dank großzügiger finanzieller Unterstützung durch das Heidelberger Institut für Theoretische Studien (HITS) die Erhöhung schrittweise von 15 € in 2016 bis auf 60 € in 2019.

LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series[17] publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers (see Fig. 4.3) supervises the conferences that are accepted for LIPIcs and is headed since May 2015 by Wolfgang Thomas.

The series published the proceedings of 19 major conferences in 2016, marking a record high since the series was started; see Fig. 4.4.

The conference *Computer Science Logic (CSL)* was re-evaluated by the LIPIcs editorial board and accepted for another five-year period (2016–2020).

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2016 has again seen several applications for LIPIcs, continuing the high interest from the previous years. Fig. 4.5 lists all conferences that have been accepted for a cooperation covering several years (typically 5 years).

The series LIPIcs as well as OASIcs are charging an article-processing charge (APC). After an evaluation of the cost recovery in 2016, the management and the supervisory board of Schloss Dagstuhl decided an increase of the APC from 15 € to 60 €. However and due to financial support by the Heidelberg Institute of Theoretical Studies (HITS), the APC will be increased step by step from 15 € in 2016 to 60 € in 2019.

[17] http://drops.dagstuhl.de/lipics

Prof. Dr. Susanne Albers Technical University Munich, Germany
Prof. Dr. Chris Hankin Imperial College London, United Kingdom
Prof. Deepak Kapur, Ph. D. University of New Mexico, US
Prof. Michael Mitzenmacher, Ph. D Harvard University, US
Prof. Madhavan Mukund, Ph. D. Chennai Mathematical Institute, India
Dr. Catuscia Palamidessi INRIA, France
Prof. Dr. Wolfgang Thomas RWTH Aachen, Germany | Chair
Pascal Weil, Ph. D CNRS, France and University Bordeaux, France
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Saarland University, Germany

Fig. 4.3
LIPIcs Editorial Board.

Dagstuhl Publishing

Vol. 46 | 19th International Conference on Principles of Distributed Systems (OPODIS 2015) http://www.dagstuhl.de/dagpub/978-3-939897-98-9
Vol. 47 | 33rd International Symposium on Theoretical Aspects of Computer Science (STACS 2016) http://www.dagstuhl.de/dagpub/978-3-95977-001-9
Vol. 48 | 19th International Conference on Database Theory (ICDT 2016) http://www.dagstuhl.de/dagpub/978-3-95977-002-6
Vol. 49 | 8th International Conference on Fun with Algorithms (FUN 2016) http://www.dagstuhl.de/dagpub/978-3-95977-005-7
Vol. 50 | 31st Conference on Computational Complexity (CCC 2016) http://www.dagstuhl.de/dagpub/978-3-95977-008-8
Vol. 51 | 32nd International Symposium on Computational Geometry (SoCG 2016) http://www.dagstuhl.de/dagpub/978-3-95977-009-5
Vol. 52 | 1st International Conference on Formal Structures for Computation and Deduction (FSCD 2016) http://www.dagstuhl.de/dagpub/978-3-95977-010-1
Vol. 53 | 15th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2016) http://www.dagstuhl.de/dagpub/978-3-95977-011-8
Vol. 54 | 27th Annual Symposium on Combinatorial Pattern Matching (CPM 2016) http://www.dagstuhl.de/dagpub/978-3-95977-012-5
Vol. 55 | 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016) http://www.dagstuhl.de/dagpub/978-3-95977-013-2
Vol. 56 | 30th European Conference on Object-Oriented Programming (ECOOP 2016) http://www.dagstuhl.de/dagpub/978-3-95977-014-9
Vol. 57 | 24th Annual European Symposium on Algorithms (ESA 2016) http://www.dagstuhl.de/dagpub/978-3-95977-015-6
Vol. 58 | 41st International Symposium on Mathematical Foundations of Computer Science (MFCS 2016) http://www.dagstuhl.de/dagpub/978-3-95977-016-3
Vol. 59 | 27th International Conference on Concurrency Theory (CONCUR 2016) http://www.dagstuhl.de/dagpub/978-3-95977-017-0
Vol. 60 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2016) http://www.dagstuhl.de/dagpub/978-3-95977-018-7
Vol. 61 | 11th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2016) http://www.dagstuhl.de/dagpub/978-3-95977-019-4
Vol. 62 | 25th EACSL Annual Conference on Computer Science Logic (CSL 2016) http://www.dagstuhl.de/dagpub/978-3-95977-022-4
Vol. 64 | 27th International Symposium on Algorithms and Computation (ISAAC 2016) http://www.dagstuhl.de/dagpub/978-3-95977-026-2
Vol. 65 | 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2016) http://www.dagstuhl.de/dagpub/978-3-95977-027-9

Fig. 4.4
LIPIcs volumes published in 2016.

COSIT	Conference on Spatial Information Theory accepted for 2017–2021
CPM	Annual Symposium on Combinatorial Pattern Matching accepted for 2016–2020
CSL	Computer Science Logic accepted for 2016–2020
ECRTS	Euromicro Conference on Real-Time Systems accepted for 2017–2021
ISAAC	International Symposium on on Algorithms and Computation accepted for 2016–2020
ITCS	Innovations in Theoretical Computer Science Conference accepted for 2017–2021
SEA	International Symposium on Experimental Algorithms accepted for 2017–2021

Fig. 4.5
Conferences accepted in 2016 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 Royal Institute of Technology Stockholm, Sweden | Subject area: Distributed embedded systems and networks
Prof. Bashir Al Hashimi University of Southampton, UK | Subject area: Energy-efficiency
Prof. Mateo Valero, Ph. D. Technical University of Catalonia | Subject area: High-performance embedded systems
Prof. Dr. Martin Fränzle Carl von Ossietzky University Oldenburg, Germany | Subject area: Hybrid systems
Prof. Dr. Samarjit Chakraborty Technical University Munich, Germany | Subject area: Multimedia applications
Prof. Dr. Gernot Heiser University of New South Wales, Australia | Subject area: Operating systems
Prof. Dr. Lothar Thiele ETH Zürich, Switzerland | Subject area: Performance and wireless sensor networks
Dr. Neil Audsley University of York, UK | Subject area: Real time
Prof. Sanjoy Baruah, 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. 4.6
LITES Editorial Board.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[18] veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die sich für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 4.6), begutachtet alle eingereichten Arbeiten. Die Zeitschrift wird gemeinsam mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*[19] der Fachgesellschaft *European Design and Automation Association (EDAA)*[20] 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 2016 wurde eine Ausgaben von *LITES* mit insgesamt 5 Artikeln veröffentlicht.

DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*[21] werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2016 wurde die zweite Ausgabe mit 14 Artefakten veröffentlicht. Diese Artefakte wurden im Rahmen der 30. European Conference on Object-Oriented Programming (ECOOP'16) – deren Konferenzband als Volume 56 in LIPIcs veröffentlicht wurde – evaluiert, wobei das Regelwerk „Artifact Evaluation for Software Conferences"[22] angewendet wurde. Jedes Artefakt wird mit einer separaten Beschreibung veröffentlicht.

Die Veröffentlichung und Bereitstellung von Forschungdaten 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"[23], der in 2016 durch zwei weitere Seminare ergänzt wurde: „Reproducibility of Data-Oriented Experiments in e-Science"[24] und „Rethinking Experimental Methods in Computing"[25].

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.

LITES: Leibniz Transactions on Embedded Systems

The *LITES*[18] 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. 4.6), reviews all submitted contributions. The journal is jointly published with the *EMbedded Systems Special Interest Group (EMSIG)*[19] of the *European Design and Automation Association (EDAA)*[20]. 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 2016, one issue of *LITES* containing 5 articles in total was published.

DARTS: Dagstuhl Artifacts Series

The *DARTS* series[21] publishes evaluated research data and artifacts. It is organized as a periodical. In 2016, one volume containing one issue with 14 artifacts was published. This second issue of the series contains the research artifacts of the 30th European Conference on Object-Oriented Programming (ECOOP'16), whose conference proceedings were published as volume 56 in the LIPIcs series. Each artifact is published with a separate description and was evaluated according to the guidelines for "Artifact Evaluation for Software Conferences"[22].

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"[23] took place which was complemented with two seminars in 2016: "Reproducibility of Data-Oriented Experiments in e-Science"[24] and "Rethinking Experimental Methods in Computing"[25].

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 focusses on conference proceedings publications.

[18] http://drops.dagstuhl.de/lites
[19] http://www.emsig.net/
[20] http://www.edaa.com/
[21] http://www.dagstuhl.de/darts

Infrastruktur 4.2 *Infrastructure*

Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 4.7. Die Bände aus den Reihen *LIPIcs* und *OASIcs* werden zudem bei Scopus[26] 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. 4.7.

Zudem unterstützen die technischen Schnittstellen die Datenakquisition durch GoogleScholar, so dass die Publikationen sichtbarer und besser recherchierbar sind.

Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 4.7. The LIPIcs and OASIcs volumes are submitted to Scopus[26] 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. 4.7.

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

[22] http://www.artifact-eval.org/
[23] http://www.dagstuhl.de/15452
[24] http://www.dagstuhl.de/16041
[25] http://www.dagstuhl.de/16111
[26] http://www.scopus.com

dblp	
Dagstuhl Reports	http://dblp.org/db/journals/dagstuhl-reports/
Dagstuhl Manifestos	http://dblp.org/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups	http://dblp.org/db/series/dfu/
OASIcs	http://dblp.org/db/series/oasics/
LIPIcs	http://dblp.org/db/series/lipics/
LITES	http://dblp.org/db/journals/lites/
DARTS	http://dblp.org/db/journals/darts/

DOAJ	
OASIcs	https://doaj.org/toc/2190-6807
LIPIcs	https://doaj.org/toc/1868-8969
LITES	https://doaj.org/toc/2199-2002

Fig. 4.7
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

Dagstuhl Publishing

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[27] ein Online-Repositorium 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 Repositorium und stärkt dadurch Forschungsergebnisse aus der Informatik innerhalb dieses multidisziplininären Repositoriums.

AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Transformation und Open Science“[28] mit organisiert, welcher bereits der vierte Workshop in Folge seit 2013 ist. Der Workshop findet am 19. und 20. Januar 2017 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt.

AG Open Access der Schwerpunkt-initiative „Digitale Information“

Die Allianz der deutschen Wissenschaftsorganisationen, zu der neben der Max-Planck-Gesellschaft, der Helmholtz-Gemeinschaft, sowie weiteren Organisationen auch die Leibniz-Gemeinschaft gehört, hat eine Schwerpunktinitiative „Digitale Information“ ins Leben gerufen, bei der auch das Thema *Open Access* als Handlungsfeld vertreten ist. Mit Dr. Marc Herbstritt wurde seitens der Leibniz-Gemeinschaft ab Juli 2013 ein Mitglied des wissenschaftlichen Stabs von Schloss Dagstuhl in die Arbeitsgruppe „Open Access“[29] berufen.

Die Mitarbeit in dieser Arbeitsgruppe erlaubt, Anforderungen aus dem Wissenschaftsumfeld der Informatik auf politischer Ebene einzubringen. Zudem erleichtert es den Austausch und die Abstimmung fortlaufender Prozesse vor dem Hintergrund der weiterhin dynamischen Umgestaltung der Publikationslandschaft hin zu Open Access.

Technisches Back-end: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)[30] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[31] 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.[32]

LeibnizOpen

The Leibniz Association has established the *LeibnizOpen*[27] 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.

Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Transformation and Open Science”[28] was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop takes place at the Leibniz Association headquarters in Berlin on January 19 and 20, 2017.

Open Access Working Group of the Priority Initiative “Digital Information”

The Alliance of German Science Organizations, to which – among others – the Max Planck Society, the Helmholtz Association and also the Leibniz Association belong, has established a priority initiative “Digital Information” where *Open Access* is handled as a core activity. Since July 2013, Dagstuhl scientific staff member Dr. Marc Herbstritt has collaborated with this working group as the delegated representative of the Leibniz Association.[29]

Such collaboration offers an opportunity to highlight the scientific requirements of the computer science discipline on a political level. Additionally, it enables and simplifies the exchange and calibration of ongoing changes in the publishing landscape towards Open Access.

Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)[30]. The general guidelines of the Dublin Core initiative[31] 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.[32]

[27] http://www.leibnizopen.de/
[28] https://www.dagstuhl.de/dagpub/journalmanagement-leibniz/2017-01-19-workshop/
[29] http://www.allianzinitiative.de/de/handlungsfelder/open_access/
[30] http://www.dagstuhl.de/drops

Dagstuhl Publishing

Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)[33] zur (digitalen) Langzeitarchivierung eingereicht.

Mirroring

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

- io-port.net: Das unter Leitung des FIZ Karlsruhe, Leibniz-Institut für Informationsinfrastruktur, organisierte Informatik-Publikations-Portal io-port.net spiegelt alle Bände der LIPIcs-Reihe.[34]
- 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.[35]

Long-term Archiving

All publications are submitted to the German National Library (D-NB)[33] for (digital) long-term archiving.

Mirroring

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

- io-port.net: The informatics publication portal organized under the auspices of io-port.net, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, mirrors all volumes of the LIPIcs series[34].
- 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.[35]

[31] http://dublincore.org/
[32] http://elib.uni-stuttgart.de/opus/doku/about.php
[33] http://www.dnb.de/DE/Netzpublikationen/Langzeitarchivierung/langzeitarchivierung_node.html
[34] http://www.io-port.net (→ Digital Library → LIPIcs)
[35] http://vesta.informatik.rwth-aachen.de/Dagstuhl/

5 Resonanz
Feedback

Resonanz unserer Seminarorganisatoren

5.1

Feedback from Seminar Organizers

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

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

Organizers of Dagstuhl Seminar 16251

16251 – Information-centric Networking and Security | Dagstuhl Seminar | http://www.dagstuhl.de/16251

We thank Schloss Dagstuhl for providing a stimulating setting for this seminar. Much progress was made over the course of the seminar and since its completion. This is mainly because of the ease of face-to-face collaboration and interaction at Dagstuhl.

Organizers of Dagstuhl Seminar 16112

16112 – From Theory to Practice of Algebraic Effects and Handlers | Dagstuhl Seminar | http://www.dagstuhl.de/16112

We are extremely happy with the outcome of the seminar and the way we organized it. An open format that gives everyone ample time outside the seminar room was significantly boosted by the unique Dagstuhl environment free of worldly distractions. We encourage future organizers to boldly try new ways of organizing meetings. There will be confusion at first, but as long as the participants are encouraged and allowed to group themselves, they will do so. If a lesson is to be taken from our seminar, it is perhaps this: let people do what they want, but also make sure they report frequently on what they are doing, preferably when they are a bit hungry.

Organizers of Dagstuhl Seminar 16372

16372 – Uncertainty Quantification and High Performance Computing | Dagstuhl Seminar | http://www.dagstuhl.de/16372

The organizers would like to express their gratitude to all participants of the Seminar. Special thanks go to the Schloss Dagstuhl team for its extremely friendly support during the preparation phase and for the warm welcome at Schloss Dagstuhl.

Organizers of Dagstuhl Seminar 16431

16431 – Computation over Compressed Structured Data | Dagstuhl Seminar | http://www.dagstuhl.de/16431

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

Resonanz in Sozialen Netzwerken

5.2 *Feedback in Social Media*

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

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

Moritz Stefaner (Truth & Beauty, Lilienthal, Germany)

16061 – Data-Driven Storytelling | Dagstuhl Seminar | https://medium.com/data-driven-storytelling/some-things-i-learned-about-data-driven-story-telling-in-schlo\T1\ss-dagstuhl-b5ecfaef0910

Thanks to Schloss Dagstuhl for being what it is, and to the workshop organizers…

Dominik Engel (FH Salzburg, Austria)

16032 – Privacy and Security in Smart Energy Grids | Dagstuhl Seminar | https://www.en-trust.at/blog/2016/01/dagstuhl-seminar-privacy-security-smart-energy-grids/

In the German speaking computer science world, Dagstuhl is a magic word – the seminars there are as renowned as the library and the obligatory picture on the stairs of Schloss Dagstuhl.

Jan Erik Moström (University of Umeå, Sweden)

16072 – Assessing Learning In Introductory Computer Science | Dagstuhl Seminar | http://mostrom.eu/2016/03/19/dagstuhl-seminar-what-a-great-idea/

To summarize: if you ever get an invitation from Dagstuhl, don't put it in your spam folder just answer "Yes" and go there.

Petra Isenberg (INRIA Saclay - Orsay, France)

16231 – Immersive Analytics | Dagstuhl Seminar | https://twitter.com/dr_pi/status/741359198731145216

yes, the seminar was great and we are very grateful for how well @dagstuhl helped to take care of our kids

Mike Croucher (Software Sustainability Institute, Edinburgh, United Kingdom)

16252 – Engineering Academic Software | Dagstuhl Seminar | http://www.walkingrandomly.com/?p=6147

I love attending seminars like this because I get to learn about all of the wonderful things that the community is up to.

Andrew Winslow (Free University of Brussels, Belgium)

16271 – Algorithmic Foundations of Programmable Matter | Dagstuhl Seminar | https://twitter.com/awinslow_cs/status/751768098676572160

Just returning from a week @dagstuhl. Easily the best place to host a research workshop I've ever been to.

John McCormack (Monash University, Caulfield, Australia)

16239 – Research Stay | Research Stay | https://twitter.com/jonmcc/status/741996251592413184

The @dagstuhl library has some real gems, including Georg Nees' formel • forbe • form

Resonanz im Fragebogen 5.3 *Seminar Survey Feedback*

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

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

16012 – Global Measurements: Practice and Experience | Dagstuhl Seminar | http://www.dagstuhl.de/16012

Very great location and place to be – delighted on staff and their service and patience.

16012 – Global Measurements: Practice and Experience | Dagstuhl Seminar | http://www.dagstuhl.de/16012

Dagstuhl is a nice cozy venue, no fuss, just focus on the interaction and discussion with people. The facilities are decent and fit the job, including the room which is very basic, but enough for the stay. The honor concept of buying drinks and snacks is superb.

16011 – Evolution and Computing | Dagstuhl Seminar | http://www.dagstuhl.de/16011

Thanks a lot to the organizers for the organization of this truly inspiring seminar. And thanks to Schloss Dagstuhl staff for the great support, the nice food, the cleanliness of all rooms et cetera. Dagstuhl is a great place to do research. I have benefited a lot from this seminar.

16012 – Global Measurements: Practice and Experience | Dagstuhl Seminar | http://www.dagstuhl.de/16012

This is one of the very few venues that gets it right! Thank you! Keep it up!

16062 – Modeling and Analysis of Semiconductor Supply Chains | Dagstuhl Seminar | http://www.dagstuhl.de/16062

A wonderful experience, congratulations to the center organization.

16072 – Assessing Learning In Introductory Computer Science | Dagstuhl Seminar | http://www.dagstuhl.de/16072

Thanks to the institutional sponsors for their support.

16072 – Assessing Learning In Introductory Computer Science | Dagstuhl Seminar | http://www.dagstuhl.de/16072

Great combination of formal and informal discussions – please have more in this area!

16072 – Assessing Learning In Introductory Computer Science | Dagstuhl Seminar | http://www.dagstuhl.de/16072

I have never been exposed to an environment like Dagstuhl before, I find it to be a wonderful opportunity for research and discussions between colleges. And of course to meet new people within the field.

16081 – Scheduling | Dagstuhl Seminar | http://www.dagstuhl.de/16081

I rarely attend regular conferences, which now became “journals that meet in a hotel”, as some people describe them. But I would never miss a workshop in Dagstuhl. I’ve initiated multiple new collaborations and quite a few of my research results originated from discussions and research conducted while in Dagstuhl.

16111 – Rethinking Experimental Methods in Computing | Dagstuhl Seminar | http://www.dagstuhl.de/16111

Dagstuhl has a winning formula. It does not need to be changed much.

16112 – From Theory to Practice of Algebraic Effects and Handlers | Dagstuhl Seminar | http://www.dagstuhl.de/16112

I got impressed with the scientific quality and seminar atmosphere. Many thanks to organizers and staff of Dagstuhl Schloss.

16131 – Language Based Verification Tools for Functional Programs | Dagstuhl Seminar | http://www.dagstuhl.de/16131

I have had a terrific time here (as always!) First time with my 2yo daughter and the support has been ABSOLUTELY OUTSTANDING! My only suggestion is to tell organizers to tell their participants that this facility exists at Dagstuhl (as more may bring their kids then :)).

16172 – Machine Learning for Dynamic Software Analysis: Potentials and Limits | Dagstuhl Seminar | http://www.dagstuhl.de/16172

Dagstuhl is wonderful! How about stopping with the tradition of handwritten abstracts.

16191 – Fresh Approaches to Business Process Modeling | Dagstuhl Seminar | http://www.dagstuhl.de/16191

The isolation encourages interaction. The good and affordable wines and beers help socialise. I like the honesty system – good way to keep costs down and engender an environment of trust.

16251 – Information-centric Networking and Security | Dagstuhl Seminar | http://www.dagstuhl.de/16251

dagstuhl is one of the very few premier seminar retreat venues. i continue to be deeply impressed by the staff, the venue, and the atmosphere. it is a privilege to meet here.

16351 – Next Generation Sequencing – Algorithms, and Software For Biomedical Applications | Dagstuhl Seminar | http://www.dagstuhl.de/16351

I was really impressed at the attention to detail in creating an atmosphere that is congenial for discussion and interaction. I think one of the best ideas was one of the simplest – shuffling the seating plan at mealtimes ensured I had a conversation with almost all attendees by the end of the meeting.

16151 – Foundations of Data Management | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16151

Compared to previous Dagstuhl seminars that I attended, this one had a much more pre-defined schedule. But it worked out very well; there was still sufficient room for discussions, etc. Very nice. I was very positively surprised how well the “group sessions” worked out. The concept was that people could move freely between the different groups. Still, I found that there were very fruitful and constructive discussions in all the groups (that I attended).

16162 – Managing Technical Debt in Software Engineering | Dagstuhl Seminar | http://www.dagstuhl.de/16162

Enough time to talk and change ideas. I can’t identify anything that I would change.

16192 – Supporting Organizational Efficiency and Agility: Models, Languages and Software Systems | Dagstuhl Seminar | http://www.dagstuhl.de/16192

The best aspect was the diversity of participants – from rather formal computer science, over information systems to management. That resulted in discussions which accounted for different perspectives on the subject – an approach which seems to be necessary to develop grounded ideas of future enterprise systems. It was also helpful that most of the participants were outstanding scholars of their field who contributed deep insights – that was the case for the few practionners as well. Unfortunately, there was only one representative of management science among the participants. As the organizers reported, they had invited more, but others had not accepted the invitation.

16232 – Fair Division | Dagstuhl Seminar | http://www.dagstuhl.de/16232

Best: the seminar is well-organized to the finest detail. The talk schedule, the randomized seats in the meals, the personalized food, the social activities – every aspect is well-planned. Worst: too many new research ideas to think of. I need at least 20 new students to study all the new research directions I came up with during the seminar...

16241 – Graph Polynomials: Towards a Comparative Theory | Dagstuhl Seminar | http://www.dagstuhl.de/16241

Best aspect: Dagstuhl is a perfect environment for thinking and talking to colleagues. The meals and lodging were so convenient that there were no interruptions or distractions from thinking and talking.

16251 – Information-centric Networking and Security | Dagstuhl Seminar | http://www.dagstuhl.de/16251

Networking and security researchers collaborating on problems of common interest. Also, there was a nice mix of European and US researchers, plus a similarly good mix of academia and industry.

16252 – Engineering Academic Software | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16252

Excellent mix of people from different communities.

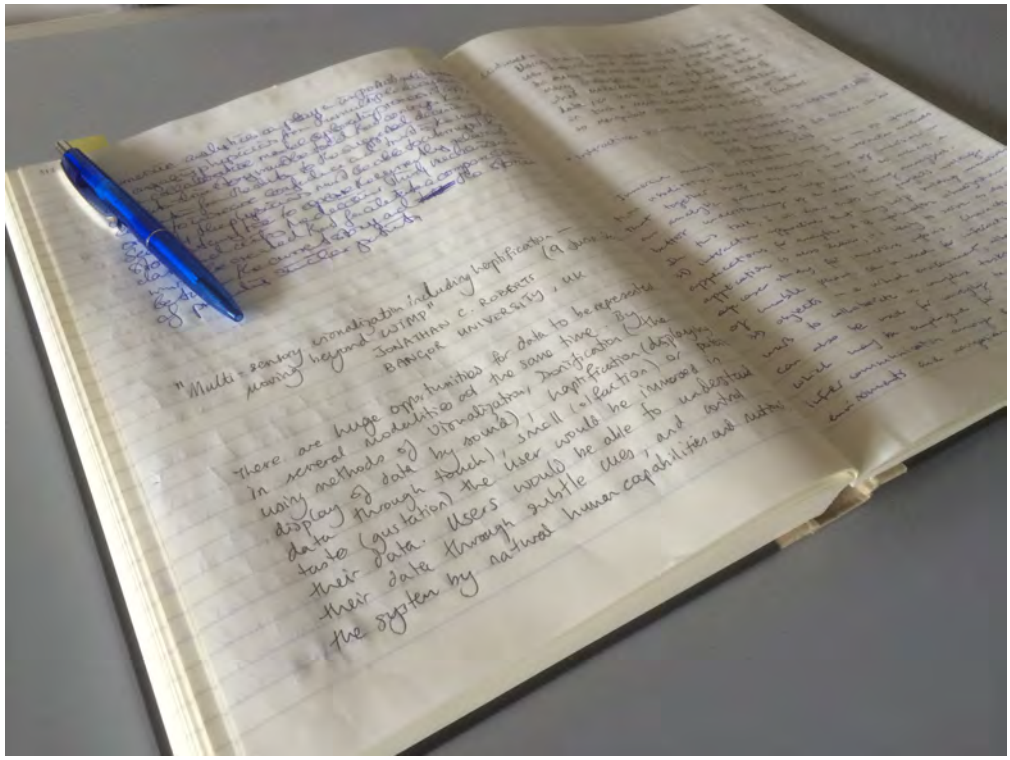

Fig. 5.1
„#ImmersiveDagstuhl [...] Just hand-written abstract on #multisensoryvis for @dagstuhl [...] This is a great #tradition“
Twitter post by 16231 Dagstuhl Seminar participant Jonathan C. Roberts. https://twitter.com/jcrbrts/status/741213059289616384. Photo courtesy of Jonathan C. Roberts.

Resonanz zur Bibliographiedatenbank dblp

5.4

Feedback on the dblp Computer Science Bibliography

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

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.

Manfred Jeusfeld (CEUR-WS Team, University of Skövde, Sweden)
dblp | https://ceurws.wordpress.com/2017/01/05/ceur-ws-enters-into-2017/

Last but not least, our thanks go to the DBLP team, who indexed CEUR-WS almost from the very beginning. Without DBLP's support, CEUR-WS would not have risen. Thank you, Michael Ley!

Carl Witt (Humboldt University, Berlin, Germany)
dblp | https://twitter.com/carl_witt/status/757914362279522304

@dblp_org is vastly useful for my literature research, thanks!

Ashish Sureka (ABB Corporate Research Center, Bangalore, India)
dblp | https://twitter.com/ashish_sureka/status/730465461452873728

DBLP snapshot data has helped us (@AiranSwati @pyNitish) conduct several Bibliometrics studies. Thanks to @dblp_org

Jochen L. Leidner (Thomson Reuters, Director of Research)
dblp | https://twitter.com/jochenleidner/status/730330162248728576

In May DBLP reached 3,333,333 publications. Congrats, Michael Ley and team! #research #science #computing

Eijiro Sumii (Tohoku University, Sendai, Japan)
dblp | https://twitter.com/esumii/status/722440236211437569

thanks for all the great service to the community!

tim 監督 (@monoids)
dblp | https://twitter.com/monoids/status/688801713638256641

feels good when a new entry finally shows up on my dblp

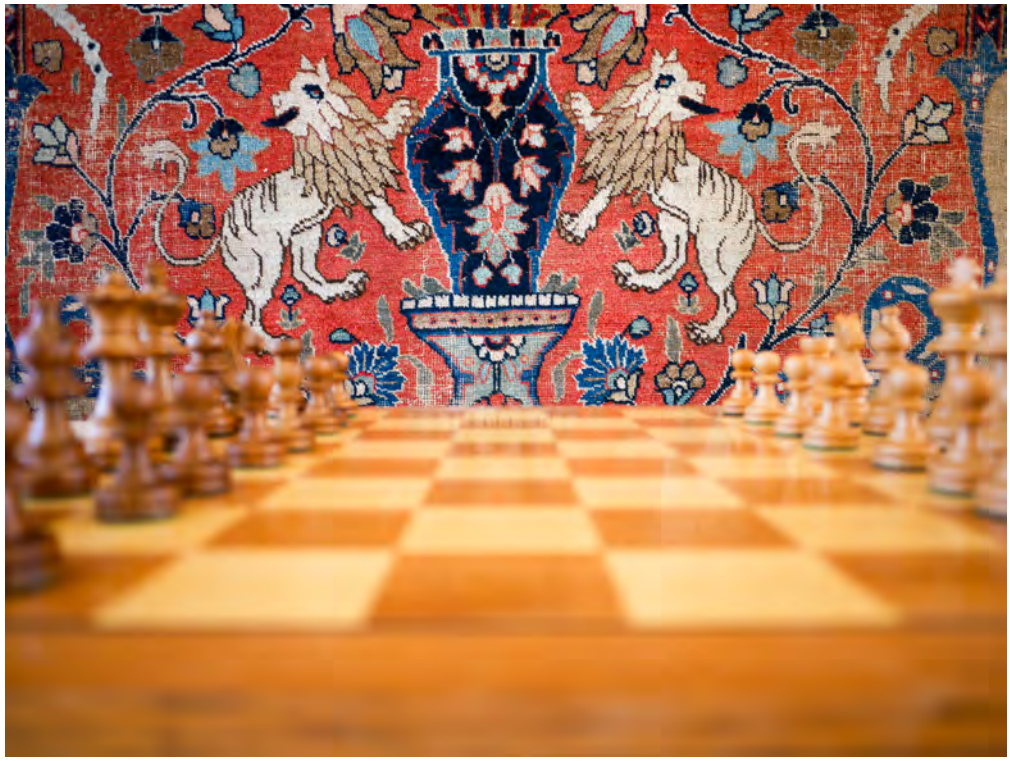

Fig. 5.2
Dagstuhl Seminar — what a great idea Blog post by 16072 Dagstuhl Seminar participant Jan Erik Moström. http://mostrom.eu/2016/03/19/dagstuhl-seminar-what-a-great-idea/. Photo courtesy of Jan Erik Moström.

6 Die Seminare in 2016
The 2016 Seminars

Applications, Interdisciplinary Work

- Analysis, Interpretation and Benefit of User-Generated Data: Computer Science Meets Communication Studies (16141)
- Assessing Learning in Introductory Computer Science (16072)
- Automotive User Interfaces in the Age of Automation (16262)
- Computational Challenges in Cooperative Intelligent Urban Transport (16091)
- Computational Music Structure Analysis (16092)
- Eyewear Computing – Augmenting the Human with Head-mounted Wearable Assistants (16042)
- Fresh Approaches to Business Process Modeling (16191)
- Immersive Analytics (16231)
- Modeling and Analysis of Semiconductor Supply Chains (16062)
- Next Generation Sequencing – Algorithms, and Software For Biomedical Applications (16351)
- Reproducibility of Data-Oriented Experiments in e-Science (16041)
- Rethinking Experimental Methods in Computing (16111)
- Supporting Organizational Efficiency and Agility: Models, Languages and Software Systems (16192)
- Uncertainty Quantification and High Performance Computing (16372)

Artificial Intelligence, Computational Linguistics

- Automated Algorithm Selection and Configuration (16412)
- Engineering Moral Agents – from Human Morality to Artificial Morality (16222)
- Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR) (16442)

Cryptography, Security, Privacy

- Assessing ICT Security Risks in Socio-Technical Systems (16461)
- Foundations of Secure Scaling (16342)
- Hardware Security (16202)
- Modern Cryptography and Security: An Inter-Community Dialogue (16051)
- Network Attack Detection and Defense – Security Challenges and Opportunities of Software-Defined Networking (16361)
- Privacy and Security in Smart Energy Grids (16032)
- Public-Key Cryptography (16371)
- Symmetric Cryptography (16021)

Data Structures, Algorithms, Complexity

- Algebraic Methods in Computational Complexity (16411)
- Algorithmic Foundations of Programmable Matter (16271)
- Algorithmic Methods for Optimization in Public Transport (16171)
- Algorithms and Effectivity in Tropical Mathematics and Beyond (16482)
- Algorithms for Optimization Problems in Planar Graphs (16221)
- Beyond-Planar Graphs: Algorithmics and Combinatorics (16452)
- Coding Theory in the Time of Big Data (16321)
- Computation over Compressed Structured Data (16431)
- Data Structures and Advanced Models of Computation on Big Data (16101)
- Evolution and Computing (16011)
- Fair Division (16232)
- Graph Polynomials: Towards a Comparative Theory (16241)
- Pattern Avoidance and Genome Sorting (16071)
- SAT and Interactions (16381)
- Scheduling (16081)
- Structure and Hardness in P (16451)
- Topological Methods in Distributed Computing (16282)

Databases, Information Retrieval, Machine Learning, Data Mining

- Data, Responsibly (16291)
- Foundations of Data Management (16151)
- Foundations of Unsupervised Learning (16382)
- Machine Learning for Dynamic Software Analysis: Potentials and Limits (16172)
- Natural Language Argumentation: Mining, Processing, and Reasoning over Textual Arguments (16161)
- New Directions for Learning with Kernels and Gaussian Processes (16481)
- Tensor Computing for Internet of Things (16152)

Distributed Computation, Networks, Architecture, Systems

- Adaptive Isolation for Predictability and Security (16441)
- Dark Silicon: From Embedded to HPC Systems (16052)
- Global Measurements: Practice and Experience (16012)
- Information-centric Networking and Security (16251)
- Network Latency Control in Data Centres (16281)
- QoE Vadis? (16472)

Geometry, Image Processing, Graphics, Visualization

- Data-Driven Storytelling (16061)
- Geometric and Graph-based Approaches to Collective Motion (16022)
- Inpainting-Based Image Compression (16462)
- Integration of Expert Knowledge for Interpretable Models in Biomedical Data Analysis (16261)
- Multidisciplinary Approaches to Multivalued Data: Modeling, Visualization, Analysis (16142)

Software Technology, Programming Languages

- Engineering Academic Software (16252)
- From Theory to Practice of Algebraic Effects and Handlers (16112)
- Integrating Process-Oriented and Event-Based Systems (16341)
- Managing Technical Debt in Software Engineering (16162)
- Programming Language Techniques for Incremental and Reactive Computing (16402)

Verification, Logic, Formal Methods, Semantics

- Concurrency with Weak Memory Models: Semantics, Languages, Compilation, Verification, Static Analysis, and Synthesis (16471)
- Language Based Verification Tools for Functional Programs (16131)
- Robustness in Cyber-Physical Systems (16362)
- Symbolic-Numeric Methods for Reliable and Trustworthy Problem Solving in Cyber-Physical Domains (16491)
- Synergies among Testing, Verification, and Repair for Concurrent Programs (16201)
- Universality of Proofs (16421)
- Well Quasi-Orders in Computer Science (16031)

6.1 Evolution and Computing

Organizers: Nick Barton, Per Kristian Lehre, and Nisheeth K. Vishnoi

Seminar No. 16011

Date: January 4–8, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.1

Participants: Ellen Baake, Nick Barton, Arnab Bhattacharyya, Erick Chastain, Duc-Cuong Dang, Harold P. de Vladar, Benjamin Doerr, Carola Doerr, Tobias Friedrich, Paulien Hogeweg, Kavita Jain, Timo Kötzing, Joachim Krug, Per Kristian Lehre, Adi Livnat, Kurt Mehlhorn, Tiago Paixao, Ioannis Panageas, Jorge Perez Heredia, Georgios Piliouras, Adam Prugel-Bennett, Jonathan L. Shapiro, Piyush Srivastava, Damian Mateusz Straszak, Dirk Sudholt, Andrew M. Sutton, Barbora Trubenova, Paul Valiant, Nisheeth K. Vishnoi, Thomas Wiehe, Carsten Witt, Xin Yao

Biological evolution has produced an extraordinary diversity of organisms, even the simplest of which is highly adapted, with multiple complex structures. Dynamic structures at even higher levels emerge from collective and social behaviour. These phenomena have traditionally been studied in population genetics, ecology and related disciplines.

However, theoretical computer scientists, endowed with a wide variety of tools, have recently made progress in describing and characterising the computational capabilities of evolution, analyzing natural algorithms, obtaining quantitative bounds for evolutionary models and understanding the role of sex in evolution. The field of evolutionary computation has found that many innovative solutions to optimisation and design problems can be achieved by simulating living processes, such as evolution via random variation and selection, or social behaviour in insects. Researchers in evolutionary computation have recently started applying techniques from theoretical computer science to analyze the optimization time of natural algorithms.

To further the connections and consolidate this burgeoning new discipline, this Dagstuhl seminar brought together participants from the population genetics, mathematical biology, theoretical computer science, and evolutionary computation communities. The seminar opened with a round of introductions, followed by five introductory talks presenting the perspectives of the disciplines attending. Benjamin Doerr introduced runtime analysis of evolutionary algorithms, Paul Valiant discussed evolution from the perspective of learning, Joachim Krug and Nick Barton introduced population genetics, and Nisheeth Vishnoi discussed evolutionary processes from the perspective of theoretical computer science. In addition to talks contributed by participants, there were several breakout sessions on topics identified during the seminar.

The organisers would like to thank the Dagstuhl team and all the participants for making the seminar a success.

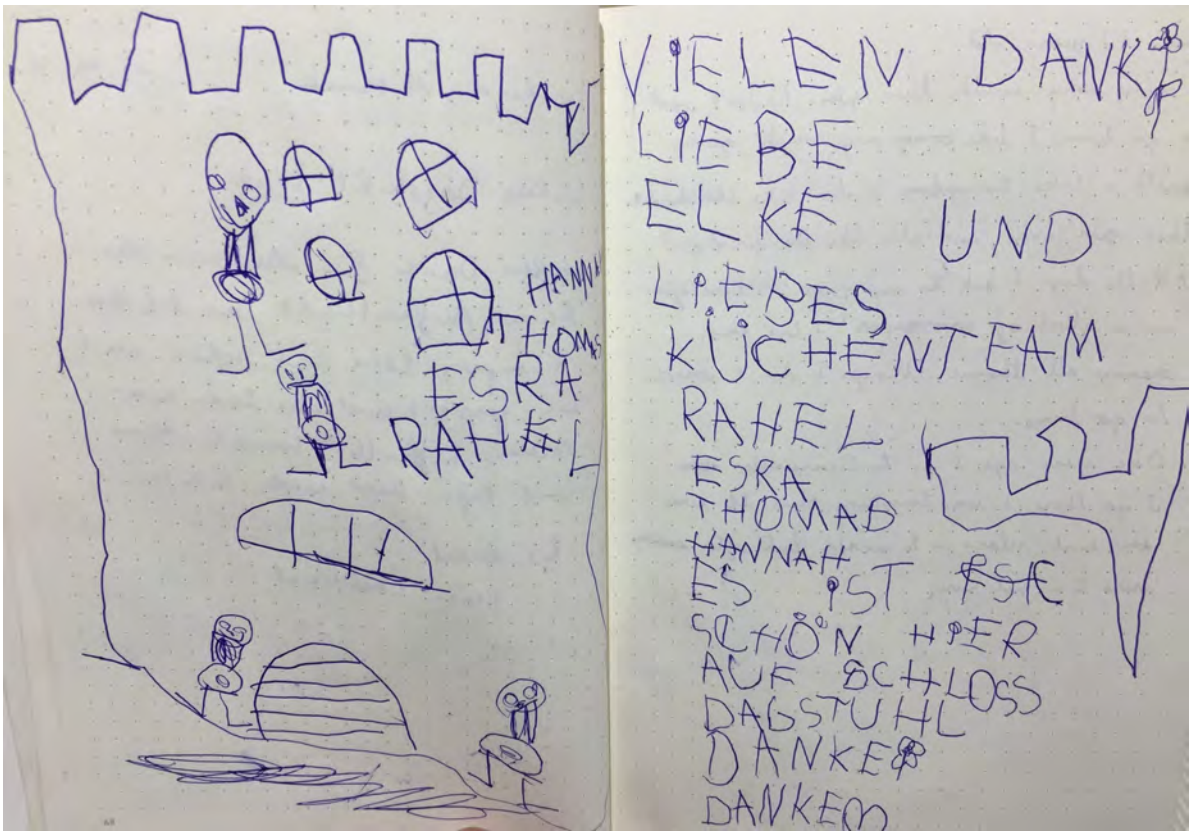

Fig. 6.1
Drawing for the Dagstuhl children's guest book.

6.2 Global Measurements: Practice and Experience

Organizers: Vaibhav Bajpai, Arthur W. Berger, Philip Eardley, Jörg Ott, and Jürgen Schönwälder

Seminar No. 16012

Date: January 4–7, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.15

Participants: Vaibhav Bajpai, Arthur W. Berger, Georg Carle, Renata Cruz Teixeira, Philip Eardley, Markus Fiedler, Phillipa Gill, Oliver Hohlfeld, Steffie Jacob Eravuchira, Daniel Karrenberg, Mirja Kühlewind, Andri Lareida, Jukka Manner, Ian Robin Marsh, Al Morton, Jörg Ott, Colin Perkins, Philipp Richter, Jair Santanna, Jürgen Schönwälder, Henning Schulzrinne, Varun Singh, Burkhard Stiller, Srikanth Sundaresan, Brian Trammell, Roland van Rijswijk-Deij

Several large-scale Internet measurement platforms have been deployed during the last years in order to understand how the Internet is performing, to observe how it is evolving, and to determine where failures or degradations occur. Examples are the CAIDA Archipelago (Ark) platform [6] (used for Internet topology discovery and detecting congestion on interdomain links), the SamKnows platform [4] (used by regulators and network operators to study network performance), the RIPE Atlas platform [3, 5] (that provides measurement services to network operators and researchers), the Netradar system [8] (for performing wireless performance measurements), and the BISmark project [9]. European collaborative research projects lately have been working on a Measurement Plane (mPlane) [10] and how to incorporate measurement results into network management systems (e.g., Leone) [2]. Related projects (e.g., Flamingo) [1] are increasingly working with measurement data from these platforms. Large-scale measurements are meanwhile also used to drive network operations or to dynamically adjust how services are delivered to customers. Content Delivery Network (CDN) providers use measurement data to optimize content caches and to tune load balancing algorithms. One key challenge is that global Internet measurement systems can generate large amounts of data that need to be processed to derive relevant information.

This seminar (#16012) was a followup of the Dagstuhl seminar on Global Measurement Frameworks (#13472) [7]. The main focus of the first seminar was an exchange of ideas on the development of global measurement infrastructures, frameworks and associated metrics. Some of this work is now further pursued in standardization bodies [4] such as the IETF Large-Scale Measurement of Broadband Performance (LMAP) working group and the Broadband Forum. The goal of this followup seminar was to focus on the experience obtained with different metrics, tools, and data analysis techniques. It provided a forum for researchers to exchange their experience with different practices to conduct global measurements. The aim was to identify what works well in certain contexts, what has proven problematic in other contexts, and identify open issues that need further research. The seminar approached this by looking at three distinct dimensions: (a) Measurement metrics, (b) data processing technologies and (c) data analysis methodologies. Some key questions were:

1. Which metrics have been found useful for measuring Quality of Experience (QoE) of certain classes of services? Which metrics have been found problematic? Is it possible to find indicators for good metrics and problematic metrics?
2. Which technologies have been found useful for storing and processing large amounts of measurement data? Which technologies were found to be problematic? Are there new promising technologies that may be used in the future? What are the specific requirements for dealing with large-scale measurement data and how do they relate to or differ from other big data applications?
3. Which data analysis techniques have been found to be useful? Which data analysis techniques have been found to be problematic? Are there any novel promising techniques that need further research and development?

Although at the seminar the participants chose to organize the discussions on more general topics than these specific questions, during the discussions most of these questions were addressed to one degree or another.

References

1 Flamingo – Management of the Future Internet. http://www.fp7-flamingo.eu. [Online; accessed 18-January-2016].

2 Leone – From Global Measurements to Local Management. http://www.leone-project.eu. [Online; accessed 18-January-2016].

3 RIPE Atlas: A Global Internet Measurement Network. *Internet Protocol Journal*, September 2015. URL: http://ipj.dreamhosters.com/wp-content/uploads/2015/10/ipj18.3.pdf.

4 V. Bajpai and J. Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *Communications Surveys Tutorials, IEEE*, 17(3):1313–1341, thirdquarter 2015. http://dx.doi.org/10.1109/COMST.2015.2418435.

5 Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. *SIGCOMM Comput. Commun. Rev.*, 45(3):35–42, July 2015. http://dx.doi.org/10.1145/2805789.2805796.

6 kc claffy. The 7th Workshop on Active Internet Measurements (AIMS7) Report. *SIGCOMM Comput. Commun. Rev.*, 46(1):50–57, January 2016. http://dx.doi.org/10.1145/2875951.2875960.

7 Philip Eardley, Marco Mellia, Jörg Ott, Jürgen Schönwälder, and Henning Schulzrinne. Global Measurement Framework (Dagstuhl Seminar 13472). *Dagstuhl Reports*, 2014. http://dx.doi.org/http://dx.doi.org/10.4230/DagRep.3.11.144.

8 S. Sonntag, J. Manner, and L. Schulte. Netradar – Measuring the wireless world. WiOpt'13, May 2013.

9 Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter De Donato. BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks. USENIX ATC'14, pages 383–394, 2014. URL: http://dl.acm.org/citation.cfm?id=2643634.2643673.

10 B. Trammell, P. Casas, D. Rossi, A. Bar, Z. Houidi, I. Leontiadis, T. Szemethy, and M. Mellia. mPlane: an Intelligent Measurement Plane for the Internet. *IEEE Communications Magazine*, 52(5):148–156, May 2014. http://dx.doi.org/10.1109/MCOM.2014.6815906.

6.3 Symmetric Cryptography

Organizers: Frederik Armknecht, Tetsu Iwata, Kaisa Nyberg, and Bart Preneel

Seminar No. 16021

Date: January 10–15, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.34

Participants: Elena Andreeva, Frederik Armknecht, Daniel J. Bernstein, Eli Biham, Alex Biryukov, Andrey Bogdanov, Anne Canteaut, Benoît Cogliati, Joan Daemen, Itai Dinur, Orr Dunkelman, Henri Gilbert, Jian Guo, Matthias Hamann, Tetsu Iwata, Jérémy Jean, Antoine Joux, Dmitry Khovratovich, Matthias Krause, Nils Gregor Leander, Jooyoung Lee, Gaëtan Leurent, Stefan Lucks, Willi Meier, Bart Mennink, Kazuhiko Minematsu, Nicky Mouha, Chanathip Namprempre, Mridul Nandi, Ivica Nikolic, Kaisa Nyberg, Jacques Patarin, Léo Paul Perrin, Bart Preneel, Christian Rechberger, Yu Sasaki, Ernst Schulte-Geers, Adi Shamir, John Steinberger, Marc Stevens, Tyge Tiessen, Meiqin Wang, Xianyun Wang, Kan Yasuda

One lesson learned from the Snowden leaks is that digital systems can never be fully trusted and hence the security awareness of citizens has increased substantially. Whenever digital data is communicated or stored, it is subject to various attacks. One of the few working countermeasures are the use of cryptography. As Edward Snowden puts it: "*Encryption works. Properly implemented strong crypto systems are one of the few things that you can rely on.*".[36]

Consequently it holds that although modern cryptography addresses a variety of security challenges, efficiently protecting the enormous amount of daily electronic communication represents a major challenge. Here, symmetric cryptography is especially highly relevant not only for academia, but also for industrial research and applications.

Although symmetric cryptography has made enormous progress in the last couple of decades, for several reasons regularly new insights and challenges are evolving. In the past, the AES competition was led by US NIST to standardize a next generation block cipher to replace DES. Similar competitions, such as the eSTREAM and the SHA-3 competition, resulted in new standard algorithms that meet public demands. The outcome of the projects are practically used in our daily lives, and the fundamental understanding of the cryptographic research community of these primitives has been increased significantly.

While this seminar concentrates in general on the design and analysis of symmetric cryptographic primitives, special focus has been put on the following two topics that we explain in more detail below:

1. Authenticated encryption
2. Even-Mansour designs

Authenticated Encryption. Today the central research question is the construction of schemes for *authenticated* encryption. This symmetric primitive efficiently integrates the protection of secrecy and integrity in a single construction. The first wave of solutions resulted in several widely used standards, including CCM and GCM standardized by NIST, and the EAX-prime standardized by ANSI. However, it turns out that these constructions are far from optimum in terms of performance, security, usability, and functionality. For instance a stream of data cannot be protected with CCM, as the length of the entire input has to be known in advance. The security of GCM heavily relies on the existence of data called a nonce, which is supposed to never be repeated. Indeed, the security of GCM is completely lost once the nonce is repeated. While it is easy to state such a mathematical assumption, experience shows that there are many practical cases where realizing this condition is very hard. For instance the nonce may repeat if a crypto device is reset with malice aforethought, or as a consequence of physical attacks on the device. Furthermore, weak keys were identified in GCM, and the security of EAX-prime is questionable.

Thus there is a strong demand for secure and efficient authenticating encryption scheme. As a consequence, the CAESAR project (Competition for Authenticated Encryption: Security, Applicability, and Robustness) has been initiated.[37] The goal of the project is to identify a portfolio of authenticated encryption schemes that (1) offer advantages over GCM/CCM and (2) are suitable for widespread adoption. The deadline of the submission was March 15, 2014, and the project attracted a total of 56 algorithms from 136 designers from all over the world. There are plenty of innovative designs with attractive features, and the final portfolio is planned to be announced at the end of 2017.

[36] See http://techcrunch.com/2013/06/17/encrypting-your-email-works-says-nsa-whistleblower-edward-snowden/.
[37] See http://competitions.cr.yp.to/caesar.html for details.

This seminar took place in the middle of the CAESAR competition; it is two years from the submission deadline and we have about two years until the announcement of the final portfolio. Therefore, it was a perfect point in time to sum up the research done so far, to exchange ideas and to discuss future directions.

Even-Mansour Designs. Another strong trend in the current symmetric key cryptography is related to the so-called *Even-Mansour designs*. This design paradigm was proposed in 1991 and can be seen as the abstraction of the framework adopted in the design of AES. This general design framework iterates r times the xor of a key and a public permutation. The design framework is highly relevant in practice, and it has been adopted in a variety of recent hash functions, block ciphers, and even in the underlying primitive of several CAESAR submissions. Despite its long history of practical use, the community has so far failed to develop a complete understanding of its security. From a theoretical viewpoint, the original proposal was accompanied with a proof of security, dealing with the case of $r = 1$ iteration.

Only 20 years after the initial proposal, in 2012, a bound was proven for the security of $r = 2$ iterations. In 2014, the question was solved to cover the general case of r iterations. However, these results only deal with the simple case of distinguishing attack on a single, unknown key setting. Its security in more advanced, yet practically relevant security models, such as the related-key setting or the chosen/known-key setting, is largely unexplored.

Another problem here is that the theoretical analysis assumes that the permutation used therein is ideal and the keys are ideally random, which is not the case for practical constructions. This implies that the theoretical results do not directly translate into the practical constructions, and the security analysis has to be repeated for each constructions.

Summing up, Evan-Mansour designs represent a fruitful and challenging area of research, that hopefully will lead to a fundamental understanding of iterated constructions and ultimately to more efficient and more secure ciphers.

Seminar Program. The seminar program consists of the presentations about the above topics, and relevant areas of symmetric cryptography, including new cryptanalytic techniques and new designs. Furthermore, there were three discussion sessions. In “discussion on attacks,” we discussed what constitutes a valid cryptographic attack in light of weak key classes, “discussion on secret agency crypto standards” was about cryptography developed by secret agencies, and there was a discussion session about the ongoing CAESAR project.

6.4 Geometric and Graph-based Approaches to Collective Motion

Organizers: Giuseppe F. Italiano, Marc van Kreveld, Bettina Speckmann, and Guy Theraulaz

Seminar No. 16022

Date: January 10–15, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.55

Participants: Martin Beye, Karl Bringmann, Kevin Buchin, Maike Buchin, Oliver Burkhard, Somayeh Dodge, Matt Duckham, Anael Engel, Ramon Escobedo, Brittany Terese Fasy, Sándor Fekete, Luca Giuggioli, Giuseppe F. Italiano, Maarten Löffler, Richard Philipp Mann, Martin Nöllenburg, Tim Ophelders, Nicholas Ouellette, Andrea Perna, Bettina Speckmann, Frank Staals, Guy Theraulaz, Goce Trajcevski, Johan van de Koppel, Marc van Kreveld, Kevin Verbeek, Robert Weibel, Carola Wenk

A trajectory is a time-stamped sequence of locations which represents the movement of entities in space. Trajectories are often created by sampling GPS locations and attaching a time-stamp, but they can also originate from RFID tags, video, or radar analysis. Huge data sets exist for entities as diverse as birds, deer, traveling humans, sports players, vehicles, and hurricanes.

During recent years analysis tools for trajectory data have been developed within the areas of GIScience and algorithms. Analysis objectives include clustering, performing similarity analysis, segmenting a trajectory into characteristic sub-trajectories, finding patterns like flocking, and several others. Since these computations are mostly spatial, algorithmic solutions have been developed in the areas of computational geometry and GIScience. Although trajectories store only the location of a single point of reference on a moving entity, this is acceptable for the common large-scale analysis tasks. However, for the study of more complex phenomena like interaction and collective motion, it is often insufficient and the basic trajectory representation must be extended.

Simultaneously, in the area of ecology the study of motion of animals has also become a topic of increasing interest. Many animal species move in groups, with or without a specific leader. The motivation for motion can be foraging, escape from predators, changing climate, or it can be unknown. The mode of movement can be determined by social interactions, energy efficiency, possibility of discovery of resources, and of course the natural environment. The more fascinating aspects of ecology include interaction between entities and collective motion. These are harder to grasp in a formal manner, needed for modelling and automated analysis.

The seminar brought together a group of enthusiastic researchers with a diverse background. To create a shared body of knowledge the seminar featured a number of survey talks that were planned early in the week. The survey talks were rather engaging: the audience learned for instance at what scale one should look at a painting of Van Gogh, how bats chase each other, what size of clumps mussels make and why, and how to interact with a computational geometer.

Probably the main research result was a momentum started up by interaction and awareness of an exciting direction of research where a lot can still be accomplished.

More specific research accomplishments included a methodology for evaluating whether fish or other animals have their movement mostly influenced by closest neighbors, and how to reconstruct movement just based on counts at different time steps.

Fig. 6.2
Drawing for the Dagstuhl children's guest book by Eric (2), son of Dagstuhl Seminar 16261 participants Elke K. Markert and Alexei Vazquez.

6.5 Well Quasi-Orders in Computer Science

Organizers: Jean Goubault-Larrecq, Monika Seisenberger, Victor Selivanov, and Andreas Weiermann

Seminar No. 16031

Date: January 17–22, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.69

Participants: Lev D. Beklemishev, Andreas R. Blass, Raphael Carroy, Raphael Chane-Yack-Fa, Matthew de Brecht, Dick De Jongh, Mirna Dzamonja, Alain Finkel, Thomas E. Forster, Willem L. Fouché, Christian Glasser, Lev Gordeev, Jean Goubault-Larrecq, Simon Halfon, Peter Hertling, Takayuki Kihara, Julia Knight, Oleg V. Kudinov, Dominique Larchey-Wendling, Chun-Hung Liu, Alberto Marcone, Roland Meyer, Kenji Miyamoto, Mizuhito Ogawa, Arno Pauly, Yann Pequignot, Maurice Pouzet, Thomas Powell, Michael Rathjen, Diana Schmidt, Heinz Schmitz, Sylvain Schmitz, Philippe Schnoebelen, Peter M. Schuster, Helmut Schwichtenberg, Luc Segoufin, Monika Seisenberger, Victor Selivanov, Dieter Spreen, Christian Sternagel, Bill Wadge, Andreas Weiermann, Klaus Weihrauch, Daniel Wessel

Computer Science, being a huge and complex conglomerate of theoretical disciplines, technological advances and social methodologies, strongly needs unifying concepts and techniques. In particular, relevant mathematical concepts and theories are required. The notion of well quasi-order (or almost-full relation, if transitivity is not required – a notion preferred by some authors) was discovered independently by several mathematicians in the 1950-s and quickly evolved to a deep theory with many applications and remarkable results. Soon afterwards, well and better quasi-orders started to appear more and more frequently in different parts of theoretical computer science such as automata theory, term rewriting, verification of infinite-state systems, computations with infinite data, and others. Accordingly, an increasing number of researchers from different fields of computer science use notions and methods of Wqo-Theory. Therefore, it seemed to be the right time to have a broad discussion on how to speedup this process and to better understand the role of well quasi-orders in theoretical computer science.

Topics of the seminar

During this seminar we concentrated on the following four topics:

1. Logic and proofs
2. Automata and formal languages
3. Topological issues
4. Verification and termination problems

Logic and proofs

Well quasi-orders, originally introduced in algebra, soon played an important role in proof theory: Higman's Lemma and Kruskal's Theorem are examples of theorems that are not provable in Peano Arithmetic. Determining the proof-theoretic strength of these (types of) theorems, as well as classifying them in terms of Reverse Mathematics, constituted an important endeavor. The concept of a WQO naturally extends to the more complex concept of a better quasi-order (BQO) which deals with infinite structures. Again, the proof theoretic strength of theorems on BQOs has been/must be investigated, and the theorems themselves can be used for more sophisticated termination problems. One of the open challenges is the strength of Fraïssé's order type conjecture. Non-constructive proofs of this type of theorems (on WQOs) include proofs using the so-called minimal-bad-sequence argument. Investigating their strengths and also their computational content, via Friedman's A-translation or Gödel's Dialectica Interpretation, has led to interesting results. To optimize these techniques so that realistic programs can be extracted from these classical proofs, using bar recursion, update recursion, selection functions, etc., is ongoing work.

Automata and formal languages

Well quasi-orders have many-fold connections to automata theory and formal language theory. In particular, there are nice characterizations of regular and context-free languages in terms of well quasi-orders, some lower levels of the concatenation hierarchies admit characterizations in terms of the subword relation and its relatives. Such characterizations sometimes help in getting new results, say on decidability of some levels of the concatenation hierarchy (Glasser, Schmitz, Selivanov). The same applies to ω-languages, though in this case the relationships are less investigated.

On the topological level, it is known that Wadge reducibility (or reducibility by functions on ω-words computable by finite automata) are well quasi-orders on the class of ω-regular finite partitions of the Cantor space. Using some variants of the Kruskal theorem on quasi-orderings of labeled trees, Selivanov was able to completely characterize the corresponding partial order, obtaining

thus a complete extension of the Wagner hierarchy from sets of finite partitions.

The mentioned relationships between Wqo-theory and formal languages are currently not well systematized, and many natural questions remain open. Further insights in this topic is essential for the development of this field.

Topological issues

An important task in computing with infinite data is to distinguish between computable and non-computable functions and, in the latter case, to measure the degree of non-computability. Usually, functions are non-computable since they are not even continuous, hence a somewhat easier and more principal task is in fact to understand the degree of discontinuity of functions. This is achieved by defining appropriate hierarchies and reducibility relations.

In classical descriptive set theory, along with the well-known hierarchies, Wadge introduced and studied an important reducibility relation on subsets of the Baire space. As shown by van Engelen et al., von Stein, Weihrauch and Hertling, this reducibility of subsets of topological spaces can be generalized in various ways to a reducibility of functions on a topological space. In this way, the degrees of discontinuity of several important computational problems were classified. The transfer from sets to functions requires some notions and results of Wqo-theory in order to define and study hierarchies and reducibilities arising in this way.

Verification and termination problems

WQOs made their debut in computer science when Don Knuth suggested that Kruskal's Theorem might find an application in proving termination of programs. This was achieved a few years later by Nachum Dershowitz and the advent of recursive path orderings. Today, it is probably the area of software verification that provides the largest number of applications of WQOs in computer science. The decidability of coverability for well-structured transition systems (WSTS) crucially relies on the very properties of well quasi-orders. WSTS include Petri nets and their extensions, and more generally affine nets. They also include lossy channel systems, weak memory models, various process algebras, data nets, certain abstractions of timed Petri nets, and certain parametrized transition systems. The verification of new classes of transition systems prompts for new classes of WQOs. In addition to this, understanding the computational complexity of the resulting verification algorithms requires a finer analysis of minimal-bad-sequence arguments and their relation to hierarchies of recursive functions (Hardy, fast growing, etc.)

Overall, our seminar attracted 44 participants (10 from Germany, 22 from other European countries, 12 from Canada, Japan, Russia, South Africa, and USA) who contributed 33 talks. In addition, we included several problem sessions where we summarized all problems mentioned in the seminar. As a result of these sessions we give a list of open problems at the end of this report. Looking at the feedback the seminar was very well received amongst the participants. Positively mentioned was that the seminar involved "people from different backgrounds" who "can still share interest", or in other words "hearing people from different research areas discuss similar questions", and that "one week is too short :-)". Thoroughly enjoyed was also our two hour long walk in the snow on Wednesday afternoon. The great success of the seminar is not only due to the participants, but also to the staff in Saarbrücken and Dagstuhl, who did a splendid job in facilitating the seminar and making our stay a very pleasant one. Special thanks go to Susanne Bach-Bernhard for all the interaction related to the organization of the seminar and to Jutka Gasiorowski for her support in producing the report.

6.6 Privacy and Security in Smart Energy Grids

Organizers: George Danezis, Stefan Katzenbeisser, Christiane Peters, and Bart Preneel

Seminar No. 16032

Date: January 17–20, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.99

Participants: Nikita Borisov, George Danezis, Benessa Defend, Dominik Engel, Zekeriya Erkin, Benedikt Gierlichs, Stefan Katzenbeisser, Florian Kerschbaum, Erwin Kooi, Klaus Kursawe, Éireann Leverett, Carlos Montes Portela, Mustafa Mustafa, Christiane Peters, Erik Poll, Bart Preneel, Ahmad-Reza Sadeghi, Kazue Sako, Matthias Schunter, Neeraj Suri, Makoto Takahashi, Pol Van Aubel, Ingrid Verbauwhede, Jos Weyers

Smart electricity grids augment the electricity distribution network with modern communications and computerized control to improve efficiency, reliability, and security of electricity distribution, and more flexible production. This initiative has been greeted by consumers and utilities not only with enthusiasm but also concern. Consumers worry about their privacy. Utilities worry about the security of their assets.

Consumer organizations across the globe protested against smart meters and smart homes collecting all their data, warning that security breaches in the databases of the utilities would expose privacy-critical data to attackers, or open to secondary uses leading to increased insurance premiums, behavioral advertising or privacy invasion. These outcries and reactions have triggered academics and industry to look into designing privacy friendly architectures for smart metering.

The seminar 16032 in particular focused on the two use cases of smart charging of electric vehicles (EVs) and distribution automation. The seminar discussed these use cases with respect to the following challenges:

- security architectures,
- secure and privacy-friendly communication, and
- hardware and software security for constrained devices in the smart grid.

Smart Charging: Charging of electric vehicles is the next big challenge for privacy and security researchers: smart charging algorithms try to minimize loads on the grid by collecting various kinds of customer data, making it easy to reserve charging spots and book charge frequencies using smart-phone apps. The main motivation behind smart charging is to save copper for cables to match the load demands, given that an electric vehicle draws as much as a full household. Cables are designed to satisfy the demands at peak times. So profiling customers helps to foresee these demands and to calculate the cost of the needed grid infrastructure. Moreover, the cable designs use prediction algorithms to optimize loads, while assigning low priority to privacy issues, security architectures, and secure communication protocols.

Distribution Automation: Another problem lies in the task of automated electricity distribution. In a smart grid, safety critical events in transformer stations can be monitored and operated remotely. Adding communication also exposes assets to new vulnerabilities and attacks. Grid components are controlled by dedicated devices that pose a challenge in terms of their storage and computation capacities. Moreover, as with any critical infrastructure, security often conflicts with safety. As a consequence security often does not play any role in the design of communication protocols and devices, supported by the argument that most devices reside in physically protected substations. However, providing such physical security is expensive and hackers do not need physical access to the grid operator sites if they are connected to the utility's IT network.

The goal of this seminar was thus (i) to raise awareness of these critical problems affecting every European citizen now or at least in the foreseeable future, and (ii) to bring together academic researchers as well as utility experts in order to start an open dialogue on smart grid privacy and security problems and potential solutions to support customers and utilities.

Fig. 6.3
Dagstuhl Seminar — what a great idea Blog post by 16072 Dagstuhl Seminar participant Jan Erik Moström. http://mostrom.eu/2016/03/19/dagstuhl-seminar-what-a-great-idea/. Photo courtesy of Jan Erik Moström.

6.7 Reproducibility of Data-Oriented Experiments in e-Science

Organizers: Juliana Freire, Norbert Fuhr, and Andreas Rauber

Seminar No. 16041

Date: January 24–29, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.108

Participants: Vanessa Braganholo, Fernando Chirigati, Christian Collberg, Shane Culpepper, David De Roure, Arjen P. de Vries, Jens Dittrich, Nicola Ferro, Juliana Freire, Norbert Fuhr, Daniel Garijo, Carole Goble, Kalervo Järvelin, Noriko Kando, Randall J. LeVeque, Matthias Lippold, Bertram Ludäscher, Mihai Lupu, Tanu Malik, Rudolf Mayer, Alistair Moffat, Kevin Page, Raul Antonio Palma de Leon, Martin Potthast, Andreas Rauber, Paul Rosenthal, Claudio T. Silva, Stian Soiland-Reyes, Benno Stein, Rainer Stotzka, Evelyne Viegas, Stefan Winkler-Nees, Torsten Zesch, Justin Zobel

In many subfields of computer science, experiments play an important role. Besides theoretical properties of algorithms or methods, their effectiveness and performance often can only be validated via experimentation. In most of these cases, the experimental results depend on the input data, settings for input parameters, and potentially on characteristics of the computational environment where the experiments were designed and run. Unfortunately, most computational experiments are specified only informally in papers, where experimental results are briefly described in figure captions; the code that produced the results is seldom available.

This has serious implications. Scientific discoveries do not happen in isolation. Important advances are often the result of sequences of smaller, less significant steps. In the absence of results that are fully documented, reproducible, and generalizable, it becomes hard to re-use and extend these results. Besides hindering the ability of others to leverage our work, and consequently limiting the impact of our field, the absence of reproducibility experiments also puts our reputation at stake, since reliability and validity of empiric results are basic scientific principles.

Reproducible results are not just beneficial to others – in fact, they bring many direct benefits to the researchers themselves. Making an experiment reproducible forces the researcher to document execution pathways. This in turn enables the pathways to be analyzed (and audited). It also helps newcomers (e.g., new students and post-docs) to get acquainted with the problem and tools used. Furthermore, reproducibility facilitates portability, which simplifies the dissemination of the results. Last, but not least, preliminary evidence exists that reproducibility increases impact, visibility and research quality.

However, attaining reproducibility for computational experiments is challenging. It is hard both for authors to derive a compendium that encapsulates all the components (e.g., data, code, parameter settings, environment) needed to reproduce a result, and for reviewers to verify the results. There are also other barriers, from practical issues – including the use of proprietary data, software and specialized hardware, to social – for example, the lack of incentives for authors to spend the extra time making their experiments reproducible.

This seminar brought together experts from various sub-fields of Computer Science as well as experts from several scientific domains to create a joint understanding of the problems of reproducibility of experiments, discuss existing solutions and impediments, and propose ways to overcome current limitations.

Beyond a series of short presentations of tools, state of the art of reproducibility in various domains and "war stories" of things not working, participants specifically explored ways forward to overcome barriers to the adoption of reproducibility. A series of break-out sessions gradually built on top of each other, (1) identifying different types of repeatability and their merits; (2) the actors involved and the incentives and barriers they face; (3) guidelines for actors (specifically editors, authors and reviewers) on how to determine the level of reproducibility of papers and the merits of reproduction papers; and (4) the specific challenges faced by user-oriented experimentation in Information Retrieval.

This led to the definition of according typologies and guidelines as well as identification of specific open research problems. We defined a set of actions to reach out to stakeholders, notably publishers and funding agencies as well as identifying follow-up liaison with various reproducibility task forces in different communities including the ACM, FORCE11, STM, Science Europe.

The key message resulting from this seminar, copied from and elaborated in more detail in the full report is:

> *Transparency, openness, and reproducibility are vital features of science. Scientists embrace these features as disciplinary norms and values, and it follows that they should be integrated into daily research activities. These practices give confidence in the work; help research*

as a whole to be conducted at a higher standard and be undertaken more efficiently; provide verifiability and falsifiability; and encourage a community of mutual cooperation. They also lead to a valuable form of paper, namely, reports on evaluation and reproduction of prior work. Outcomes that others can build upon and use for their own research, whether a theoretical construct or a reproducible experimental result, form a foundation on which science can progress. Papers that are structured and presented in a manner that facilitates and encourages such post-publication evaluations benefit from increased impact, recognition, and citation rates.

Experience in computing research has demonstrated that a range of straightforward mechanisms can be employed to encourage authors to produce reproducible work. These include: requiring an explicit commitment to an intended level of provision of reproducible materials as a routine part of each paper's structure; requiring a detailed methods section; separating the refereeing of the paper's scientific contribution and its technical process; and explicitly encouraging the creation and reuse of open resources (data, or code, or both).

6.8 Eyewear Computing – Augmenting the Human with Head-Mounted Wearable Assistants

Organizers: Andreas Bulling, Ozan Cakmakci, Kai Kunze, and James M. Regh

Seminar No. 16042

Date: January 24–29, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.160

Participants: Andreas Bulling, Ozan Cakmakci, Rita Cucchiara, Steven K. Feiner, Kristen Grauman, Scott Greenwald, Sabrina Hoppe, Masahiko Inami, Shoya Ishimaru, Moritz Kassner, Koichi Kise, Kiyoshi Kiyokawa, Kai Kunze, Yin Li, Paul Lukowicz, Päivi Majaranta, Walterio W. Mayol-Cuevas, René Mayrhofer, Masashi Nakatani, Will Patera, Thies Pfeiffer, James M. Rehg, Philipp M. Scholl, Linda B. Smith, Gábor Sörös, Thad Starner, Julian Steil, Yusuke Sugano, Yuji Uema

Computing devices worn on the human body have a long history in academic and industrial research, most importantly in wearable computing, mobile eye tracking, and mobile mixed and augmented reality. In contrast to traditional systems, body-worn devices are always with the user and therefore have the potential to perceive the world and reason about it from the user's point of view. At the same time, given that on-body computing is subject to ever-changing usage conditions, on-body computing also poses unique research challenges.

This is particularly true for devices worn on the head. As humans receive most of their sensory input via the head, it is a particularly interesting body location for simultaneous sensing and interaction as well as cognitive assistance. Early egocentric vision devices were rather bulky, expensive, and their battery lifetime severely limited their use to short durations of time. Building on existing work in wearable computing, recent commercial egocentric vision devices and mobile eye trackers, such as Google Glass, PUPIL, and J!NS meme, pave the way for a new generation of "smart eyewear" that are light-weight, low-power, convenient to use, and increasingly look like ordinary glasses. This last characteristic is particularly important as it makes these devices attractive for the general public, thereby holding the potential to provide a research and product platform of unprecedented scale, quality, and flexibility.

While hearing aids and mobile headsets became widely accepted as head-worn devices, users in public spaces often consider novel head-attached sensors and devices as uncomfortable, irritating, or stigmatising. Yet with the advances in the following technologies, we believe eyewear computing will be a very prominent research field in the future:

- Increase in storage/battery capacity and computational power allows users to run eyewear computers continuously for more than a day (charging over night) gathering data to enable new types of life-logging applications.
- Miniaturization and integration of sensing, processing, and interaction functionality can enable a wide array of applications focusing on micro-interactions and intelligent assistance.
- Recent advances in real-life tracking of cognitive activities (e.g. reading, detection of fatigue, concentration) are additional enabling technologies for new application fields towards a quantified self for the mind. Smart eyewear and recognizing cognitive states go hand in hand, as naturally most research work in this field requires sensors.
- Cognitive scientists and psychologists have now a better understanding of user behavior and what induces behavior change. Therefore, smart eyewear could help users in achieving behaviour change towards their long term goals.

Eyewear computing has the potential to fundamentally transform the way machines perceive and understand the world around us and to assist humans in measurably and significantly improved ways. The seminar brought together researchers from a wide range of computing disciplines, such as mobile and ubiquitous computing, head-mounted eye tracking, optics, computer vision, human vision and perception, privacy and security, usability, as well as systems research. Attendees discussed how smart eyewear can change existing research and how it may open up new research opportunities. For example, future research in this area could fundamentally change our understanding of how people interact with the world around them, how to augment these interactions, and may have a transformational impact on all spheres of life – the workplace, family life, education, and psychological well-being.

6.9 Modern Cryptography and Security: An Inter-Community Dialogue

Organizers: Kristin Lauter, Radu Sion, and Nigel P. Smart

Seminar No. 16051

Date: January 31 to February 5, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.207

Participants: Raad Bahmani, Daniel J. Bernstein, Konstantin Beznosov, Alex Biryukov, Allison Bishop, Alexandra Boldyreva, Nikita Borisov, Ferdinand Brasser, Christian Cachin, Bogdan Carbunar, Melissa Chase, Jung Hee Cheon, Marc C. Dacier, George Danezis, Yevgeniy Dodis, Maria Dubovitskaya, Dieter Gollmann, Christian Grothoff, Krista Grothoff, Nadia Heninger, Aaron Michael Johnson, Stefan Katzenbeisser, Florian Kerschbaum, Yongdae Kim, Tanja Lange, Kristin Lauter, Yehuda Lindell, Sarah Meiklejohn, Refik Molva, Moni Naor, Claudio Orlandi, Kenneth G. Paterson, Adrian Perrig, Giuseppe Persiano, Andreas Peter, Benny Pinkas, Martina Angela Sasse, Vitaly Shmatikov, Radu Sion, Nigel P. Smart, Gene Tsudik, Avishai Wool

The seminar aimed to bring together communities with different backgrounds and form a bridge between them.

The outcomes ranged from a series of bridging exercises where participants summarized the current thoughts in existing areas; these included areas such as

- Hardware Attacks: Where we summarized the known attacks in this space.
- Computing on Encrypted Data: Various aspects of this were discussed, including Secure Guard Extensions (SGX), Searchable Symmetric Encryption (SSE), Multi Party Computation (MPC), and Fully Homomorphic Encryption (FHE).

We then went on to discuss more technical aspects, rather than just summarizing work,

- Cyberphysical Systems and IoT: Where the research challenges of performing work in this new area were discussed. A reliance on practical experimental was noted in the current research landscape.
- Mass Surveillance, Trapdoors, Secure Randomness: The recent "backdooring" of the DUAL–EC random number generator formed the background of this discussion. The seminar examined different aspects of this area, both in preventing, creating and detecting backdoors.
- Anonymous Payment Systems: This was a rather broad discussion which examined a number of issues around payments in general, and how cryptography could solve address these issues.

We also discussed aspects related to the process of research in this field. In particular focusing on the problem of the lack of expository writing. Here we identified a number of disincentives in the research culture which prevents the creation of more discursive writing and expository articles. A number of solutions both existing, and proposed, were discussed to solve this issue. In another small breakout we discussed the lack of incentives to work on the underlying hard problems upon which our security infrastructure rests.

In summary the seminar found more problems with our current research trends, than solutions.

6.10 Dark Silicon: From Embedded to HPC Systems

Organizers: Hans Michael Gerndt, Michael Glaß, Sri Parameswaran, and Barry L. Rountree

Seminar No. 16052

Date: January 31 to February 3, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.1.224

Participants: Axel Auweter, Sergio Bampi, Andrea Bartolini, Kirk W. Cameron, Pietro Cicotti, Isaias Alberto Compres Urena, Jonathan Eastep, Siddharth Garg, Hans Michael Gerndt, Michael Glaß, Per Gunnar Kjeldsberg, Michael Knobloch, Tulika Mitra, David Montoya, Wolfgang E. Nagel, Michael Niemier, Santiago Pagani, Sri Parameswaran, Tapasya Patki, Barry L. Rountree, Martin Schulz, Andrey Semin, Kathleen Shoga, Jürgen Teich

Topic

Dark Silicon

Semiconductor industry is hitting the utilization wall and puts focus on parallel and heterogeneous many-core architectures. While continuous technological scaling enables the high integration of 100s-1000s of cores and, thus, enormous processing capabilities, the resulting power consumption per area (the power density) increases in an unsustainable way. With this density, the problem of Dark Silicon will become prevalent in future technology nodes: It will be infeasible to operate all on-chip components at full performance at the same time due to the thermal constraints (peak temperature, spatial and temporal thermal gradients etc.).

Recent research work on power management for Dark Silicon aims at efficiently utilizing the TDP (Thermal Design Power) budget to maximize the performance or to allocate full power budget for boosting single-application performance by running a single core at the maximum voltage or multiple cores at nominal level for a very short time period. Control-based frameworks are proposed to find the optimal trade-off between power and performance of many-core systems under a given power budget. The controllers are coordinated to throttle down the power when the system exceeds the TDP and to assign the task to the most suitable core to get the optimal performance. The work on near-threshold computing (NTC) enables operating multiple cores at a voltage close to the threshold voltage. Though this approach favors applications with thread-level parallelism at low power, it severely suffers from errors or inefficiency due to process variations and voltage fluctuations. On the other hand, the computational sprinting approach leverages Dark Silicon to power-on many extra cores for a very short time period (100s of millisecond) to facilitate sub-second bursts of parallel computations through multi-threading but thereby wasting a significant amount of energy due to leakage current. When doing so, it consumes power that significantly exceeds the sustainable TDP budget. Therefore, these cores are subsequently power-gated after the computational sprint. Alternate methods are Intel's Turbo Boost and AMD's Turbo CORE technologies that leverage the temperature headroom to favor high-ILP applications by increasing the voltage/frequency of a core while power-gating other cores. These techniques violate the TDP constraint for a short period (typically in terms of 10s of seconds) until the critical temperature is reached and then switches to a nominal operation. However, in case of dependent workloads, boosting of one core may throttle the other due to thermal coupling (i.e. heat exchange between different cores sharing the same die). Therefore, these boosting techniques lack efficiency in case dependent tasks of an application mapped to two different cores or, in general, for multiple concurrently executing applications with distinctive/dependent workloads.

State-of-the-art boosting techniques assume a chip with only 10–20 cores (typically 16) and accordingly a full chip temperature violation for short time. However, in a large-scale system (with 100s–1000s cores), temperature hot spots may occur on certain chip portions far before the full chip's average temperature exceeds the critical temperature. Therefore, a chip may either get damaged before reaching the full chip critical temperature or TDP needs to be pessimistically designed. Advanced power management techniques are required to overcome these challenges in large-scale environments.

HPC – Dark Power

The energy consumption of HPC systems is steadily growing. The costs for energy in the five year lifetime of large scale supercomputers already almost equal the cost of the machine. It is a necessity to carefully tune systems, infrastructure and applications to reduce the overall energy consumption. In addition, the computing centers running very big systems face the problem of limited power provided by the energy providers and

of the requirement for an almost constant power draw from the grid. The big machines, especially future exascale systems, are able to use more power if they are run at highest performance of all components than can be provided by the energy company. Thus, a carefully optimized power distribution is necessary to make most efficient use of the provided power. The second aspect is the requirement of an almost constant power draw: Sudden changes from 20 MW to 10 MW for example, will be dangerous for the components of the power grid. In addition, the contracts with the energy companies force the centers use the same power all the time by charging more, if it drops below or exceeds certain limits. These challenges also require a careful and flexible power and resource management for HPC systems.

For a certain class of high-end supercomputer, there is a standard pattern of power consumption: During burn-in (and perhaps while getting a result to go onto the top-500 list) the machine will run dozens or hundreds of instances of Linpack. This code is quite simple and often hand-optimized, resulting in an unusually well-balanced execution that manages to keep vector units, cache lines and DRAM busy simultaneously. The percent of allocated power often reaches 95 % or greater, with one instance in recent memory exceeding 100 % and blowing circuit breakers. After these initial runs, however, the mission-critical simulation codes begin to execute and they rarely exceed 60 % of allocated power. The remaining 40 % of electrical capacity is dark: just as unused and just as inaccessible as dark silicon. While we would like to increase the power consumption (and thus performance) of these simulation codes, a more realistic solution in the exascale timeframe is hardware overprovisioning. This solution requires buying more compute resources than can be executed at maximum power draw simultaneously. For example, if most codes are expected to use 50 % of allocated power, the optimal cluster would have twice as many nodes.

Making this a feasible design requires management of power as a first-class resource at the level of the scheduler, the run-time system, and on individual nodes. Hardware power capping must be present. Given this, we can theoretically move power within and across jobs, using all allocated power to maximize throughput. The purpose of this seminar is to find this optimal level.

Hybrid (Design-time & Run-time) Resource Management

Today's complex applications need to exploit the available parallelism and heterogeneity of – non-darkened – cores to meet their functional and non-functional requirements and to gain performance improvements. From a resource management's point of view, modern many-core systems come with significant challenges: (a) Highly dynamic usage scenarios as already observable in today's "smart devices" result in a varying number of applications with different characteristics that are running concurrently at different points in time on the system. (b) Due to the constraints imposed by the power density, the frequency at which cores can be operated as well as their availability as a whole, are subject to change. Thus, resource management techniques are required that enable a resource assignment to applications that satisfies their requirements but at the same time can consider the challenging dynamics of modern many-cores as a result of Dark Silicon.

Traditional techniques to provide a binding or pinning of applications to processor that are optimal and predictable with respect to performance, timing, energy consumption, etc. are typically applied at design time and result in a kind of static system design. Such a static design may, on the one hand, be too optimistic by assuming that all assigned resources are always available or it may require for a kind of over-allocation of cores to compensate for worst-case scenarios, e.g., a frequent unavailability of cores due to Dark Silicon. Hence, the dynamic effects imposed in Dark Silicon require for novel modeling techniques already at design time.

Approaches that focus on pure run-time resource management are typically designed with flexibility in mind and should inherently be able to dynamically react to changing applications as well as to the described effects of Dark Silicon. But, future run-time resource management should not only react to a possible violation of a maximum power-density constraint, but also be able to proactively avoid such situations. The latter is an important aspect of the system's dependability as well. At the same time, such dynamic resource management is also required to regard the applications' requirements. Here, a careful consideration on whether pure run-time management strategies enable the amount of predictability of execution qualities required by some applications becomes necessary.

A recent research direction focuses on hybrid (design-time and run-time) approaches that explore this field of tension between a high predictability of design-time approaches and the dynamic adaptivity of run-time resource management. In such approaches, design-time analysis and optimization of the individual applications is carried out to capture information like core allocation, task binding, or message routing and predict resulting quality numbers like timeliness, energy consumption, or throughput. This information is then passed to the run-time resource management that then dynamically selects between the pre-optimized application embeddings. Such strategies may not only be able to achieve application requirements even in such highly dynamic scenarios, but could even balance the requirements of the individual applications with the system's requirements – in particular the maximum power density. On the other hand, coarse-grained resource management as required for core allocation etc. may be considered to happen on a longer time scale. The effects of Dark Silicon are instead on a smaller time scale with temperature almost immediately following changing workloads, thus, requiring for an intervention of the resource-management infrastructure. Therefore, novel concepts are required that enable a fine-grained resource management in the presence of Dark Silicon – both in the context of abstraction layer and time scale – without sacrificing the required efficiency but also predictable realization of application requirements via coarse-grained resource management.

Goals

Traditionally, resource management techniques play an important role in both domains – targeting very different systems. But, as outlined before, resource management may be the key to tackle the problem of dark silicon that both communities face. The aim of this seminar is to give an overview of the state of the art in the area of both embedded and HPC. It will make both groups aware of similarities and differences. Here, the competences, experiences, and existing solutions of both communities shall stimulate discussions and co-operations that hopefully manifest in innovative research directions for many-core resource management in the dark silicon era.

Overview of Contributions

This seminar presentations on the state-of-the-art in power and energy management in HPC and on techniques mitigating the Dark Silicon problem in embedded systems. In a joint session

commonalities and differences as well as collaboration potential in the area of Dark Silicon were explored. This subsection gives an overview of the topics covered by the individual speakers in the seminar. Please refer to the included abstracts to learn more about individual presentations.

The HPC-related presentations where started with an overview presentation by Barry Rountree from the Lawrence Livermore National Laboratory. He introduced the field of HPC and of exascale systems. The new challenge is that these systems will be power limited and the hardware is overprovisioned. Techniques increasing the efficient usage of the available power need to be developed. Exascale systems will be heterogeneous, even systems with homogeneous cores become heterogeneous due to production variability which takes effect under power limits. Careful distribution of power among jobs and within jobs as well as application and system configurations for jobs will be important techniques for these power limited and overprovisioned systems.

Axel Auweter added to this introduction deep insights into the electricity market in Germany, its complex price structure, and the challenges for German compute centers to act successfully on that market.

An introduction from the embedded field to Dark Silicon was given by Sri Parameswaran from the University of New South Wales. The continuous decrease in feature size without an appropriate decrease in the threshold voltage leads to increased power density. Between 50 % and 90 % of dark silicon is expected in future chips. Mitigation techniques are energy reduction techniques as well as spatial and temporal dimming of cores. Considerable energy reduction can be achieved from heterogeneity on various levels, e.g., heterogeneous cores and the DarkNoC approach.

Dark Silicon due to Power Density

Several techniques were presented to mitigate the effect of power density. Santiago Pagani presented *spatial and temporal dimming of cores* to make best use of the thermal distribution on the chip. He and Andrey Semin talked also about *boosting* the core frequency to exceed the power limit for a short time period to speedup computation. Sergio Bampi presented *near threshold computing* as a potential solution based on further lowering the threshold voltage. Michael Niemier explored the potential of *new transistor technology* to mitigate the Dark Silicon effect.

Dark Silicon due to Limited Power

Mitigation techniques in this field are quite similar in mobile computing and HPC, although the overall objective is a bit different. While in mobile computing the minimal power required to meet the QoS requirements of applications is the goal, in HPC it is to go as fast as possible with the available power, may be considering energy efficiency and system throughput as well.

The following approaches relevant for mobile computing and HPC were presented: *Heterogeneity* in various hardware aspects can be used to reduce the energy consumption of computations. Siddarth Garg and Tulika Mitra covered *performance heterogeneity* in scheduling tasks for big/little core combinations. Tulika Mitra and Andrea Bartolini talked about using *function heterogeneity*, e.g. accelerators, in mobile computing and HPC to increase energy efficiency. The *Heterogeneous Tile Architecture* was introduced in the presentations of Sri Parameswaran and Santiago Pagani as a general architecture enabling exploitation of heterogeneity to mitigate the Dark Silicon effect.

Another approach is to determine the most efficient *application and system configuration*. *Static tuning* of parameters, such as the power budget of an application, were presented by Michael Knobloch and Tapasya Patki. *Dynamic tuning* techniques were covered in the presentations of Michael Gerndt, Martin Schulz, and Per Gunnar Kjeldsberg. Jonathan Eastep introduced the GEO run-time infrastructure for distributed machine-learning based power and performance management.

Kirk Cameron highlighted the unexpected effects of changing the core frequency due to non-linear dependencies. Jürgen Teich talked about *Invasive Computing* providing dynamic resource management not only for improving certain non-functional application aspects but also for increasing the predictability of those aspects.

Wolfgang Nagel and Sri Parameswaran presented *energy efficient network architectures*. They covered heterogeneous on-chip network architectures and wireless communication within compute clusters.

Approximate computing was presented by Sergio Bampi. It allows trading off accuracy and energy. Pietro Cicotti covered in his presentation *data movement optimization* within a CPU to save energy.

Application and system monitoring is a pre-requisite for many of the above techniques. Michael Knobloch, Wolfgang Nagel, and Kathleen Shoga presented application and system monitoring techniques based on software as well as hardware instrumentation. Many compute centers are installing infrastructures to gather sensor values from the whole facility to enable future analysis. In addition to performance and energy measurements for application, higher level information about the application characteristics is useful in taking tuning decisions. Tapasay Patki presented *application workflows* as a mean to gather such information.

Besides these generally applicable techniques, some presentations covered also techniques that are specific to HPC installations with their batch processing approach and large compute systems.

Andrea Bartolini highlighted in his presentation the holistic multiscale aspect of power-limited HPC. The application, the compute system, and the *cooling infrastructure* have to be seen as a complex integrated system. *Power-aware scheduling*, presented by Tapasya Patki and Andrea Bartolini, can significantly improve the throughput of power-limit HPC systems and *moldable jobs* can improve the effect of power-aware scheduling significantly. Isaias Compres presented *Invasive MPI*, an extension of MPI for programming moldable application.

Conclusion

At the end of the seminar a list of takeaway messages was collected based on working-group discussions followed by an extensive discussion of all participants:

1. Dark silicon is a thermal problem in embedded and a power problem in HPC. HPC can cool down while in the embedded world you can't. Therefore HPC can power up everything if they have enough power. But the costs for providing enough power for rare use cases have to be rectified.
2. Better tools are required on both sides to understand and optimize applications.
3. Better support for optimizations is required through the whole stack from high level languages down to the hardware.
4. In both communities run-time systems will get more important. Applications will have to be written in a way that run-time systems can work effectively.
5. Task migration is of interest to both groups in combination with appropriate run-time management techniques.
6. Embedded also looks at specialized hardware designs while HPC has to use COTS. In HPC, the machine architecture might be tailored towards the application areas. Centers are specialized for certain customers.

7. Heterogeneity on architecture level is important to both groups for energy reduction.
8. Better analyzable programming models are required, providing composable performance models.
9. HPC will have to live with variability. The whole tuning step has to change since reproducibility will no longer be given.
10. Hardware-software co-design will get more important for both groups.
11. Both areas will see accelerator-rich architectures. Some silicon has to be switched off anyway, thus these can be accelerators that might not be useful for the current applications.

6.11 Data-Driven Storytelling

Organizers: Sheelagh Carpendale, Nicholas Diakopoulos, Nathalie Henry Riche, and Christophe Hurter

Seminar No. 16061

Date: February 7–12, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.1

Participants: Fereshteh Amini, Benjamin Bach, Lyn Bartram, Dominikus Baur, Gordon Bolduan, Jeremy Boy, Matthew Brehmer, Sheelagh Carpendale, Fanny Chevalier, Paolo Ciuccarelli, Nicholas Diakopoulos, Marian Dörk, Steven M. Drucker, Tim Dwyer, Jason Dykes, Kennedy Elliott, Christina Elmer, Yuri Engelhardt, Xaquin Gonzalez Veira, Theresia Gschwandtner, Nathalie Henry Riche, Jessica Hullman, Samuel Huron, Christophe Hurter, Ulrike Köppen, Robert Kosara, Bongshin Lee, Giuseppe Santucci, Jonathan Schwabish, John T. Stasko, Moritz Stefaner, Alice Thudt, Melanie Tory, Barbara Tversky, Jarke J. van Wijk, Jagoda Walny, Stefan Wehrmeyer, Benjamin Wiederkehr, Jo Wood

Data visualization is the "use of computer-supported, interactive, visual representations of data to amplify cognition" [5]. Visualization can play a crucial role for exploring data and for communicating information as "a picture is worth a thousand words". Early research in this field focused on producing static images and quantifying the perception of different visual encodings [6] in these visual representations. The vast majority of research since then focused on designing and implementing novel interfaces and interactive techniques to enable data exploration. Major advances in visual analytics and big data initiatives concentrated on integrating machine learning and analysis methods with visual representations to enable powerful exploratory analysis and data mining [10]. As interactive visualizations play an increasing role in data analysis scenarios, they also started to appear as a powerful vector for communicating information. Stories supported by facts extracted from data analysis proliferate in many different forms from animated infographics and videos [2] to interactive online visualizations on news media outlets. We argue that it is now time for the visualization research community to understand how these powerful interactive visualizations play a role in communicating information. We define this line of research as data-driven storytelling.

The popularity of javascript web technology and the availability of the D3 toolkit [3] enabled a wider range of people to create data visualizations. Being able to easily share interactive data visualizations on the web also increased the democratization of interactive visualizations. Coupled with the emphasis on data science, these advances raise new practices such as data journalism. Data journalists gather and explore available datasets to extract relevant insights, often conveying their stories via interactive data visualizations [1,9]. The popularity of data-driven stories on New York Times especially, revealed the potential of interactive visualizations as a powerful communication tool [7].

Central to our vision of the convening was that the vast majority of research on data visualization to date has focused on designing and implementing novel interfaces and interactive techniques to enable data exploration. Major advances in visual analytics and big data initiatives have concentrated on integrating machine learning and analysis methods with visual representations to enable powerful exploratory analysis and data mining. But just as interactive visualization plays an important role in data analysis scenarios it is also becoming increasingly important in structuring the communication and conveyance of insights and stories in a compelling format. Visual data-driven stories have proliferated in many different forms, from talks [8], to animated infographics and videos [1, 7, 9], to interactive online visualizations.

Data-driven storytelling is also compelling for a wide range of applications. In enterprise scenarios, the output of data analysis (often reports and slide-based presentations) has to be conveyed to decision makers. In scientific research, interactive visualizations are increasingly used to convey data-driven discoveries to peers or used to communicate complex findings to a broader audience. In education scenarios, interactive visualizations are used by teachers to explain mathematical concepts or to illustrate biological or physical mechanisms. Many questions arise as interactive visualizations are used beyond data exploration by experts, for communication purposes to a broader audience. Research on understanding of static images in cognitive psychology and perception must be extended to encompass more advanced techniques (videos and interactive applications). Visualization literacy, defined as the ability to extract, interpret, and make meaning from information presented in the form of an (interactive) data visualization is also a crucial component for data-driven storytelling research. Assessing the visualization literacy of an audience and developing techniques to better teach how to decode interactive visualizations has started to attract the attention of our research community [4] However a plethora of research remains to be done. For example,

research on how visualizations can lie [11] or at least how they may introduce bias in the reader's mind has focused on static visual representations but has not yet been extended to other medium. Similarly it is crucial for advancing researches in visualization to assess the role data-driven storytelling can play in easing the comprehension of a messages or in increasing their memorability.

The visualization research community needs to reflect on data-driven storytelling and to develop a research agenda to investigate how advanced data-driven stories are understood by the audience, identify factors that makes them compelling as well as factors that can introduce bias in their perception. By learning from master storytellers from other fields (journalism, design, art and education) strategies to craft successful stories, our community will be able to reflect on these questions and eventually build novel consuming tools that engage a broad audience while minimizing perception bias, as well as build novel authoring tools to craft high quality data-driven stories.

One domain where there has been extensive and practical progress on the question of data-driven storytelling is data journalism. News sites like FiveThirtyEight or the New York Times' The Upshot have seen a recent surge of attention and interest as a means of communicating data-driven news to the public. By carefully structuring the information and integrating explanation to guide the consumer, journalists help lead users toward a valid interpretation of the underlying data. Because of the rapid and practical progress of data-driven storytelling in the domain of journalism, our seminar sought to put some of the top practitioners from that field together with computer science researchers to discuss the challenges and opportunities of data-driven communication.

The Dagstuhl seminar was structured to leverage the interdisciplinarity of the attendees by first tapping into a divergent design thinking process meant to enumerate the range of issues that are relevant to data-driven stories. Hundreds of index cards and sticky notes were sacrificed as participants generated ideas (see Figure 6.4).

We then clustered these ideas to arrived at a set of key themes, including:

- Techniques and Design Choices for Storytelling
- Exploration and Explanation
- From Analysis to Communication
- Audience
- Evaluation
- Devices and Gadgets
- Ethics

Groups of participants formed around common interests and each of these major themes were then the focus of discussion. Each work group was geared towards developing an outline and plan to produce a written chapter for a forthcoming edited book on the topic of data-driven storytelling. Some groups met for a day or two and then reformed around other topics, whereas other groups spent the entire week going deep in exploring a single topic. And as if the daytime activities weren't enough, additional evening breakout groups formed around additional topics of interest like Education in Data Visualization, Urban Visualization, and the Technology Stack for data-driven stories.

In-between the intense, small group sessions the entire group came together daily for five-minute lightning talks on a wide array of relevant topics. These stimulating talks primed the group for approaching data-driven storytelling from different perspectives and were an entertaining and informative way to share creative ideas or results in small and easily digestible nuggets. Among the more than 25 lightning talks, topics ranged from storytelling with timelines, to mobile visualization, the use of data comics, visual literacy, affect and color, data-story design workflows, and even the visualization of data through cuisine.

Outcomes

Our initial goal of the seminar was to have groups work intensively on their chosen topic(s) so that an outline and work-plan could be developed to write a contributing chapter to a book on data-driven storytelling. The book is underway and will have contributions on each of the main themes outlined above, as well as an introductory chapter by the editors / organizers of the Dagstuhl seminar. Moreover, our creative contributors at the seminar produced other outputs as well: curated lists of example data driven stories, as well as of storytelling techniques were created and will be published online, and a blog has pulled together some of the formative impressions of participants (https://medium.com/data-driven-storytelling).

Below we briefly summarize the expected contents of each of the chapters that will form the book.

Techniques and Design Choices for Storytelling. This chapter will discuss techniques and design choices for visual storytelling grounded in a survey of over 60 examples collected from various online news sources and from award-winning visualization and infographic design work. These design choices represent a middle ground between low-level visualization and interaction techniques and high-level narrative devices or structures. The chapter will define several classes of design choices: embellishment, explanation, exploration, navigation, story presentation, emphasis, focus, and annotation. Examples from the survey for each class of design choices will be provided. Finally, several case studies of examples from the survey that make use of multiple design choices will be developed.

Exploration and Explanation in Data-Driven Stories. This chapter will explore the differences between and integration of exploration and explanation in visual data-driven storytelling. Exploratory visualizations allow for a lot of freedom which can include changing the visual representation, the focus of what is being shown and the sequence in which the data is viewed. They allow readers to find their own stories in the data. Explanatory stories include a focused message which is usually more narrow and guides the reader often in a linear way. Advantages and disadvantages of exploration and explanation as well as dimensions that help to describe and classify data-driven stories will be developed. The space is described by identifying freedom, guidance regarding representation, focus and sequence as well as interpretation as important dimensions of data-driven storytelling and existing systems are characterized along these dimensions. Recommendations will be developed for how to integrate both aspects of exploration and explanation in data-driven stories.

From Analysis to Communication: Supporting the Lifecycle of a Story. This chapter will explore how tools can better support the authoring of rich and custom data stories with natural / seamless workflows. The aim is to understand the roles and limitations of analysis / authoring tools within current workflow practices and use these insights to suggest opportunities for future research and design. First, the chapter will report a summary of interviews with practitioners at the Dagstuhl seminar; these interviews aim to understand current workflow practices for analysis and authoring, the tools used to support those practices, and pain points in those processes. Then the chapter will reflect on design implications that may improve tool support for the authoring process as well as research opportunities related to such tool support. A strong theme is the interplay between analytical

and communicative phases during both creation and consumption of data-driven stories.

The Audience for Data-Driven Stories. Creators of data-driven visual stories want to be as effective as possible in communicating their message. By carefully considering the needs of their audience, content creators can help their readers better understand their content. This chapter will describe four separate characteristics of audience that creators should consider: expertise and familiarity with the topic, the medium, data, and data visualization; expectations about how and what the story will deliver; how the reader uses the interface such as reading, scrolling, or other interactivity; and demographic characteristics of the audience such as age, gender, education, and location. This chapter will discuss how these audience goals match the goals of the creator, be it to inform, persuade, educate, or entertain. Then it will discuss certain risks creators should recognize, such as confusing or offending the reader, or using unfamiliar jargon or technological interfaces. Case studies from a variety of fields including research, media, and government organizations will be presented.

Evaluating Data-Driven Storytelling. The study of data-driven storytelling requires specific guidelines, metrics, and methodologies reflecting their different complex aspects. Evaluation is not only essential for researchers to learn about the quality of data-driven storytelling but also for editorial rooms in media and enterprises to justify the required resources the gathering, analyzing and presentation of data. A framework will be presented that takes the different perspectives of author, audience and publisher and their correspondent criteria into account. Furthermore it connects them with the methods and metrics to provide a roadmap for what and how to measure if these resulting data-driven stories met the goals. In addition, the chapter will explore and define the constraints which might limit the metrics and methods available making it difficult to reach the goals.

Devices and Gadgets for Data Storytelling. This chapter will discuss the role of different hardware devices and media in visual data driven storytelling. The different form factors offer different affordances for data storytelling affecting their suitability to the different data storytelling settings. For example, wall displays are well suited to synchronous co-located presentation, while watches and virtual reality headsets work better for personal consumption of pre-authored data stories.

Ethics in Data-Driven Visual Storytelling. Is the sample representative, have we thought of the bias of whoever collected or aggregated the data, can we extract a certain conclusion from the dataset, is it implying something the data doesn't cover, does the visual device, or the interaction, or the animation affect the interpretation that the audience can have of the story? Those are questions that anyone that has produced or edited a data-driven visual story has, or at least should have, been confronted with. After introducing the space, and the reasons and implications of ethics in this space, this chapter will look at the risks, caveats, and considerations at every step of the process, from the collection/acquisition of the data, to the analysis, presentation, and publication. Each point will be supported by an example of a successful or flawed ethical consideration.

Conclusion

The main objective of this Dagstuhl seminar was to develop an interdisciplinary research agenda around data-driven storytelling as we seek to develop generalizable findings and tools to support the use of visualization in communicating information. Productive group work converged to delineate several research opportunities moving forward:

- The need for interfaces that enable the fluid movement between exploratory and communicative visualization so that storytelling workflow is seamless and powerful.
- The need to develop typologies of visual storytelling techniques and structures used in practice so that opportunities for supporting these techniques can be sought through computing approaches.
- The need to develop evaluation frameworks that can assess storytelling techniques and tools both scientifically and critically.
- The need for design frameworks that can guide the structure of visual information for experiences across different output devices, both existing and future.
- The need to understand the audience and their role in co-constructing meaning with the author of a data-driven story.
- The need for ethical frameworks that should guide tool development for visual data-driven communication.

These opportunities were productively enumerated at the Dagstuhl seminar and are in the process of being written up as chapters in our book on data-driven storytelling.

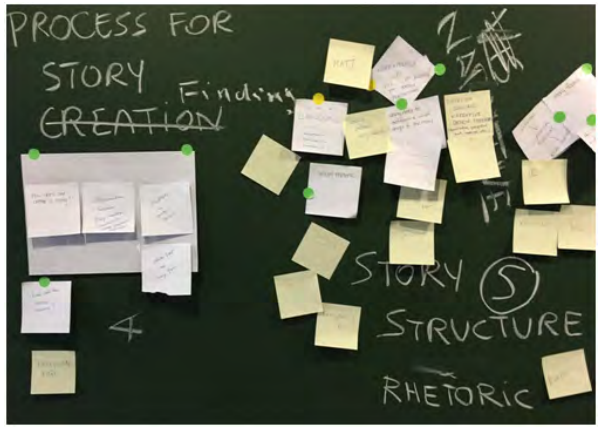

Fig. 6.4
Converging on topical groups from hundreds of individual ideas.

References

1 All the medalists: Men's 100-meter freestyle: Racing against history. Web, 2012. http://www.nytimes.com/interactive/2012/08/01/sports/olympics/racing-against-history.html?_r=1&.

2 Fereshteh Amini, Nathalie Henry Riche, Bongshin Lee, Christophe Hurter, and Pourang Irani. Understanding data videos: Looking at narrative visualization through the cinematography lens. In *Proceedings of CHI: SIGCHI Conference on Human Factors in Computing System*, pages 1459–1468, New York, NY, USA, 2015. ACM.

3 Michael Bostock, Vadim Ogievetsky, and Jeffrey Heer. D3 data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, December 2011.

4 Jeremy Boy, Sara Johansson Fernstad, Martin Turner, Simon Walton, David Ebert, Jean-Daniel Fekete, Andy Kirk, and Mario Romeo. Eurovis 2014 workshop: Towards visualization literacy. visualization literacy workshop. EuroVis 2014, 2014. https://www.kth.se/profile/178785/page/eurovis-2014-workshop-towards-visualiza/.

5 Stuart K. Card, Jock D. Mackinlay, and Ben Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers Inc. San Francisco, CA, USA., 1999.

6 William C. Cleveland and Marylyn E. McGill. *Dynamic Graphics for Statistics (1st ed.)*. CRC Press, Inc., Boca Raton, FL, USA., 1999.

7 Instance of data storytelling to explain the subprime mortgage crisis in the usa. Web, 2014. http://www.bloomberg.com/dataview/2014-02-25/bubble-to-bust-to-recovery.html.

8 Hans Rosling. The best stats you've ever seen. TED Talks, 2006. https://youtu.be/hVimVzgtD6w.

9 The wealth report 2013: Examining high net worth individuals from around the world. Web, 2013. http://www.elitehavenssales.com/news/47/The-Wealth-Report-2013-Examining-High-Net-Worth-Individuals-from-around-the-world.

10 J. J. Thomas and K. A. Cook. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE CS Press, USA., 2005.

11 Edward Tufte. *The Visual Display of Quantitative Information, Second Edition*. Graphics Press, USA., 1991.

6.12 Modeling and Analysis of Semiconductor Supply Chains

Organizers: Chen-Fu Chien, Hans Ehm, John Fowler, and Lars Mönch

Seminar No. 16062

Date: February 7–12, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.28

Participants: Chen-Fu Chien, Stéphane Dauzère-Pérès, Ton de Kok, Hans Ehm, Kenneth Fordyce, José M. Framinán, Cathal Heavey, Raphael Herding, Jesus Jimenez, Adar Kalir, Sebastian Knopp, Peng-Chieh Lee, Peter Lendermann, Iris Lorscheid, Scott J. Mason, Leon F. McGinnis, Hubert Missbauer, Lars Mönch, Irfan Ovacik, Thomas Ponsignon, Oliver Rose, Can Sun, Israel Tirkel, Reha Uzsoy, Gerald Weigert, Jei-Zheng Wu

Complex manufacturing processes are the heart of semiconductor manufacturing. A semiconductor chip is a highly miniaturized, integrated circuit (IC) consisting of thousands of components. Semiconductor manufacturing starts with thin discs, called wafers, (typically) made of silicon. A large number of usually identical chips can be produced on each wafer by fabricating the ICs layer by layer in a wafer fabrication facility (wafer fab). The corresponding step is referred to as the Fab step. Next, electrical tests that identify the individual dies that are likely to fail when packaged are performed in the Probe facility. An electronic map of the condition of each die is made so that only the good ones will be used. The probed wafers are then sent to an Assembly facility where the good dies are put into an appropriate package. The assembled dies are sent to a test facility where they are tested to ensure that only good products are sent to customers. The tested devices are then sent to regional warehouses or directly to customers. Wafer fabrication and probe are often called the front-end and assembly and test are called the back-end.

Supply chain management (SCM) problems have become more and more important in the last decade. This has been caused by the fact that front-end operations are often performed in highly industrialized nations, while back-end operations are typically carried out in countries where labor rates are cheaper. Moreover, there are centers of competencies (e.g. bumping) that may consist of only a few process steps that may be done in a different company owned facility or remotely by a subcontractor. These centers of competencies speed up innovations and reduce costs, but increase the complexity of SCM.

The semiconductor industry is capital intensive with the cost of an entire wafer fab up to nearly $10 billion US caused primarily by extremely expensive machines, some up to $100 million US each. The manufacturing process is very complex due to the reentrant flows in combination with very long cycle times and the multiple sources of uncertainty involved. Capacity expansions are very expensive and time-consuming. This kind of decision is based on demand forecasts for the next years. Because of the rapidly changing environment, the demand is highly volatile. Consequently, the forecast is rarely accurate. The semiconductor industry is an extreme field for SCM solutions from an algorithmic as well as from a software and information systems point of view. The huge size of the supply chains involved, the pervasive presence of different kinds of uncertainties and the rapid pace of change leads to an environment that places approaches developed in other industries under major stress. Modeling and analysis approaches that are successful in this industry are likely to find applications in other areas, and to significantly advance the state of the art in their fields (cf. [1]).

The purpose of this seminar was to bring together researchers from different disciplines including information systems, computer science, industrial engineering, operations research, and supply chain management whose central interest is in modeling, analyzing, and designing complex and large-scale supply chains as in the semiconductor industry. Moreover, practitioners from the semiconductor industry who have frequently articulated their perception that academic research does not always address the real problems faced by the industry brought in their domain knowledge to make sure that progress towards applicability and feasibility would be made during this seminar. The seminar had 26 attendees from ten different countries (see participant list at the end of the report). We had participants from leading semiconductor companies Infineon Technologies and Intel Corp. as well as researchers who work closely with ST Microelectronics, Globalfoundries, and Taiwan Semiconductor Manufacturing Company (TSMC).

A primary purpose of the workshop was to extend the scope of the academic research community from single wafer fabs to the entire semiconductor supply chain. We show the principle architecture of the planning and control system of a semiconductor supply chain in Figure 6.5.

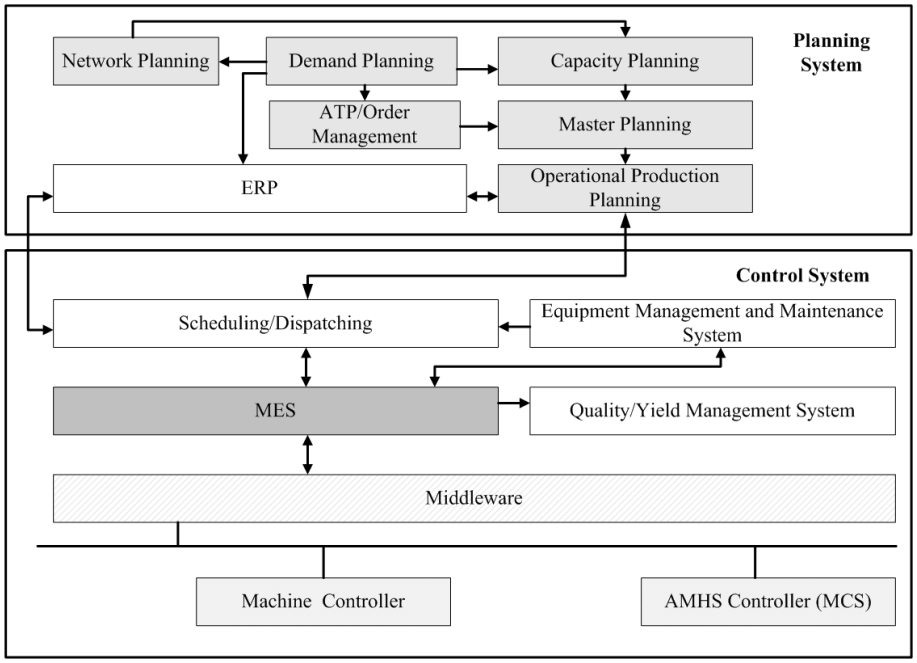

Fig. 6.5
Planning and Control System of a Semiconductor Supply Chain (adapted from [2]).

Seminar Objectives

The first objective of the seminar consisted of developing a research agenda for semiconductor supply chain modeling and analysis topics. This includes innovative modeling approaches for supply chain network planning, demand planning, master planning, and detailed production planning and scheduling in semiconductor supply chains. But it also includes ideas on how to design the related future information systems.

The research agenda was developed around the following two main topics:

- Topic 1: Novel planning and scheduling approaches that can deal with the complexity and stochasticity of the semiconductor supply chain:
  - Many planning approaches on the SC-level are based on (distributed) hierarchical and generally deterministic approaches to deal with the sheer complexity of the semiconductor supply chain. The role of anticipation of lower level behavior in upper level decision-making is still not well understood and has to be studied in more detail. Because a semiconductor supply chain contains many different, often autonomous decision-making entities including humans, negotiation approaches are typical in such distributed hierarchical systems for planning and control. It should be researched how such negotiation approaches can be automated and which decisions should be made by humans.
  - The overall cycle times in a typical semiconductor supply chain are on the order of 10 to 15 weeks. Therefore lead times have to be modeled in planning formulations. Using lead times as exogenous parameters in planning formulations leads to a well-known circularity because the cycle time depends in a nonlinear manner on the resource utilization which is a result of the release decisions made by the planning approach. Different types of clearing functions have to be researched in the semiconductor supply chain context.
  - Approaches to demand planning that take the product life cycle into account have to be studied. The interaction of demand planning and supply chain planning has to be investigated.
  - Different ways to anticipate stochasticity including robust optimization, approximate dynamic programming, and stochastic programming have to be researched in the semiconductor supply chain context.
  - Different ways to appropriately deal with stochasticity including rolling planning techniques and inventory holding strategies have to be studied.
  - Generation of scenarios and other distribution parameters for planning problems in supply chains using data mining techniques have to be researched.
  - Because of the complexity of supply chains, long computing times still hinder the usage of analytic solution approaches especially for what-if analysis. The role of state-of-the-art computing techniques including parallel computing on Graphics Processing Units (GPU) machines or Cloud computing techniques in decision-making for semiconductor supply chains has to be investigated.
- Topic 2: Future information systems and supply chain management in the semiconductor industry:
  - Understanding the limitations of today's packaged software for supply chain management in the semiconductor industry.
  - Proposing alternative software solutions including software agents and service-oriented computing for planning and scheduling applications in the supply chain context.
  - Integration concepts for state-of-the-art computing techniques to get models that are computationally tractable and address the different uncertainties encountered in this industry.
  - Approaches to embed real time simulation techniques in current and future information systems to support decision-making in semiconductor supply chains.
  - Understanding the interaction of human agents with information systems.

The implementation of ERP, APS, and MES systems in semiconductor supply chains provides both an opportunity and the need for development of supply-chain wide integrated production planning and scheduling solutions. Therefore, we think that the second topic is important and should be also addressed in the research agenda. Research related only to the first main topic is not sufficient.

The second objective of the seminar consisted of identifying the core elements of a conceptual reference model for planning and control of a supply chain in the semiconductor industry that can be used for analysis and performance assessment purposes and to foster a common understanding in the research community both in academia and industry. This included specifying reference planning and control activities, the major information flows, and their interaction with a reference system of a physical supply chain. Due to the inherent complexity of semiconductor supply chains it requires simulation of the physical supply chain to understand the interactions between the planning and control components and the physical supply chain, to find solution approaches to problems and to verify them in the risk-free simulation environment before implementing them. There are widely accepted reference (simulation) models for single wafer fabs, mainly developed in the Measurement and Improvement of Manufacturing Capacity (MIMAC) project (led by one of the organizers of this Dagstuhl seminar) 20 years ago that are still used by many academic researchers working with the semiconductor industry.

Existing reference models on the planning and control level like the Supply Chain Operations (SCOR) reference model and the supply chain planning (SCP) matrix are too generic to be useful for detailed analysis and have to be refined considerably to cover the important domain-specific aspects of semiconductor supply chains.

The Process

In the opening session, the organizers welcomed the participants and acknowledged Infineon Technologies as a sponsor of the seminar. Next, the participants each introduced themselves. This was followed by an overview of the goals and objectives of the seminar and a detailed review of the seminar program including the ground rules for interactions.

The remainder of the day on Monday consisted of four industry overview talks (by Hans Ehm, Kenneth Fordyce, Chen-Fu Chien, and Irfan Ovacik) and a review of the literature related to modeling an analysis of semiconductor supply chains (by Lars Mönch and Reha Uzsoy). Tuesday and half a day on Wednesday were devoted to presentations and discussions about the various elements of the semiconductor supply chain planning and control systems shown in Figure 6.5 above.

Wednesday afternoon was the excursion that was enjoyed by the participants. Thursday was devoted to 3 breakout sessions.

The first set of breakout sessions had four groups focus on the individual elements in Figure 6.5 and one group focus on a semiconductor supply chain reference model. The second set of breakouts had three groups consider the interaction between various elements in Figure 6.5, one group talked about the incorporation of humans in the supply chain, and one discussed how to go from the reference model to a specific semiconductor supply chain model instance.

The final Friday set of breakouts included three groups that discussed process models of multiple elements from Figure 6.5 and the flow of information needed between the elements to provide core elements of a reference model. Another group discussed the role of agents in a semiconductor company's supply chain. The final breakout group discussed the level of detail needed in a top down reference model. Friday consisted of a discussion on the required core elements of a reference model for semiconductor supply chains and a wrap-up session.

Key Take Aways

There were a number of key findings and areas for future research that were identified in the seminar. We will first summarize some of the key findings and will follow this with some areas for future research.

One of the first findings was that the participants generally agreed that the different elements in Figure 6.5 are reasonably well understood by both the industrial and academic communities, but the interactions between the elements are less well understood. Having said this, a number of the software solutions for the elements are not geared toward the complexities of the semiconductor industry (e.g. ATP/APS systems are generally focused on profit maximization and ignore many of the system complexities). Second, it appears that there are still limitations in solution approaches in practice such as: capacity generally is expressed without regard to mix; fixed lead times are generally still assumed despite research done on clearing functions for planning; and ignoring all but production lots when developing plans. Third, as indicated above both the industrial and academic participants generally agree that the integration of the decisions made by the different elements is often fairly ad hoc and could/should be improved. Finally, the participants generally agreed that there does not currently exist an adequate reference model for the semiconductor supply chain. In fact, there is not even a reasonable set of data sets that describe instances of the semiconductor supply chain such as the MIMAC datasets at the factory level. There is some indication that a reference model and incorporating human behavior of the various decision makers on the supply chain level will help to better understand supply chains producing and containing semiconductors.

In addition to the findings mentioned above, several areas for future research were identified. An overarching idea was that the future research should focus more on formulation of appropriate models because this is fundamentally more important than the actual solution techniques chosen. Some of the future research areas are included below:

- Using event-driven process chains (EPCs) to model/visualize planning processes.
- Developing better integration of various decisions made in the elements of Figure 6.5.
- Combining rolling horizon strategies with demand forecast evolution models.
- Incorporating sustainability aspects into supply chain models.
- Developing stochastic model versions of current deterministic models.
- Incorporating the behavior of human decision makers (this will be useful, but challenging).
- Exploring the use of different simulation paradigms (systems dynamics, agent-based, hybrid models, reduced simulation models) to model and analyze semiconductor supply chains.

Next Steps

As a way to further the discussion of and collaboration on the topics of the seminar, Prof. Lars Mönch, Prof. Chen-Fu Chien, Prof. Stéphane Dauzère-Pérès, Hans Ehm, and Prof. John Fowler are guest editing a special issue of the *International Journal of Production Research* (IJPR) entitled *Modeling and Analysis of*

Semiconductor Supply Chains. The deadline for submission is September 1, 2016. This date was selected to allow time for ideas created by the participants of the seminar to be incorporated into papers for the special issue. The Call for Papers can be found at the following address:

http://explore.tandfonline.com/cfp/est/semiconductor-supply-chains-call

Acknowledgements. The seminar organizers would like to thank Infineon Technologies AG for their support of the seminar. The seminar also would not have been nearly as productive without the active contribution of every attendee, and for that the organizers are extremely grateful.

References

1 Chien, C.-F., Dauzère-Pérès, S., Ehm, H., Fowler, J. W., Jiang, Z., Krishnaswamy, S., Mönch, L., Uzsoy, R. (2011): Modeling and Analysis of Semiconductor Manufacturing in a Shrinking World: Challenges and Successes. *European Journal of Industrial Engineering*, 5(3), 254–271, 2011. http://dx.doi.org/10.1504/EJIE.2011.041616

2 Mönch, L., Fowler, J. W., Mason, S. J. (2013): *Production Planning and Control for Semiconductor Wafer Fabrication Facilities: Modeling, Analysis, and Systems*. Springer Operations Research/Computer Science Interfaces, New York, Vol. 52, 2013. ISBN 978-1-4899-9901-6.

6.13 Pattern Avoidance and Genome Sorting

Organizers: Michael Albert, Miklós Bóna, István Miklós, and Einar Steingrimsson

Seminar No. 16071

Date: February 14–19, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.65

Participants: Michael Albert, David Bevan, Miklós Bóna, Mathilde Bouvel, Marilia Braga, Brona Brejova, Robert Brignall, Cedric Chauve, Anders Claesson, Péter L. Erdös, Niklas Eriksen, Pedro Feijão, Guillaume Fertin, Sylvie Hamel, Vít Jelínek, Anthony Labarre, Marie-Louise Lackner, Martin Lackner, Megan Martinez, István Miklós, Jay Pantone, Adeline Pierrot, Yann Ponty, Svetlana Poznanovikj, Manda Riehl, Bruce Sagan, David Sankoff, Rebecca Smith, Einar Steingrimsson, Jens Stoye, Krister Swenson, Eric Tannier, Vincent Vatter, Stéphane Vialette, Tomáš Vinař

The seminar took place from February 14, 2016, to February 19, 2016. It had 36 participants, who were researchers in theoretical computer science, combinatorics, and molecular biology. It was a geographically diverse group, with participants coming from the US, Canada, Brazil, Germany, Iceland, the United Kingdom, Sweden, France, Slovakia, Hungary and New Zealand. The seminar featured 18 talks, three of which were hourlong talks, and an open problem session.

Numerous collaborative research efforts have been started. Here is a sampling.

Megan Martinez and Manda Riehl worked on a bijection between LP matchings (one of the RNA matchings described in Vincent Vatter's talk) and Klazar's nesting equivalent matchings. They studied a paper by Klazar and Aziza Jefferson's dissertation and made progress on the bijection.

István Miklós, Péter Erdős and Miklós Bóna worked on proving a log-convexity conjecture related to ordered degree sequences of bipartite graphs.

Brona Brejova and Manda Riehl discussed two potential future projects related to gene and species tree reconciliation. The most probable starting point is a project involving gene and species trees where a gene is allowed to duplicate a string inside itself. This situation was not allowed in previous models, however it seems that as long as the specific breakpoints are not reused from this insertion, a modification of the previous algorithms could still be effective.

Jay Pantone, David Bevan and Miklós Bóna collaborated on asymptotic enumeration of a balanced urns and balls model that was seen to be a step towards finding a better upper bound for a pattern avoidance enumeration problem.

We have all the reasons to believe that this, and many other joint research efforts that started during this seminar will lead to new results that would not have been possible without the seminar. Therefore, we strongly believe that the seminar was a success that we would like to repeat at some point in the future.

6

6.14 Assessing Learning In Introductory Computer Science

Organizers: Michael E. Caspersen, Kathi Fisler, and Jan Vahrenhold

Seminar No. 16072

Date: February 14–19, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.78

Participants: Michael E. Caspersen, Holger Danielsiek, Brian Dorn, Katrina Falkner, Sally Fincher, Kathi Fisler, Mark Guzdial, Geoffrey L. Herman, Lisa C. Kaczmarczyk, Andrew J. Ko, Michael Kölling, Shriram Krishnamurthi, Raymond Lister, Briana Morrison, Jan Erik Moström, Andreas Mühling, Anthony Robins, Rolf Schulmeister, Carsten Schulte, R. Benjamin Shapiro, Beth Simon, Juha Sorva, Martijn Stegeman, Heike Theyssen, Jan Vahrenhold, Mirko Westermeier, Steven A. Wolfman

The goal of the seminar was to focus on several broadly applicable learning outcomes for first year university computer science courses, looking at what it would take to understand and assess them in multiple pedagogic contexts.

In preparation for the seminar, we surveyed participants to get an understanding of a what could be a common denominator of CS1/2 learning outcomes, using the outcomes from the ACM CC 2013 curriculum as a starting point. We asked participants to (a) identify ones that are covered in their institution's CS1/2 courses, and (b) to identify ones that they have either experience or interest in investigating further. Participants also suggested objectives that were not included in CC 2013.

Of these candidate outcomes, we studied a subset during the seminar, as voted by the participants. We used breakout sessions to get small groups of participants to focus on individual outcomes, reporting on what is known about each outcome, its underlying challenges and/or relevant underlying theory, how to best assess it, and what sorts of research questions should be asked to advance educational research on that outcome. We had three separate sets of breakout sessions, so each participant had the chance to work on three outcomes in detail during the week. The discussion of some sessions was continued in a following session.

Rather than have most individual participants give talks, we ran three speed-dating poster sessions on the first afternoon: each person got to put up a poster on some outcome that they have studied, so others could see the research of other attendees.

In addition, we had three invited presentations focussing on workload and determinants of study success (Schulmeister), types of prior knowledge and their relation to study success (Theyssen), and Concept Inventories (Kaczmarczyk and Wolfman). The abstracts of these presentations are included in this report.

6.15 Scheduling

Organizers: Nikhil Bansal, Nicole Megow, and Clifford Stein

Seminar No. 16081

Date: February 21–26, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.97

Participants: Fidaa Abed, Susanne Albers, Antonios Antoniadis, Yossi Azar, Nikhil Bansal, Sanjoy K. Baruah, Vincenzo Bonifaci, Niv Buchbinder, Marek Chrobak, Bouke Cloostermans, Liliana Cucu-Grosjean, Robert Davis, Christoph Dürr, Thomas Erlebach, Anupam Gupta, Magnus M. Halldórsson, Sungjin Im, Klaus Jansen, Christos Kalaitzis, Samir Khuller, Amit Kumar, Retsef Levi, Alberto Marchetti-Spaccamela, Monaldo Mastrolilli, Nicole Megow, Rolf H. Möhring, Benjamin J. Moseley, Seffi Naor, Kirk Pruhs, Thomas Rothvoss, Jiri Sgall, Hadas Shachnai, David Shmoys, René Sitters, Frits C. R. Spieksma, Clifford Stein, Ola Svensson, Marc Uetz, Suzanne van der Ster, Rob van Stee, Jose Verschae, Tjark Vredeveld, Andreas Wiese, Gerhard J. Woeginger, Prudence W. H. Wong

This fourth meeting in a series of Dagstuhl "Scheduling" seminars had two major objectives. Firstly, it offered a forum for presenting recent scheduling results of high impact and new techniques which may be useful for solving important and long-standing open problems. The second major objective was to debate and explore future research directions, discuss important open problems, and foster new collaborations with a particular attention to interactions with application areas, both in academia and industry.

The organization of the meeting differed from the previous Dagstuhl "Scheduling" seminars by not inviting a different community to interact. Despite (or perhaps because of) the success of the cross-discipline events, there was an explicit desire to dedicate a seminar explicitly to recent advances and new research trends within the algorithmics/math programming scheduling community. This setting allowed for very high technical level talks and deep discussions on recent scheduling results, new techniques, and discussions on important open problems. The program included 15 invited main talks, 10 short spot-light talks, open problem sessions in the beginning of the week, and ample unstructured time for research and interaction. The overall atmosphere among the 45 participants was very interactive and oriented towards solving problems (also initiated by the few well-chosen application-driven talks) within new collaborations.

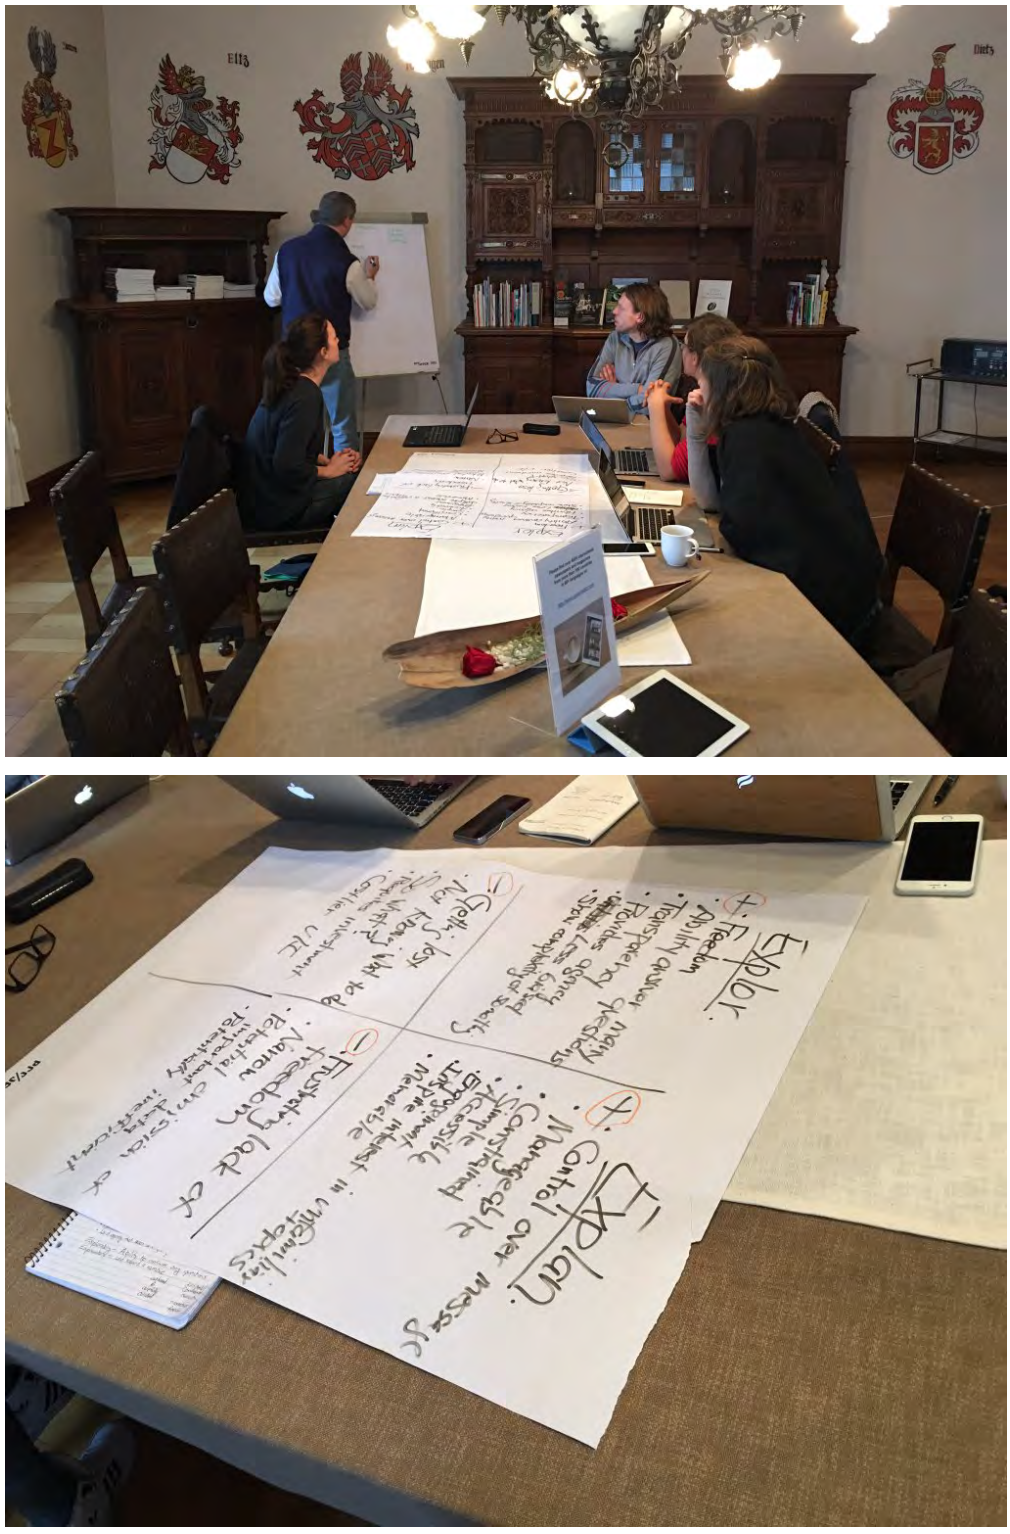

Fig. 6.6
"Explanation vs Exploration discussions #dds2016 awesome setting @dagstuhl" Twitter post by 16061 Dagstuhl Seminar participant Nathalie Henry Riche. https://twitter.com/nathriche/status/697010341264416768. Photos courtesy of Nathalie Henry Riche.

6.16 Computational Challenges in Cooperative Intelligent Urban Transport

Organizers: Caitlin Doyle Cottrill, Jan Fabian Ehmke, Franziska Klügl, and Sabine Timpf

Seminar No. 16091

Date: February 28 to March 4, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.119

Participants: Niels Agatz, Ana Lucia Bazzan, Catherine Cleophas, Caitlin Doyle Cottrill, Sybil Derrible, Ivana Dusparic, Jan Fabian Ehmke, Cecilia Gomes, Benjamin Heydecker, Andreas Hotho, Benjamin Kickhöfer, Franziska Klügl, Tobias Kretz, Ronny Kutadinata, Marco Mamei, Dirk Christian Mattfeld, Thomas Leo McCluskey, Andrea Prati, Daniele Quercia, Jörg-Rüdiger Sack, Jörn Schlingensiepen, Monika Sester, Piyushimita Vonu Thakuriah, Kevin Tierney, Sabine Timpf, Eric van Berkum, Ronald Van Katwijk, Laszlo Zsolt Varga, Giuseppe Vizzari, Ouri E. Wolfson

Following the history of two Dagstuhl seminars on Computational Issues in Transportation in 2010 and 2013, the organizers of this follow-up seminar concentrated on upcoming, data-driven challenges in the area of urban transport. In recent years, urban transportation networks have become more diverse, with a growing mix of public and private operators providing disaggregated services and information. The resulting multitude of transportation options includes non-traditional modes and services such as car and bike sharing in addition to established public transport and individual car options. So far, it is challenging to combine detailed operational data automatically arising from these services, since these data are generated both from service operation and from the users of services via crowdsourcing. The seminar aimed to discuss how data sources can be made available for individual planning and system-wide coordination of urban transportation using an approach from distributed computing, i.e., getting all involved parties to cooperate in providing relevant spatial and temporal information in a timely fashion. It was not clear how to derive reliable information for planning and control approaches, or how to adapt optimization methodologies to make urban transportation more cooperative and intelligent.

The aims of the seminar were to extend the existing network in disciplines such as Computational Traffic Science, Optimization, Autonomic Computing and Artificial Intelligence for discussing computational challenges in cooperative intelligent urban transportation, mesh communities by collecting suggestions for (partial) solutions for burning issues in urban transportation and discussing the prerequisites for merging into interdisciplinary approaches, document the state of the art and current computational challenges in cooperative intelligent transportation.

To this end, an interdisciplinary group from areas such as computer science, geography, applied optimization and traffic engineering met at Dagstuhl. The number of attendees was advantageous for group discussions, not too small for breakout groups but also not too large for meaningful discussions in the plenum.

We started on Sunday evening with a game ("Cards Against Urbanity – special issue for this seminar") specifically designed for this event by Ms. Cottrill. The game was a great success as icebreaker and helped bringing together the participants with their various backgrounds. Monday was opened with a keynote by Vonu Thakuriah, who discussed examples, prospects and challenges of emerging forms of data in transportation research and applications. The participants introduced themselves, bringing a significant object describing their relationship with the seminar's topic.

For the remaining seminar time, the participants were asked to contribute to the seminar's content by one of the following options: they could give an overview talk of an emerging area (20 minutes), a research statement on what they have been working on in their particular area (5 minutes), and they were asked to come together in groups that were defined dynamically on Monday afternoon. The resulting abstracts can be found in this report. Based on the participants' interests, groups discussing the topics of online simulation, pedestrian behavior, autonomous transportation, smart cities, and benchmark data emerged. On Wednesday afternoon, the participants went on a 'field trip' to the retail lab by DFKI in St. Wendel, where the future of retail can be explored hands-on. Since there was a significant interest in the provision of benchmark data for urban transport, there was a special session and group work devoted to this topic on Thursday afternoon. Friday morning was meant for collecting the results of the group work and collecting open challenges for future seminars.

Summarizing, the seminar identified computational challenges to cooperative intelligent urban transport, among others notably research on opportunistic groups in public transport (i.e., people sharing tickets and or trajectories in an ad-hoc fashion), freight pods attached to light rail (i.e., mixing of freight

and passenger transportation), define a common language for sharing complex knowledge and real-time data in smart cities and creating benchmark datasets for different modelling purposes and at different scales. We think that the seminar was quite successful in extending the existing networks by bringing together researchers from many different disciplines relevant for the future of urban transport. Some of the groups are planning to write proposals for the appropriate EU calls coming out in October, while others have started to work on position papers describing the state of the art as well as resulting future challenges of the field.

6.17 Computational Music Structure Analysis

Organizers: Meinard Müller, Elaine Chew, Juan Pablo Bello

Seminar No. 16092

Date: February 28 to March 4, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.2.147

Participants: Stefan Balke, Juan Pablo Bello, Frédéric Bimbot, Carmine Emanuele Cella, Elaine Chew, Ching-Hua Chuan, Roger B. Dannenberg, Matthew Davies, Christian Dittmar, Sebastian Ewert, Mary Farbood, Masataka Goto, Dorien Herremans, Andre Holzapfel, Rainer Kleinertz, Frank Kurth, Cynthia C. S. Liem, Brian McFee, Meinard Müller, Oriol Nieto, Mitchell Ohriner, Hélène Papadopoulos, Geoffroy Peeters, Christopher Raphael, Martin Rohrmeier, Mark Sandler, Xavier Serra, Jordan Smith, Anja Volk, Christof Weiß, Geraint A. Wiggins

Introduction

One of the attributes distinguishing music from other types of multimedia data and general sound sources are the rich, intricate, and hierarchical structures inherently organizing notated and performed music. On the lowest level, one may have sound events such as individual notes, which are characterized by the way they sound, i.e., their timbre, pitch and duration. Such events form larger structures such as motives, phrases, and chords, and these elements again form larger constructs that determine the overall layout of the composition. This higher structural level is specified in terms of musical parts and their mutual relations. The general goal of *music structure analysis* is to segment or decompose music into patterns or units that possess some semantic relevance and then to group these units into musically meaningful categories.

While humans often have an intuitive understanding of musical patterns and their relations, it is generally hard to explicitly describe, quantify, and capture musical structures. Because of different organizing principles and the existence of temporal hierarchies, musical structures can be highly complex and ambiguous. First of all, a temporal segmentation of a musical work may be based on various properties such as homogeneity, repetition, and novelty. While the musical structure of one piece of music may be explained by repeating melodies, the structure in other pieces may be characterized by a certain instrumentation or tempo. Then, one has to account for different musical dimensions, such as melody, harmony, rhythm, or timbre. For example, in Beethoven's Fifth Symphony the "fate motive" is repeated in various ways – sometimes the motive is shifted in pitch, sometimes only the rhythmic pattern is preserved. Furthermore, the segmentation and structure will depend on the musical context to be considered; in particular, the threshold of similarity may change depending on the timescale or hierarchical level of focus. For example, the recapitulation of a sonata may be considered a kind of repetition of the exposition on a coarse temporal level even though there may be significant modifications in melody and harmony. In addition, the complexity of the problem can depend on how the music is represented. For example, while it is often easy to detect certain structures such as repeating melodies in symbolic music data, it is often much harder to automatically identify such structures in audio representations. Finally, certain structures may emerge only in the aural communication of music. For example, grouping structures may be imposed by accent patterns introduced in performance. Hence, such structures are the result of a creative or cognitive process of the performer or listener rather then being an objective, measurable property of the underlying notes of the music.

Main Topics and Questions

In this seminar, we brought together experts from diverse fields including psychology, music theory, composition, computer science, music technology, and engineering. Through the resulting interdisciplinary discussions, we aimed to better understand the structures that emerge in composition, performance, and listening, and how these structures interrelate. For example, while there are certain structures inherent in the note content of music, the perception and communication of structure are themselves also creative acts subject to interpretation. There may be some structures intended by the composer or improviser, which are not fully communicated by symbolic descriptions such as musical score notation. The performer, if different from the composer, then must interpret structures from the score, and decide on the prosodic means by which to convey them. When a listener then tries to make sense of the performed piece, that act of sense-making, of constructing structure and meaning from an auditory stream is also a creative one. As a result, different people along this communication chain may come up with different solutions, depending on their experiences, their musical backgrounds, and their current thinking or mood.

Based on our discussions of various principles and aspects that are relevant for defining musical patterns and structures, the following questions were raised.

- How can ambiguity in notions such as repetition, similarity, grouping, and segmentation be handled and modeled?
- In which way do these notions depend on the music style and tradition?
- How can one account for the relations within and across different hierarchical levels of structural patterns?
- How can long-term structures be built up from short-term patterns, and, vice versa, how can the knowledge of global structural information support the analysis of local events?
- How can information on rhythm, melody, harmony, timbre, or dynamics be fused within unifying structural models?
- How can the relevance of these aspects be measured?
- How do computational models need to be changed to account for human listeners?

By addressing such fundamental questions, we aimed for a better understanding of the principles and model assumptions on which current computational procedures are based, as well as the identification of the main challenges ahead.

Another important goal of this seminar was to discuss how computational structure analysis methods may open up novel ways for users to find and access music information in large, unstructured, and distributed multimedia collections. Computational music structure analysis is not just an end in itself; it forms the foundation for many music processing and retrieval applications. Computational methods for structuring and decomposing digitized artifacts into semantically meaningful units are of fundamental importance not only for music content but also for general multimedia content including speech, image, video, and geometric data. Decomposing a complex object into smaller units often constitutes the first step for simplifying subsequent processing and analysis tasks, for deriving compact object descriptions that can be efficiently stored and transmitted, and for opening up novel ways for users to access, search, navigate, and interact with the content. In the music context, many of the current commercially available services for music recommendation and playlist generation employ *context-based* methods, where textual information (e. g., tags, structured metadata, user access patterns) surrounding the music object are exploited. However, there are numerous data mining problems for which context-based analysis is insufficient, as it tends to be low on specifics and unevenly distributed across artists and styles. In such cases, one requires *content-based* methods, where the information is obtained directly from the analysis of audio signals, scores and other representations of the music. In this context, the following questions were raised.

- How can one represent partial and complex similarity relations within and across music documents?
- What are suitable interfaces that allow users to browse, interact, adapt, and understand musical structures?
- How can musical structures be visualized?
- How can structural information help improve the organizing and indexing of music collections?

Participants, Interaction, Activities

In our seminar, we had 31 participants, who came from various locations around the world including North America (8 participants from the U.S.), Asia (2 participants from Japan), and Europe (21 participants from Austria, France, Germany, Netherlands, Portugal, Spain, United Kingdom). Many of the participants came to Dagstuhl for the first time and expressed enthusiasm about the open and retreat-like atmosphere. Besides its international character, the seminar was also highly interdisciplinary. While most of the participating researchers are working in the fields of music information retrieval, we have had participants with a background in musicology, cognition, psychology, signal processing, and other fields. This led to the seminar having many cross-disciplinary intersections and provoking discussions as well as numerous social activities including playing music together. One particular highlight of such social activities was a concert on Thursday evening, where various participant-based ensembles performed a wide variety of music including popular music, jazz, and classical music. Some of the performed pieces were original compositions by the seminar's participants.

Overall Organization and Schedule

Dagstuhl seminars are known for having a high degree of flexibility and interactivity, which allows participants to discuss ideas and to raise questions rather than to present research results. Following this tradition, we fixed the schedule during the seminar asking for spontaneous contributions with future-oriented content, thus avoiding a conference-like atmosphere, where the focus tends to be on past research achievements. After the organizers have given an overview of the Dagstuhl concept and the seminar's overall topic, we started the first day with self-introductions, where all participants introduced themselves and expressed their expectations and wishes for the seminar. We then continued with a small number of ten-minute stimulus talks, where specific participants were asked to address some critical questions on music structure analysis in a nontechnical fashion. Each of these talks seamlessly moved towards an open discussion among all participants, where the respective presenters took over the role of a moderator. These discussions were well received and often lasted for more than half an hour. The first day closed with a brainstorming session on central topics covering the participants' interests while shaping the overall schedule and format of our seminar. During the next days, we split into small groups, each group discussing a more specific topic in greater depth. The results and conclusions of these parallel group sessions, which lasted between 60 to 90 minutes, were then presented to, and discussed with, the plenum. Furthermore, group discussions were interleaved with additional stimulus talks spontaneously given by participants. This mixture of presentation elements gave all participants the opportunity for presenting their ideas to the plenum while avoiding a monotonous conference-like presentation format. Finally, on the last day, the seminar concluded with a session we called "self-outroductions" where each participant presented his or her personal view of the main research challenges and the seminar.

Conclusions and Acknowledgment

Having the Dagstuhl seminar, our aim was to gather researchers from different fields including information retrieval, signal processing, musicology and psychology. This allowed us to approach the problem of music structure analysis by looking at a broad spectrum of data analysis techniques (including signal processing, machine learning, probabilistic models, user studies), by considering different domains (including text, symbolic, image, audio representations), and by drawing inspiration from creative perspectives of the agents (composer, performer, listener) involved. As a key result of this seminar, we achieved some significant progress towards understanding, modeling, representing, extracting, and exploiting musical structures. In particular, our seminar contributed to further closing the gap between music theory, cognition, and the computational sciences.

The Dagstuhl seminar gave us the opportunity for having interdisciplinary discussions in an inspiring and retreat-like atmosphere. The generation of novel, technically oriented scientific

contributions was not the focus of the seminar. Naturally, many of the contributions and discussions were on a rather abstract level, laying the foundations for future projects and collaborations. Thus, the main impact of the seminar is likely to take place in the medium to long term. Some more immediate results, such as plans to share research data and software, also arose from the discussions. As measurable outputs from the seminar, we expect to see several joint papers and applications for funding.

Beside the scientific aspect, the social aspect of our seminar was just as important. We had an interdisciplinary, international, and very interactive group of researchers, consisting of leaders and future leaders in our field. Many of our participants were visiting Dagstuhl for the first time and enthusiastically praised the open and inspiring setting. The group dynamics were excellent with many personal exchanges and common activities. Some scientists expressed their appreciation for having the opportunity for prolonged discussions with researchers from neighboring research fields – some thing that which is often impossible during conference-like events.

In conclusion, our expectations of the seminar were not only met but exceeded, in particular with respect to networking and community building. We would like to express our gratitude to the Dagstuhl board for giving us the opportunity to organize this seminar, the Dagstuhl office for their exceptional support in the organization process, and the entire Dagstuhl staff for their excellent service during the seminar. In particular, we want to thank Susanne Bach-Bernhard, Roswitha Bardohl, Marc Herbstritt, and Sascha Daeges for their assistance during the preparation and organizing of the seminar.

Fig. 6.7
Dagstuhl Seminar — what a great idea Blog post by 16072 Dagstuhl Seminar participant Jan Erik Moström. http://mostrom.eu/2016/03/19/dagstuhl-seminar-what-a-great-idea/. Photo courtesy of Jan Erik Moström.

6.18 Data Structures and Advanced Models of Computation on Big Data

Organizers: Alejandro Lopez-Ortiz, Ulrich Carsten Meyer, Markus E. Nebel, and Robert Sedgewick

Seminar No. 16101

Date: March 6–11, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.3.1

Participants: Deepak Ajwani, Helmut Alt, Alexandr Andoni, Martin Aumüller, Timo Bingmann, Gerth Stølting Brodal, Andrej Brodnik, Martin Dietzfelbinger, Anne Driemel, Fabian Dütsch, Guy Even, Rolf Fagerberg, Martin Farach-Colton, Simon Gog, Mordecai Golin, Goetz Graefe, Torben Hagerup, Herman J. Haverkort, John Iacono, Riko Jacob, Tsvi Kopelowitz, Moshe Lewenstein, Alejandro Lopez-Ortiz, Jérémie Lumbroso, Conrado Martinez, Kurt Mehlhorn, Ulrich Carsten Meyer, Friedhelm Meyer auf der Heide, Ian Munro, Markus E. Nebel, Elisabeth Neumann, John D. Owens, Manuel Penschuck, Seth Pettie, Rajeev Raman, Alejandro Salinger, Peter Sanders, Robert Sedgewick, Francesco Silvestri, He Sun, Jan Vahrenhold, Sebastian Wild

About the Seminar

Data structures provide ways of storing and manipulating data and information that are appropriate for the computational model at hand. Every such model relies on assumptions that we have to keep questioning. The aim of this seminar was to exchange ideas for new algorithms and data structures, and to discuss our models of computations in light of recent technological advances. This Dagstuhl seminar was the 12th in a series of loosely related Dagstuhl seminars on data structures.

Topics

The presentations covered both advances in classic fields, as well as new models for recent trends in computing, in particular the appearance of big-data applications.

The talks by Brodal, Penschuck, Silvestri, and Vahrenhold covered methods in the *external-memory model* that models the situation that data does no longer fit into internal memory. This limit can be pushed a bit further by using *succinct* data structures, which use only as much memory as absolutely necessary. Such methods were covered in the talks of Hagerup, Raman, and Gog. If the task is to generate large random instances, Even showed that one can delay generation of large parts until they really become requested.

Big-data applications rely on *parallel* computation to speed up processing. Bingmann announced the creation of a new framework to simplify developing such applications. Brodnik presented a parallel string-searching algorithm. Since such methods are often used in a distributed setting, the cost of *communication* can become dominating. Sanders discussed several algorithms from this point of view.

Iacono and Mehlhorn reported on recent advances in the long-standing open problem of *dynamic optimality* of binary search trees (BSTs). The classic problem of finding *optimal* static *BSTs* was taken up by Munro: it becomes significantly harder if the objective is to minimize the number of *binary* comparisons instead of the classic ternary comparisons.

Wild used the connection between BSTs and recursion trees of Quicksort to analyze Quicksort on inputs with equal keys, including multi-way partitioning Quicksort. The latter was discussed in detail by Aumüller who presented a novel analysis for comparison-optimal partitioning.

Neumann introduced a new randomized dictionary implementation based on jumplists. Kopelowitz showed a much simplified solution to the file-maintenance problem.

In the context of large sparse graphs, Andoni, Fagerberg, and Sun showed how to exploit special structure in the input for algorithmic applications. Pettie showed how to efficiently answer connectivity queries in graphs when vertices can be deleted.

The seminar also enjoyed contributions on new algorithms: two innovative applications of hashing were presented by Silvestri and Jacob; Meyer auf der Heide applied the primal-dual approach for *online algorithms* to online leasing problems. Driemel reported on clustering methods for time series.

The theory-focused talks were complemented by broader perspectives from practice: Ajwani presented his vision for future communication tools that are supported by context-sensitive agents, and Sedgewick sketched his views on the future of higher education. Finally, Salinger summarized the approaches taken by SAP to include data-specific algorithms directly in their HANA database system.

New models of computation were also discussed. Owens explained how the architecture of graphic cards calls for different approaches to design data structures; Dütsch discussed the cost of virtual address translation in several algorithms. Finally, Farach-Colton and Graefe challenged the claim that data structures are independent of the application they are used in: they showed intriguing examples where the context a data structure was applied in entailed unforeseen additional requirements.

6

6.19 Rethinking Experimental Methods in Computing

Organizers: Daniel Delling, Camil Demetrescu, David S. Johnson, and Jan Vitek

Seminar No. 16111

Date: March 13–18, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.3.24

Participants: Umut A. Acar, José Nelson Amaral, David A. Bader, Judith Bishop, Ronald F. Boisvert, Marco Chiarandini, Markus Chimani, Emilio Coppa, Daniel Delling, Camil Demetrescu, Amer Diwan, Dmitry Duplyakin, Eric Eide, Erik Ernst, Sebastian Fischmeister, Norbert Fuhr, Paolo G. Giarrusso, Andrew V. Goldberg, Matthias Hagen, Matthias Hauswirth, Benjamin Hiller, Richard Jones, Tomas Kalibera, Marco Lübbecke, Catherine C. McGeoch, Kurt Mehlhorn, J. Eliot B. Moss, Ian Munro, Petra Mutzel, Luís Paquete, Mauricio Resende, Peter Sanders, Nodari Sitchinava, Peter F. Sweeney, Walter F. Tichy, Petr Tuma, Dorothea Wagner, Roger Wattenhofer

This seminar is dedicated to the memory of our co-organiser and friend David Stifler Johnson, who played a major role in fostering a culture of experimental evaluation in computing and believed in the mission of this seminar. He will be deeply missed.

David Stifler Johnson, 1945–2016.

The pervasive application of computer programs in our modern society is raising fundamental questions about how software should be evaluated. Many communities in computer science and engineering rely on extensive experimental investigations to validate and gain insights on properties of algorithms, programs, or entire software suites spanning several layers of complex code. However, as a discipline in its infancy, computer science still lags behind other long-standing fields such as natural sciences, which have been relying on the scientific method for centuries.

There are several threats and pitfalls in conducting rigorous experimental studies that are specific to computing disciplines. For example, experiments are often hard to repeat because code has not been released, it relies on stacks of proprietary or legacy software, or the computer architecture on which the original experiments were conducted is outdated. Moreover, the influence of side-effects stemming from hardware architectural features are often much higher than anticipated by the people conducting the experiments. The rise of multi-core architectures and large-scale computing infrastructures, and the ever growing adoption of concurrent and parallel programming models have made reproducibility issues even more critical. Another major problem is that many experimental works are poorly performed, making it difficult to draw any informative conclusions, misdirecting research, and curtailing creativity.

Surprisingly, in spite of all the common issues, there has been little or no cooperation on experimental methodologies between different computer science communities, who know very little of each others efforts. The goal of this seminar was to build stronger links and collaborations between computer science sub-communities around the pivotal concept of experimental analysis of software. Also, the seminar allowed exchange between communities their different views on experiments. The main target communities of this seminar were algorithm engineering, programming languages, operations research, and software engineering, but also people from other communities were invited to share their experiences. Our overall goal was to come up with a common foundation on how to evaluate software in general, and how to reproduce results. Since computer science is a leap behind natural sciences when it comes to experiments, the ultimate goal of the seminar was to make a step forward towards reducing this gap. The format of the seminar alternated talks intended for a broad audience, discussion panels, and working sessions in groups.

The organisers would like to thank the Dagstuhl team and all the participants for making the seminar a success. A warm acknowledgement goes to Amer Diwan, Sebastian Fischmeister, Catherine McGeoch, Matthias Hauswirth, Peter Sweeney, and Dorothea Wagner for their constant support and enthusiasm.

6.20 From Theory to Practice of Algebraic Effects and Handlers

Organizers: Andrej Bauer, Martin Hofmann, Matija Pretnar, and Jeremy Yallop

Seminar No. 16112

Date: March 13–18, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.3.44

Participants: Sandra Alves, Kenichi Asai, Robert Atkey, Clément Aubert, Andrej Bauer, Edwin Brady, Xavier Clerc, Stephen Dolan, Andrzej Filinski, Philipp Haselwarter, Martin Hofmann, Patricia Johann, Yukiyoshi Kameyama, Ohad Kammar, Oleg Kiselyov, Daan Leijen, Sam Lindley, Conor McBride, Gordon Plotkin, Matija Pretnar, Amr Hany Shehata Saleh, Gabriel Scherer, Tom Schrijvers, Alex Simpson, KC Sivaramakrishnan, Sam Staton, Niki Vazou, Niels Voorneveld, Leo White, Jeremy Yallop

Being no strangers to the Dagstuhl seminars we were delighted to get the opportunity to organize Seminar 16112. Our seminar was dedicated to algebraic effects and handlers, a research topic in programming languages which has received much attention in the past decade. There are strong theoretical and practical aspects of algebraic effects and handlers, so we invited people from both camps. It would have been easy to run the seminar as a series of disconnected talks that would take up most of people's schedules – we have all been to such seminars – and run the risk of disconnecting the camps as well. We decided to try a different format, and would like to share our experience in this executive summary.

On the first day we set out to identify topics of interest and organize working groups around them. This did not work, as everybody wanted to be in every group, or was at least worried they would miss something important by choosing the wrong group. Nevertheless, we did identify topics and within them ideas began to form. At first they were very general ideas on the level of major research projects, but soon enough people started asking specific questions that could be addressed at the seminar. Around those questions small groups began to form. Out of initial confusion came self-organization.

We had talks each day in the morning, with the schedule planned two days ahead, except for the first day which started by a tutorial on algebraic effects and handlers. We left the afternoons completely free for people to work in self-organized groups, which they did. The organizers subtly made sure that everybody had a group to talk to. In the evening, just before dinner, we had a "show & tell" session in which groups reported on their progress. These sessions were the most interesting part of the day, with everyone participating: some showing what they had done so far, and others offering new ideas. Some of the sessions were accompanied by improvised short lectures.

Work continued after dinner and late at night. One of the organizers was shocked to find, on his way to bed, that the walls of a small seminar room were completely filled with type theoretic formulas, from the floor to the ceiling. He was greatly relieved to hear that the type theory was not there to stay permanently as the Dagstuhl caretakers painted the walls with a special "whiteboard" paint. They should sell the paint by the bucket as a Dagstuhl souvenir.

We are extremely happy with the outcome of the seminar and the way we organized it. An open format that gives everyone ample time outside the seminar room was significantly boosted by the unique Dagstuhl environment free of worldly distractions. We encourage future organizers to boldly try new ways of organizing meetings. There will be confusion at first, but as long as the participants are encouraged and allowed to group themselves, they will do so. If a lesson is to be taken from our seminar, it is perhaps this: let people do what they want, but also make sure they report frequently on what they are doing, preferably when they are a bit hungry.

6.21 Language Based Verification Tools for Functional Programs

Organizers: Marco Gaboardi, Suresh Jagannathan, Ranjit Jhala, and Stephanie Weirich

Seminar No. 16131

Date: March 28 to April 1, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.3.59

Participants: Andreas Martin Abel, Amal Ahmed, Andrew W. Appel, Lennart Augustsson, Edwin Brady, Iavor Diatchki, Richard A. Eisenberg, Jean-Christophe Filliâtre, Cormac Flanagan, Marco Gaboardi, Deepak Garg, Michael Greenberg, Reiner Hähnle, Cătălin Hriţcu, Suresh Jagannathan, Ranjit Jhala, Gowtham Kaki, Gabriele Keller, Naoki Kobayashi, Ekaterina Komendantskaya, Martin Lange, K. Rustan M. Leino, Conor McBride, Jan Midtgaard, Chih-Hao Luke Ong, Dominic Orchard, Brigitte Pientka, Ruzica Piskac, Nadia Polikarpova, Scott Smith, Matthieu Sozeau, Wouter Swierstra, Tachio Terauchi, Sam Tobin-Hochstadt, Hiroshi Unno, David Van Horn, Niki Vazou, Stephanie Weirich, Nobuko Yoshida

The web, multi-core and "big-data" revolutions have been largely built on higher-order programming constructs pioneered in the Functional Programming community. Despite the increasing importance of such programs, there are relatively few tools that are focussed on ensuring that functional programs possess crucial correctness properties. While language based verification for imperative and first-order programs has been studied for decades yielding important ideas like Floyd-Hoare Logics, Abstract Interpretation, and Model Checking. It is only relatively recently, that researchers have proposed language based verification tools e.g. advanced type systems, contract systems, model checking and higher-order program analyses for functional and higher-order programs.

We organised this seminar to bring together the different schools of researchers interested in software reliability, namely, the designers and implementers of functional programming languages, and experts in software verification, in order create a larger community of researchers focused on this important goal, to let us compare the strengths and limitations of different approaches, to find ways to unite both intellectually, and via tools the complementary advantages of different techniques, and to devise challenging open problems and application areas where verification may be most effective. To this end, the seminar comprised a program of 30 talks from the leading experts on the above topics, and breakout sessions on:

1. Integrating formal methods tools in the curriculum
2. Hands on Tool Tutorials
3. User Interaction
4. Types and Effects

6.22 Analysis, Interpretation and Benefit of User-Generated Data: Computer Science Meets Communication Studies

Organizers: Thorsten Quandt, German Shegalov, Helle Sjøvaag, and Gottfried Vossen

Seminar No. 16141

Date: April 3–8, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.1

Participants: Christian Baden, David Domingo, Martin Emmer, Raul Ferrer Conill, Johann-Christoph Freytag, Elisabeth Günther, Krishna P. Gummadi, Ari Heinonen, Jukka Huhtamäki, Seth C. Lewis, Alexander Löser, Wiebke Loosen, Truls Pedersen, Thorsten Quandt, Tatjana Scheffler, Ralf Schenkel, German Shegalov, Helle Sjøvaag, Hendrik Stange, Eirik Stavelin, Martin Theobald, Heike Trautmann, Damian Trilling, Gottfried Vossen, Rodrigo Zamith

The success of the Internet as a communication technology and tool for human interaction in countless contexts, including production and trade, has had a dramatic impact on modern societies. With diffusion rates nearing one hundred percent in most societal groups, there is virtually no one whose life is not influenced by online communication – either directly or indirectly. Every day, private end users and business users act and interact online, producing immense amounts of data. Many disciplines, including computer science, computer linguistics, psychology, and communication studies, have identified 'big data' generated by online users as a research field. As a result, big data has become a somewhat over-hyped catch-all term for many different types of data, which are analyzed through varying methods for multiple purposes. This ranges from an analysis of (unstructured) Twitter or Facebook content to rule-structured texts as found in the professional media (i.e., news websites). The implication of value generated through sheer size of data sets is misleading, though – much of the value is based on the nature of these data sets as being user-generated, either on purpose or inevitably (and often unknowingly) as behavioral traces of actions with divergent aims.

Big data sets generated by human users pose some challenges to the scientific disciplines that are interested in them: Typically, computer scientists have the knowledge and tools to access, extract and process big data sets. However, the analysis and interpretation of such data mirrors the interactions of users who produced the data and is not following a purely technological logic. In other words, such data has a human/social component, and in order to interpret and understand it, social-scientific theories and methods are helpful. Social scientists, however, typically do not specialize in the practicalities of online technologies and of programming. While they have theoretical approaches and empirical methods available that can be helpful in the analysis of user generated content – and this is especially true for communication scholars who specialize in the analysis of (online) media content –, their possibilities to access and process data are limited (as this is not core to their field yet).

Consequently, both disciplinary approaches will not be able to fully address the challenges of analyzing big data based on user (inter)action from the perspective of their own 'silo'. A combination of the two approaches seems fruitful, as each discipline may help in solving the problems of the other, and the sum will be more than its parts – leading to a better understanding of social interaction and human communication in a digitized world. This seminar will bring together both computer scientists interested in the analysis of (large-scale) user-generated data, and communication scholars interested in computer-assisted acquisition and processing of such data. It is intended to start a fruitful dialogue on potential approaches, methods, uses and benefits of a cooperation between the two disciplines, and it will also include the input of practitioners in the field of media and business who will offer valuable insights into practical use cases.

Fig. 6.8
Drawing for the Dagstuhl children's guest book.

6.23 Multidisciplinary Approaches to Multivalued Data: Modeling, Visualization, Analysis

Organizers: Ingrid Hotz, Evren Özarslan, and Thomas Schultz

Seminar No. 16142

Date: April 3–8, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.16

Participants: Burak Acar, Maryam Afzali-Deligani, Bernhard Burgeth, Tom Dela Haije, Anders Eklund, Aasa Feragen, Luc Florack, Andrea Fuster, Hans Hagen, Ingrid Hotz, Andrada Ianuş, Cheng Guan Koay, Chunlei Liu, Georgeta Elisabeta Marai, Rodrigo Moreno, Jaya Sreevalsan Nair, Jos B.T.M. Roerdink, Evren Özarslan, Benoit Sherrer, Gerik Scheuermann, Thomas Schultz, Alexandra Tobisch, Xavier Tricoche, Gözde Ünal, Anna Vilanova Bartroli, Yue Zhang, Valentin Zobel

Topics and Motivation

This seminar is the sixth in a series of Dagstuhl Seminars devoted to the use of tensor fields and other higher order descriptors, including higher-order tensors or Spherical Harmonics, to model intricate multivalued data that arises in modern medical imaging modalities, as well as in simulations in engineering and industry. Even though the literature on image analysis, visualization, as well as statistical hypothesis testing and machine learning is quite rich for scalar or vector-valued data, relatively little work has been performed on these disciplines for tensors and higher-order descriptors.

Applications wherein such descriptors can be employed to process multivalued data range from neuroimaging to image analysis and engineering. Diffusion Weighted Magnetic Resonance Imaging (DW-MRI), an MRI modality which makes it possible to visualize and quantify structural information about white matter pathways in the brain *in vivo,* is one of the driving technologies, but tensors have also shown their usefulness as feature descriptors for segmentation and grouping in image analysis, including structure tensors and tensor voting. Applications in solid mechanics, civil engineering, computational fluid dynamics and geology require the processing of tensor fields as part of domain-specific modeling, simulation, and analysis (e.g. stress-strain relationships, inertia tensors, permittivity tensor).

The Dagstuhl seminar provides a unique platform by facilitating scientific exchange between key researchers in seemingly diverse applications. Despite these disciplines' commonalities in terms of the tools employed, it would be very unlikely that these scientists would attend the same conference as the theme of most conferences is defined by a specific application. By bringing together specialists in visualization, image processing, statistics, and numerical mathematics, the Dagstuhl seminar provides new impulses for methodological work in those areas.

Organization of the Seminar

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 26 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. For the traditional social event on Wednesday, we went on a hike, which was joined by almost all participants, and offered additional welcome opportunities for interaction.

Three breakout sessions were organized in the afternoons, and another one in the evening, so that none of them took place in parallel, and everyone had the opportunity to visit all groups relevant to him or her. The topics of the four groups were formed by clustering topics brought up in the round of introductions, and were denoted as:

- Visual encodings and the interface between theory and applications
- Models and geometry
- Topological methods
- Multi-field and tensor group analysis

Depending on the interests of the participants, the breakout groups differed in nature, ranging from the collection of open questions

and discussions on future directions of the field to spontaneous tutorial-style presentations. Notes taken during these sessions, and the main results of two of them are summarized in this report.

Outcomes

The participants all agreed that the meeting was successful and stimulating, and we plan to publish another Springer book documenting the results of the meeting. Participants have pre-registered thirteen 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 2017.

It was voted that the group will apply for another meeting in this series. In addition to the current organizers Thomas Schultz (University of Bonn, Germany) and Evren Özarslan (Linköpings Universitet, Sweden), Andrea Fuster (TU Eindhoven, The Netherlands) and Eugene Zhang (Oregon State University, USA) agreed to help apply for the next event.

Acknowledgments

The organizers thank all the attendees for their contributions and extend special thanks to the team of Schloss Dagstuhl for helping to make this seminar a success. As always, we enjoyed the warm atmosphere, which supports formal presentations as well as informal exchange of ideas.

6.24 Foundations of Data Management

Organizers: Marcelo Arenas, Richard Hull, Wim Martens, Tova Milo, and Thomas Schwentick

Seminar No. 16151

Date: April 10–15, 2016 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.6.4.39

Participants: Serge Abiteboul, Marcelo Arenas, Pablo Barcelo, Meghyn Bienvenu, Iovka Boneva, Angela Bonifati, Diego Calvanese, Claire David, Daniel Deutch, Thomas Eiter, Ronald Fagin, Floris Geerts, Georg Gottlob, Torsten Grust, Paolo Guagliardo, Eyke Hüllermeier, Richard Hull, Benny Kimelfeld, Phokion G. Kolaitis, Maurizio Lenzerini, Leonid Libkin, Carsten Lutz, Wim Martens, Tova Milo, Filip Murlak, Frank Neven, Dan Olteanu, Magdalena Ortiz, Reinhard Pichler, Andreas Pieris, Juan L. Reutter, Sudeepa Roy, Thomas Schwentick, Jianwen Su, Dan Suciu, Jens Teubner, Jan Van den Bussche, Stijn Vansummeren, Victor Vianu, Domagoj Vrgoc, Ke Yi

The focus of Foundations of Data Management (traditionally termed Database Theory) is to provide the many facets of data management with solid and robust mathematical foundations. The field has a long and successful history and has already grown far beyond its traditional scope since the advent of the Web.

The recent push towards Big Data, including structured, unstructured and multi-media data, is transforming and expanding the field at an unusually rapid pace. However, for understanding numerous aspects of Big Data, a robust research exploration into the principled foundations is still lacking. This transformation will call upon the Database Theory community to substantially expand its body of tools, techniques, and focal questions and to much more fully embrace several other disciplines, most notably statistics and probability theory, natural language processing, data analytics, emerging hardware and software supports for computation, and data privacy and security.

Big Data is not the only force that is driving expansion and transformation for the Foundations of Data Management. With the increasing digitization of diverse industries, including "smarter cities", education, healthcare, agriculture and others, many diverse kinds of data usage at large scales are becoming crucial. The push towards data-centric business processes, which are especially important for knowledge-worker driven processes, raise fundamentally new questions at the intersection of data and process. And increasing adoption of semantic web and other ontology-based approaches for managing and using meta-data push the boundaries of traditional Knowledge Representation.

The purpose of this Dagstuhl Perspectives Workshop was to explore the degree to which principled foundations are crucial to the long-term success and effectiveness of the new generation of data management paradigms and applications, and to understand what forms of research need to be pursued to develop and advance these foundations.

For this workshop we brought together specialists from the existing database theory community, and from adjoining areas, such as Machine Learning, Database Systems, Knowledge Representation, and Business Process Management, to understand the challenge areas that might be resolved through principled foundations and mathematical theory.

More specifically, during this workshop we worked on:

- Identifying areas, topics and research challenges for Foundations of Data Management in the forthcoming years, in particular, areas that have not been considered as Database Theory before but will be relevant in the future and of which we expect to have papers at PODS and ICDT, the main conferences in the field.
- Outlining the techniques that will be most fruitful as starting points for addressing the new foundational challenges in Data Management.
- Characterising the major challenge areas in Big Data where a principled, mathematically-based approach can provide important contributions.
- Finding research goals in neighbouring areas that may generate synergies with our own.

The workshop consisted of eight invited tutorials on selected topics: (1) Managing Data at Scale, (2) Uncertainty and Statistics in Foundations of Data Management, (3) Human in the Loop in Data Management, (4) Machine Learning and Data Management, (5) Data-Centric Business Processes and Workflows, (6) Ethical Issues in Data Management, (7) Knowledge Representation, Ontologies, and Semantic Web, and (8) Classical DB Questions on New Kind of Data. The abstracts of these talks can be found below in the document.

There were also seven working groups on theory-related topics, which identified the most relevant research challenges for Foundations of Data Management in the forthcoming years, outlined the mathematical techniques required to tackle such

problems, and singled out specific topics for insertion in a curriculum for the area. The topics of these working groups were: (1) Imprecise Data, (2) Unstructured and Semi-structured Data, (3) Process and Data, (4) Data Management at Scale, (5) Data Management and Machine Learning, (6) Knowledge-Enriched Data Management, and (7) Theory and Society. There was also a working group on curriculum related issues, that collected and enriched the information provided by the working groups about the design of a curriculum on Foundations of Data Management. Each one of these groups worked for two consecutive hours in different days. Workshop participants had to participate in at least two working groups, although most of the people participated in four of them. Summaries of the discussions held in each one of these working groups can be found below in the document.

During the first day of the workshop, there were also five working groups that analysed several community-related aspects. In particular: (1) Attraction of women and young members, (2) cross-fertilization with neighbouring areas, (3) relationship to industry, (4) impact of our research, and (5) the publishing process. The discussion within some of these working groups gave rise to the creation of specific tasks to be accomplished by the community in the following years. These tasks will be coordinated by the councils of PODS and ICDT, the two main conferences in the field.

This Dagstuhl Report will be accompanied by a Dagstuhl Manifesto, in which the outcome of the different working groups will be explained in more detail and several strategies for the development of our field will be proposed.

6.25 Tensor Computing for Internet of Things

Organizers: Evrim Acar, Animashree Anandkumar, Lenore Mullin, Sebnem Rusitschka, and Volker Tresp

Seminar No. 16152

Date: April 10–13, 2016 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.6.4.57

Participants: Evrim Acar, Kareem Aggour, Animashree Anandkumar, Rasmus Bro, Ali Taylan Cemgil, Edward Curry, Lieven De Lathauwer, Hans Hagen, Souleiman Hasan, Denis Krompaß, Gerwald Lichtenberg, Benoit Meister, Lenore Mullin, Morten Mørup, Axel-Cyrille Ngonga-Ngomo, Ivan Oseledets, Renato Pajarola, Vagelis Papalexakis, Christine Preisach, Achim Rettinger, Sebnem Rusitschka, Volker Tresp, Bülent Yener

In April 2016, Dagstuhl hosted a Perspectives Workshop on Tensor Computing for the Internet of Things. The prior year, industrial researchers had formulated the challenges of gaining insights from multi-dimensional sensory data coming from large-scale connected energy, transportation networks or manufacturing systems. The sheer amount of streaming multi-aspect data was prompting us to look for the most suitable techniques from the machine learning community: multi-way data analysis. Hence, we organized a three-day interactive workshop with two separate questions bringing two formerly distinct communities together: (i) How can we assure performance and reliability given the increasing complexity and data of an always-on connected world? (ii) Can we exploit the power of tensor algebra to solve high-dimensional large-scale machine learning problems that such a world poses?

The workshop focused on the Internet of Things (IoT), i.e. devices, which have the capability to sense, communicate, and more so, control their environments. These devices are increasingly becoming a part of complex, dynamic, and distributed systems of electricity or mobility networks, hence our daily lives. Various sensors enable these devices to capture multiple aspects of their surroundings in real-time. For example, phasor measurement units capture transient dynamics and evolving disturbances in the power system in high-resolution, in a synchronized manner, and in real-time. Another example is traffic networks, where a car today can deliver about 250 GB of data per hour from connected electronics such as weather sensors within the car, parking cameras and radars. Experts estimate that the IoT will consist of almost 50 billion objects by 2020 [2], which will trigger the Era of Exascale computing necessitating the management of heat and energy of computing in concert with more and more complex processor/network/memory hierarchies of sensors and embedded computers in distributed systems. Crucial for the extraction of relevant information is the format in which the raw data from such systems is represented. Crucial for viable efficiency of information extraction in IoT is which operations are used guaranteeing various attributes of resource use and management. Tensors can be viewed as data structures or as multilinear operators.

The goal of the workshop was to explore tensor representations and computing as the basis for machine learning solutions for the IoT. Tensors are algebraic objects which describe linear and multilinear relationships, and can be represented as multi-dimensional arrays. They often provide a natural and compact representation for multidimensional data. In the recent years, tensor and machine learning communities – mainly active in the data-rich domains such as neuroscience, social network analysis, chemometrics, knowledge graphs etc. – have provided a solid research infrastructure, reaching from the efficient routines for tensor calculus to methods of multi-way data analysis, i.e., tensor decompositions, to methods for consistent and efficient estimation of parameters of the probabilistic models.

Some tensor-based models have the intriguing characteristic that if there is a good match between the model and the underlying structure in the data, the models are much better interpretable than alternative techniques. Their interpretability is an essential feature for the machine learning techniques to gain acceptance in the rather engineering heavy fields of automation and control of cyber-physical systems. Many of these systems show intrinsically multilinear behavior, which is appropriately modeled by tensor methods and tools for controller design can use these models. The calibration of sensors delivering data and the higher resolution of measured data will have an additional impact on the interpretability of models.

Various presentations on tensor methods by established researchers from different application domains underscored that tensor methods are reaching a maturity tipping point. However, knowledge of usage characteristics of tensor models is scattered.

Discussions of the currently independent perspectives on the usage of tensor methods showed convergence potential which we will detail in the Dagstuhl Manifesto. During our discussions based on the presentations of the IoT industrial researchers, it quickly became clear that we would need benchmark challenges for cyber-physical systems and benchmark data in order to be able to replicate the successes in machine learning for neuroscience, image processing or chemometrics, for example.

The tensor computing community will equally benefit from the new types of data, requirements and characteristics of IoT, which can lead to techniques that increase success rates of previous applications, as was the case with the challenges of social network data analysis leading to better tensor models/algorithms that can analyze data sets with missing entries, now used in many other fields in addition to social network analysis. Additionally, as opposed to standardized machine learning techniques, tensor computing currently lacks a common language and the homogeneity to flexibly exchange models. Hence, a hub platform bringing data and domain knowledge of cyber-physical systems together with a variety of practitioners of tensor computing would enhance increasing coherence of terms, best practices in data acquisition and structuring methods as well as model benchmarking, cataloging, and exchange of methods.

Furthermore, industrial researchers from IoT, automation and control domains highlighted their view that tensor computing methods are currently still inaccessible to the majority of the industrial practitioners even though there has been a considerable progress in developing tools for tensor computing. Matlab extensions to enable the use of tensor analysis are quite mature [1] [3] [4]. Matlab is widely used by control and automation practitioners. Python ecosystem for machine learning practitioners is very quickly adopting extensions for enabling tensor operations [5] [6]. However, both are mainly for prototyping and ultimately do not fulfill the need for a unified framework for industrial grade development and deployment of models in highly distributed cyber-physical systems. Interestingly, just five months prior to our workshop, Tensorflow [7], a numerical computation library aiming at capturing structures in multidimensional data as well as supporting both prototyping and production level algorithms was open sourced. Tensorflow can run on server clusters as well as embedded systems such as smart phones [8]. Another framework, unifying both batch and streaming data analysis, is Apache Spark [9]. Spark provides seamless scalability of software code to run on multiple machines. Recently there have been deployments of tensor methods on the Spark platform.

As a multidisciplinary community we believe that we will be able to formulate requirements and provide support in developing improvements for unifying frameworks. The required skill set is quite rare: we are in need of software developers that can create reliable high-performant code for both server-side distributed training on massive amounts of data and deployment of trained models in embedded distributed system. Heterogeneous processor architectures are predominant in cyber-physical systems. Either these software developers should be data scientists proficient in tensor computing and very good at communicating with domain experts or we need tooling such that data scientists and domain experts can collaboratively model data for cyber-physical systems. We will detail these discussions in the Manifesto: We believe that it is feasible to create such tooling that automates the generation of reliable, secure code, which accounts for the adaptive logic of devices interacting with their dynamic physical environment – but also through which there is a direct feedback between data scientist, domain or control expert, and the adaptive control device.

The Manifesto, which will be published on http://www.dagstuhl.de/16152/, will include a roadmap of how we as a newly formed multidisciplinary community want to start with a knowledge hub on tensors, and iterate through data grand challenges from IoT pilots, results dissemination, into what may one day become collaborative modeling hub for learning cyber-physical systems.

References

1 Tamara G. Kolda, Brett W. Bader and others. Matlab tensor toolbox version 2.6. http://www.sandia.gov/~tgkolda/TensorToolbox/index-2.6.html

2 Statista. Internet of things (iot): number of connected devices worldwide from 2012 to 2020 (in billions). http://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/

3 Ivan Oseledets. Tt (tensor train) toolbox version 2.2.2. https://github.com/oseledets/TT-Toolbox

4 Laurent Sorber Marc Van Barel Nico Vervliet, Otto Debals and Lieven De Lathauwer. Tensor lab user guide, release 3.0. http://www.tensorlab.net/userguide3.pdf

5 Ivan Oseledets. Python implementation of the tt-toolbox. https://github.com/oseledets/ttpy

6 Maximilian Nickel. scikit-tensor: Python library for multilinear algebra and tensor factorizations. https://github.com/mnick/scikit-tensor

7 Google. Tensorflow: An open source software library for numerical computation using data flow graphs. https://www.tensorflow.org/

8 Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zhang. Tensorflow: A system for large-scale machine learning. *CoRR*, abs/1605.08695, 2016

9 Apache Spark. Apache spark: A fast and general engine for large-scale data processing. http://spark.apache.org/

6.26 Natural Language Argumentation: Mining, Processing, and Reasoning over Textual Arguments

Organizers: Elena Cabrio, Graeme Hirst, Serena Villata, and Adam Wyner

Seminar No. 16161

Date: April 17–22, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.80

Participants: Laura Alonso Alemany, Michal Araszkiewicz, Kevin D. Ashley, Pietro Baroni, Floris Bex, Katarzyna Budzynska, Elena Cabrio, Claire Cardie, Federico Cerutti, Giorgos Flouris, Nancy L. Green, Iryna Gurevych, Ivan Habernal, Graeme Hirst, Eduard H. Hovy, Anthony Hunter, Diane J. Litman, Bernardo Magnini, Robert Mercer, Marie-Francine Moens, Smaranda Muresan, Vincent Ng, Sebastian Padó, Fabio Paglieri, Alexis M. Palmer, Andreas Peldszus, Ariel Rosenfeld, Patrick Saint-Dizier, Jodi Schneider, Edwin Simpson, Noam Slonim, Jan Šnajder, Manfred Stede, Benno Stein, Simone Teufel, Francesca Toni, Leon van der Torre, Serena Villata, Vern R. Walker, Zhe Yu

Philosophers and, in more recent years, theorists working largely within Artificial Intelligence have developed formal computational models of argumentation, how it works, and what makes an argument valid or invalid. This work has made substantial progress in abstract, formal models to represent and reason over complex argumentation structures and inconsistent knowledge bases. Relatively little research, however, has applied these computational models to naturally occurring argumentation in text; nor have Computational Linguistics and Natural Language Processing substantially examined argumentation in text. Moreover, much of the work to date has studied only domain-specific texts and use-cases. Examples include finding the specific claims made in a scientific paper and distinguishing argumentation from narrative in legal texts.

But there are many uses and applications for automatic processing of the argumentative aspects of text, such as summarizing the argument of a complex court decision, helping a writer to structure an argument, and processing a large corpus of texts, such as blogs or consumer comments, to find arguments within it. To identify and integrate arguments across a corpus is a very significant problem. To address the issues, solve problems, and build applications, tools must be developed to analyze, aggregate, synthesize, structure, summarize, and reason about arguments in texts. Such tools would enable users to search for particular topics and their justifications, to trace through the argument, and to systematically and formally reason about the relations among arguments. However, to do so requires more linguistic sophistication and newer techniques than currently found in NLP. Moreover, NLP approaches must be connected to computational models of argument. The issues and problems have started to receive attention from both communities; for example, legal documents, on-line debates, product reviews, newspaper articles, court cases, scientific articles, and other kinds of text have all been the subject of recent NLP research on argumentation mining and have been tied to computational models.

Because argumentation is an inherently cross-disciplinary topic involving philosophy, psychology, communications studies, linguistics, and computer science, where different interpretations, analyses, and uses of arguments are proposed and applied, for progress in building NLP tools for argumentation there needs to be progress not only within each domain, but in bridging between these various disciplines, Natural Language Processing, and the computational models. This seminar aimed to help build this bridge by bringing together researchers from different disciplines, with the following goals:

- To understand better the specific kinds of tasks that NLP can carry out in argumentation.
- To establish a set of domain-specific and cross-domain use-cases that will guide the direction of research in the field.
- To understand better how computational argumentation tasks are tied – or not tied – to their specific domains, such as scientific papers, legal argumentation, and political discussions, looking for new cross-domain generalizations.
- To understand better the technical challenges to success in each of these tasks, and to discuss how the challenges can be addressed.
- To develop and explicate specific challenge problems for the integration of argumentation theory and NLP that are beyond the state of the art (but not too much so), and in which success would have the greatest effect on the field.
- To provide prototype solutions that address issues in the integration of NLP and argumentation theory, and to outline follow-on development.
- To propose or provide preliminary solutions to common open challenges in natural language argumentation (among others: argument retrieval in text, argument summarization, identification of semantic relations among arguments), profiting

from the cross-fertilization between researchers coming from the research areas of NLP and formal argumentation.

The seminar was held on 17–22 April 2016, with 40 participants from 14 different countries. The event's seven sessions included 30 talks, two tutorials and a hands-on "unshared" task. The program included several plenary presentations and discussions in smaller working groups. The presentations addressed a variety of topics, as argument mining applied to legal argumentation and to writing support. Collective discussions were arranged for most of these topics, as well as plans for a future interdisciplinary research agenda involving experts from social sciences and psychology.

As a result of the seminar, a number of challenges and open issues have been highlighted:

- At this stage of maturity of the research area, it is difficult to choose good (possibly new) challenges and to define the task(s) to be addressed by automated systems
- Similarly, it is also challenging to precisely define and accomplish annotation task(s) to establish benchmarks and gold standards to test such automated systems
- It is essential to the fruitful development of the research area establish an Interdisciplinary outreach, involving social sciences, psychology, and economics.

Addressing these issues and other questions is now on the agenda of the Argument Mining research community.

6.27 Managing Technical Debt in Software Engineering

Organizers: Paris Avgeriou, Philippe Kruchten, Ipek Ozkaya, and Carolyn Seaman

Seminar No. 16162

Date: April 17–22, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.110

Participants: Areti Ampatzoglou, Francesca Arcelli Fontana, Rami Bahsoon, Ayse Basar Bener, Frank Buschmann, Alexandros Chatzigeorgiou, Zadia Codabux, Bill Curtis, Steven D. Fraser, Alfredo Goldman, Christine Hofmeister, Johannes Holvitie, Clemente Izurieta, Andreas Jedlitschka, Sven Johann, Heiko Koziolek, Philippe Kruchten, Jean-Louis Letouzey, Antonio Martini, John D. McGregor, Mehdi Mirakhorli, J. David Morgenthaler, Robert Nord, Ipek Ozkaya, Fabio Queda Bueno da Silva, Klaus Schmid, Carolyn Seaman, Andriy Shapochka, Forrest Shull, Will Snipes, Damian Andrew Tamburri, Graziela Tonin, Guilherme Horta Travassos

The term *technical debt* refers to delayed tasks and immature artifacts that constitute a "debt" because they incur extra costs in the future in the form of increased cost of change during evolution and maintenance. The technical debt metaphor provides an effective mechanism for communicating design trade-offs between developers and other decision makers. When managed effectively, technical debt provides a way to gauge the current maintainability of a system and correct the course when that level is undesirable. While other software engineering disciplines – such as software sustainability, maintenance and evolution, refactoring, software quality, and empirical software engineering – have produced results relevant to managing technical debt, none of them alone suffice to model, manage, and communicate the different facets of the design trade-off problems involved in managing technical debt.

Despite recent progress by the research community in understanding technical debt, increased attention by tool vendors on assessing technical debt through code conformance checking, and collaboration with industry in sharing data and challenges, there are several open questions about the role of technical debt in software development. The goal of this seminar was to establish a common understanding of key concepts of technical debt and build a road map for future work in this area to address these open questions.

How do we define and model technical debt? The software engineering community is converging on defining *technical debt* as making technical compromises that are expedient in the short term, but that create a technical context that increases complexity and cost in the long term. While the conceptual roots of technical debt imply an idealized, deliberate decision-making process and rework strategy as needed, we now understand that technical debt is often incurred unintentionally and catches software developers by surprise. Hence, it is mostly observed during maintenance and evolution. Technical debt as a metaphor serves as a strong communication mechanism, but the community now understands that technical debt is also a software development artifact. This overloaded nature creates confusion, especially for newcomers to the field. In addition, there is a Well-defined benchmarks provide a basis for evaluating new approaches and ideas. They are also an essential first step toward creating an empirical basis on which work in this area can grow more effectively. Effective and well-accepted benchmarks allow researchers to validate their work and tailor empirical studies to be synergistic. Technical debt's evolving definition and its sensitivity to context have inhibited the development of benchmarks so far. An ideal benchmark for technical debt research would consist of a code base, architectural models (perhaps with several versions), and known technical-debt items (TD items). New approaches to identify technical debt could be run against these artifacts to see how well the approaches reveal TD items. Industry needs guidance for how and what data to collect and what artifacts they can make available to enable progress in understanding, measuring, and managing technical debt. risk of associating anything detrimental to software systems and development processes with technical debt. This risk necessitates crisply defining both technical debt and related concepts.

How do we manage technical debt? Managing technical debt includes recognizing, analyzing, monitoring, and measuring it. Today many organizations do not have established practices to manage technical debt, and project managers and developers alike are longing for methods and tools to help them strategically plan, track, and pay down technical debt. A number of studies have examined the relationship between software code quality and technical debt. This work has applied detection of "code smells" (low internal code quality), coupling and cohesion, and dependency analysis to identify technical debt. However, empirical examples collected from industry all point out that the most significant technical debt is caused by design trade-offs,

which are not detectable by measuring code quality. Effective tooling to assist with assessing technical debt remains a challenge for both research and industry.

How do we establish an empirical basis and data science for technical debt? Well-defined benchmarks provide a basis for evaluating new approaches and ideas. They are also an essential first step toward creating an empirical basis on which work in this area can grow more effectively. Effective and well-accepted benchmarks allow researchers to validate their work and tailor empirical studies to be synergistic. Technical debt's evolving definition and its sensitivity to context have inhibited the development of benchmarks so far. An ideal benchmark for technical debt research would consist of a code base, architectural models (perhaps with several versions), and known technical-debt items (TD items). New approaches to identify technical debt could be run against these artifacts to see how well the approaches reveal TD items. Industry needs guidance for how and what data to collect and what artifacts they can make available to enable progress in understanding, measuring, and managing technical debt.

Seminar Format

In this seminar, we brought together researchers, practitioners, and tool vendors from academia and industry who are interested in the theoretical foundations of technical debt and how to manage it from measurement and analysis to prevention. Before the seminar, the organizers created a blog where attendees could post positions and start discussions to facilitate seeding of ideas.

Before the seminar, the organizers grouped discussions and blog entries into relevant themes that included creating a common definition and conceptual model of technical debt, measurement and analysis of technical debt, management of technical debt, and a research road map for managing technical debt.

Our goal was to make this seminar a working week; hence we had a dynamic schedule. We did not feature any long talks. Each day had three types of sessions. There was a plenary session for "lightning talks," in which each presenter had 10 minutes for presentation and questions on each day except for the last day of the seminar. The second type of session was for breakout discussions. Breakout sessions focused on themes that emerged from the blog and the goals of the seminar. Participants first discussed these in randomly assigned small groups in order to maximize cross-pollination of ideas. Last, we had plenary discussion sessions to collate and summarize the discussions during the breakouts. At the end of each day, the organizers asked for feedback and adjusted the flow of the following day accordingly. As a result, we dedicated the fourth day of the seminar to an "un-conference" format in which the discussion topic emerged based on the interests and votes of the attendees.

The Definition of Technical Debt and a Conceptual Model

At the conclusion of the seminar, attendees agreed on the following working definition of technical debt, which we refer to as the 16162 definition of technical debt:

> In software-intensive systems, technical debt is a collection of design or implementation constructs that are expedient in the short term, but set up a technical context that can make future changes more costly or impossible. Technical debt presents an actual or contingent liability whose impact is limited to internal system qualities, primarily maintainability and evolvability.

A significant outcome of the week was the recognition that, similar to other complex software engineering artifacts, technical debt is best described through multiple viewpoints. Concepts related to technical debt in particular should be discussed based on two related viewpoints:

1. the viewpoint describing the properties, artifacts, and elements related to technical debt items
2. the viewpoint articulating the management- and process-related activities to perform, or the different states that debt may go through

Figure 6.9 shows the initial conceptual model that served as the starting point for discussions. This model helped the group converge on key concepts. Mismatches occurred when the discussions focused on causes that may or may not be input to measurement and analysis. The dynamic view is intended to articulate these aspects.

The technical debt associated with a software-intensive system is composed of a set of TD items, and this technical debt is one of many concerns associated with a system. TD items have both causes and consequences. The cause of technical debt can be a process, a decision, an action (or lack thereof), or an event that triggers the existence of that TD item, such as schedule pressure, unavailability of a key person, or lack of information about a technical feature.

The consequences of a TD item are many: technical debt can effect the value of the system, the costs of future changes, the schedule, and system quality. The business objectives of the sponsoring organization developing or maintaining the software system are affected in several ways: through delays, loss of quality for some features of the system, and difficulties in maintaining the system operations (continuance).

A TD item is associated with one or more concrete, tangible artifacts of the software development process, primarily the code, but also to some extent the documentation, known defects, and tests associated with the system.

To keep with the financial metaphor, the cost impact of technical debt can be seen as composed of principal and interest. The principal is the cost savings gained by taking some initial approach or shortcut in development (the initial principal, often the initial benefit) or the cost that it would now take to develop a different or better solution (the current principal).

The interest is comprised of costs that add up as time passes. There is recurring interest: additional cost incurred by the project in the presence of technical debt, due to reduced velocity (or productivity), induced defects, and loss of quality (maintainability is affected). And there are accruing interests: the additional cost of the developing new software depending on not-quite-right code (evolvability is affected).

This view summarizing the elements related to technical debt, however, does not capture causes that may or may not be input to measurement and analysis, the activities that need to be conducted to manage technical debt, and the states debt may go through. Another view is intended to articulate these aspects.

This definition and the model serve as the starting point for the community to build on and improve.

Research Road Map

One outcome of the seminar was a broad agenda for future work in technical debt research. While this road map needs to be fleshed out in the future with more detailed research questions and problem statements, it lays out three areas that require attention. First is the identification of a core concept – value – that is central to the technical debt metaphor and that needs definition and operationalization. Second is a recognition that there is an important context to technical debt that should be studied. There are attributes of the context of any particular instance of technical

debt in a real environment that must be understood. But there are also other phenomena that are related to technical debt that should be studied, such as other types of "debt." Third, the road map lays out the community's basic infrastructure needs, which will enable further collaboration and progress in this area. The research road map that arose out of the discussions at Dagstuhl is described in more detail in the full report.

Follow-up Work

At the seminar, participants recognized that a carefully considered conceptual model and research road map would be useful outputs for the broader community interested in managing technical debt. Hence, more comprehensive explanation of a conceptual model and the research road map are planned as publications in appropriate venues once the community has a chance to vet the ideas further. The blog established before the seminar will continue to facilitate this interaction.

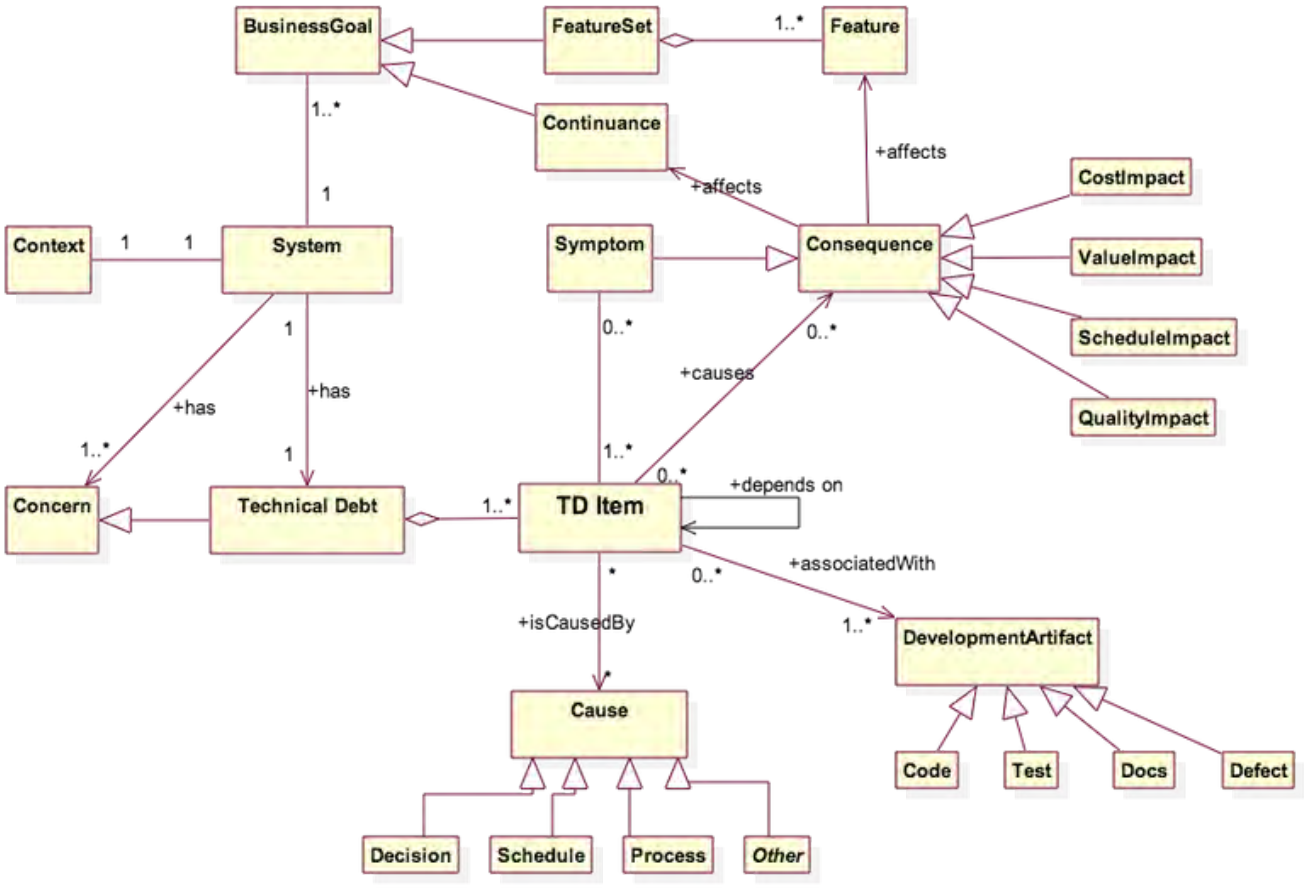

Fig. 6.9
Conceptual Model for Technical Debt.

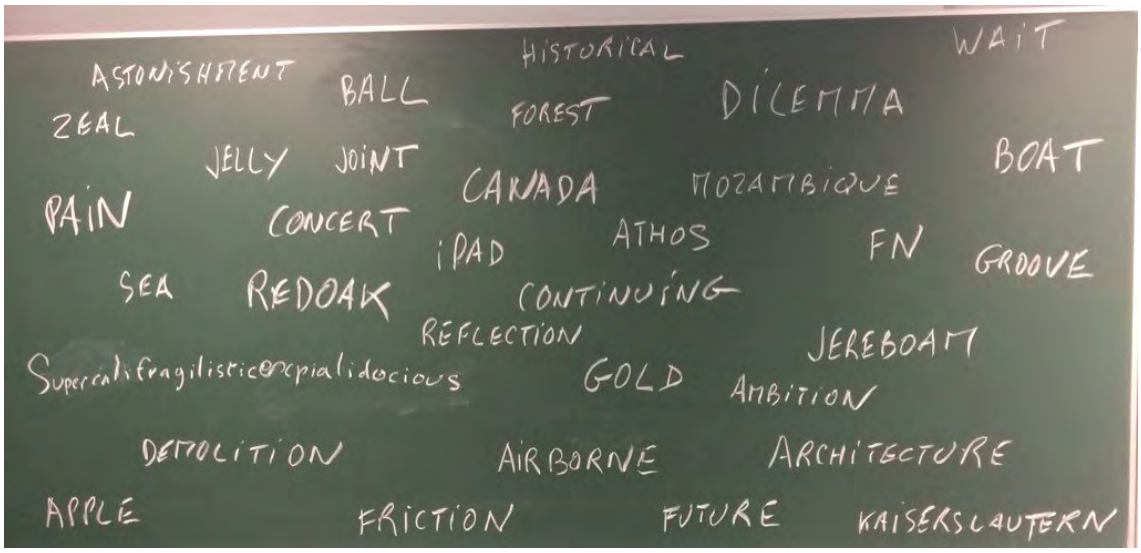

Fig. 6.10
"@dagstuhl workshop on managing technical debt participant clues" Twitter post by 16162 Dagstuhl Seminar participant Ipek Ozkaya. https://twitter.com/ipekozkaya/status/723474431536930816. Photo courtesy of Ipek Ozkaya.

6.28 Algorithmic Methods for Optimization in Public Transport

Organizers: Leo G. Kroon, Anita Schöbel, and Dorothea Wagner

Seminar No. 16171

Date: April 24–29, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.139

Participants: Ralf Borndörfer, Valentina Cacchiani, Francesco Corman, David De Almeida, Markus Friedrich, Marc Goerigk, Rob Goverde, Jonas Harbering, Mark Hickman, Dennis Huisman, Giuseppe F. Italiano, Natalia Kliewer, Leo G. Kroon, Allan Larsen, Jesper Larsen, Marco Laumanns, Janny Leung, Marco Lübbecke, Carlo Mannino, Gabor Maroti, Lingyun Meng, Juan Antonio Mesa, Matus Mihalak, Pitu Mirchandani, Rolf H. Möhring, Matthias Müller-Hannemann, Karl Nachtigall, Morten Nyhave Nielsen, Dario Pacciarelli, Thomas Pajor, Narayan Rangaraj, Marcella Sama, Anita Schöbel, Leena Suhl, Johanna Törnquist Krasemann, Norio Tomii, Lucas Veelenturf, Dorothea Wagner, Christos Zaroliagis

Public transport systems are highly complex systems, due to their technical and organizational complexity, and due to the large numbers of passengers that are transported each day. The quality of the services provided to the passengers is on the one hand the result of the quality and robustness of the underlying plans, such as the timetable and the vehicle and crew schedules. On the other hand, in real-time the quality of the service is the result of the complex interactions between the real-time logistic management of the public transport system and the information to and guidance of the passengers.

Both in the planning stage and in real-time, dealing with these problems requires handling large amounts of data, solving complex combinatorial optimization problems, and dealing with uncertainty. Preferably, the optimization models aim to improve the robustness of the public transport system, so that the system is less vulnerable to disturbances.

In addition, due to the use of smart cards and smart phones, it becomes technically possible to give personalized real-time traffic advice for passengers to guide them to their destinations, even in disturbed situations. In addition, the use of these devices makes huge amount of data available, which can improve decisions in real-time control and in disruption management as well as in the planning stage.

In this seminar, researchers from algorithm engineering and operations research worked together with researchers with an engineering background and participants from practice. The common goal was to improve methods for planning and scheduling of public transportation.

Among others, some specific topic which were covered were

- *Scheduling* of public transport. Several new applications and new ideas on algorithms for public transport scheduling were presented.
- *Integration of planning stages*. Suggestions were developed on how the traditional approach of sequential planning can be replaced by integrated approaches.
- *Robustness and recoverability*. Here several talks discussed methods on how to react to different kinds of disturbances, or how to make schedules more robust.
- *Real-time control*. Real-time control measures which can be taken to get back to the plan as soon as possible were proposed and discussed.
- *Routing in public transport*. For the important issue of routing passengers in public transport, also needed for timetable information systems, several algorithms and new approaches were presented and discussed.
- *Applications and case studies*. Among others the situation in Mumbai, India, was presented and discussed and representatives of several public transport operators sketched the planning process in their companies and pointed out open questions for further research.
- *Future technologies* were another important issue. The participants discussed the potential of new technologies and identified algorithmic challenges for their future utilization.

The seminar started with an introductory round in which every participant presented him- or herself with three slides. It was a good start to get to know each other. In the following days, nearly all participants contributed with talks. There were also two panel discussions, one with the other Dagstuhl group on learning algorithms, and another one on future technologies. The participants discussed and identified challenging algorithmic problems in this field.

Leo Kroon, the main organizer of this Dagstuhl seminar, died unexpectedly on 14 September 2016. We are shocked and very sad about his sudden death. Leo was a great scientist and a wonderful person. We will never forget him.

6.29 Machine Learning for Dynamic Software Analysis: Potentials and Limits

Organizers: Amel Bennaceur, Dimitra Giannakopoulou, Reiner Hähnle, and Karl Meinke

Seminar No. 16172

Date: April 24–27, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.4.161

Participants: Andreas Abel, Dalal Alrajeh, Amel Bennaceur, Pavol Bielik, Radu Grosu, Roland Groz, Reiner Hähnle, Falk Howar, Bengt Jonsson, Karl Meinke, Mohammad Reza Mousavi, Daniel Neider, Zvonimir Rakamaric, Alessandra Russo, Bernhard Steffen, Frits Vaandrager, Sicco Verwer, Neil Walkinshaw, Andrzej Wasowski

Machine learning of software artefacts is an emerging area of interaction between the machine learning (ML) and software analysis (SA) communities. Increased productivity in software engineering hinges on the creation of new adaptive, scalable tools that can analyse large and continuously changing software systems. For example: agile software development using continuous integration and delivery can require new documentation models, static analyses, proofs and tests of millions of lines of code every 24 hours. These needs are being addressed by new SA techniques based on machine learning, such as learning-based software testing, invariant generation or code synthesis.

Machine learning is a powerful paradigm for SA that provides novel approaches to automating the generation of models and other essential artefacts. However, the ML and SA communities are traditionally separate, each with its own agenda. This Dagstuhl Seminar brought together top researchers active in these two fields who can present the state of the art, and suggest new directions and collaborations for future research. We, the organisers, feel strongly that both communities have much to learn from each other, and the seminar focused strongly on fostering a spirit of collaboration.

The first day was dedicated to mutual education through a series of tutorials by leading researchers in both ML and SA to familiarise everyone with the terminology, research methodologies, and main approach of each community. The second day was dedicated to brainstorming and focused discussion in small groups, each of which supported by one of the organisers acting as a facilitator. At the end of the day a plenary session was held for each group to share a summary of their discussions. The participants also reflected and compared their findings. The morning of the third day was dedicated to the integration of the groups and further planning.

Acknowledgements. The organisers would like to express their gratitude to the participants and the Schloss Dagstuhl team for a productive and exciting seminar.

6.30 Fresh Approaches to Business Process Modeling

Organizers: Richard Hull, Agnes Koschmider, Hajo A. Reijers, and William Wong

Seminar No. 16191

Date: May 8–13, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.1

Participants: Banu Aysolmaz, Fernanda Baião, Achim D. Brucker, Artur Caetano, Soren Debois, Marlon Dumas, Kathrin Figl, Avigdor Gal, Jens Gulden, Alexander Herwix, Thomas Hildebrandt, Richard Hull, Christian Janiesch, Udo Kannengiesser, Julius Köpke, Agnes Koschmider, Ralf Laue, Henrik Leopold, Tamara Mchedlidze, Massimo Mecella, Hamid Reza Motahari Nezhad, Andreas Oberweis, Alexander Paar, Jan Recker, Hajo A. Reijers, Flávia Maria Santoro, Tijs Slaats, Chris Snijders, Minseok Song, Christian Stahl, Jianwen Su, Victoria Torres, Han van der Aa, Barbara Weber, Ingo Weber, William Wong, Liang Zhang, Michael zur Muehlen

Business Process Management (BPM) has significantly advanced and gained high popularity in industry. However, it remains an open issue why tools frequently are used for business process modeling that are not mainly implemented for this purpose. Often, macros for Microsoft Visio or Microsoft Excel form the first choice to capture the flow of business activities. One reason why these tools might be used is the low training effort and the fast creation of a quick model, which can be generated with these tools. Another reason for the "lower" preference of BPM software tools might be their inability to respond to changes in technology and working styles, e.g. the shift towards "agile" processes and the "flattening" of workforce hierarchies that bring more stakeholders into contact with a much broader array of processing steps than before.

A central question is whether the BPM community should create an entirely new paradigm for process modeling. One can think of more intuitive drawing conventions that laymen would use, and of models of an entirely different kind (i.e. not process-centric and not data- or case-centric) that still bear the possibility to support modern and future business processes.

The purpose of this seminar was to bring together a cross-disciplinary group of academic and industrial researchers to foster a better understanding of how to ease the access to, and applicability of, business process modeling. We discussed business process modeling approaches against emerging trends such as Internet of Things, the need for incremental and agile creation of new processes, and the need for workers to understand and participate in multiple contextual levels (e.g. transactional, business goals, strategic directions) while performing processes. The seminar also considered how new technologies, such as modern tools for UI design could be applied to support fundamental shifts in how processes are modeled and how humans are involved with their execution.

6.31 Supporting Organizational Efficiency and Agility: Models, Languages and Software Systems

Organizers: Tony Clark, Ulrich Frank, and Vinay Kulkarni

Seminar No. 16192

Date: May 8–13, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.31

Participants: Balbir Barn, Christoph Brandt, Tony Clark, Jan L. G. Dietz, Elmar Dorner, Gregor Engels, Peter Fettke, Hans-Georg Fill, Ulrich Frank, Stijn Hoppenbrouwers, John Krogstie, Vinay Kulkarni, Andreas Leue, Florian Matthes, Andreas L. Opdahl, Henderik Proper, Dirk Riehle, Kurt Sandkuhl, Gerhard Schwabe, Stefan Strecker, Reinhard Wilhelm, Robert Winter, Peter Zencke

Organizations are complex systems that need to respond to a variety of changes while operating in a dynamic environment. They involve multiple stakeholders each having a domain-specific perspective that relies on concepts and languages relative to individual information-centric processes, which may lead to undesirable side-effects such as scattered and fractured knowledge about goals, strategies, operational processes etc.

Organizations are increasingly penetrated by software: Processes and resources are digitized, decision making relies on data provided by software systems, and transactions with external stakeholders are performed by machines. On the one hand, the omnipresence of digital systems creates the opportunity for further automation: The more structures and processes that constitute organizations are represented in software, the greater the scope for computer-supported management. On the other hand, this omnipresence creates a substantial challenge: Many organizations lack the competence to cope with the further increasing complexity of IT infrastructures. This includes the problem of assessing the business impact of IT investment and of assigning IT costs appropriately.

In addition to these problems, organizations face a tremendous challenge: The digital transformation will eliminate many existing business models. It will enable new products and services and it may require organizations to substantially change the way they do business. Only, if organizations are prepared to cope with this challenge, will they be able to benefit from the digital transformation instead of suffering from it.

A key aspect of the digital transformation is automation. While the potential for further automation through software is especially obvious in industrial production, other areas such as administrative work, management, and professional training are more and more dominated by machines. Therefore, there is need for new ways of supporting enterprise agility through the use of integrated computer-based systems

This seminar analyses how organizations can be supported not only with managing their resources and processes efficiently, but also with coping with the digital transformation, a topic which is subject of various research fields including: Management Science (a rationalist perspective); Organisational Studies (including Psychology and Sociology); Information Systems; Software Engineering (including modelling and meta-modelling, big-data and self-adaptive systems); Requirements Engineering. Even though there is an obvious correspondence of foundational assumptions, there is hardly any exchange between these fields: an issue that the seminar aims to address.

Against this background, the seminar is based on the following assumptions:

- Organizations are prepared for change only if they account for the challenges related to adapting their software systems as well as the peculiarities of social change.
- Research on organizational change in general, on designing organizational software systems in particular, recommends not only ideas of how to make organizations more efficient, but of how to make them a better place to work and live in. Otherwise it will be hardly possible to develop advanced conceptions of future organizations that may serve as an orientation for change. Without respective considerations efficiency remains a fairly meaningless concept.
- Support for organizational efficiency and change recommends cross-disciplinary collaboration. While all three research streams outlined above focus on important aspects, none of them is sufficient on its own.
- Support for organisational decision making is currently very difficult due to the tacit nature of knowledge that must be reified and processed using advanced technologies.

6.32 Synergies among Testing, Verification, and Repair for Concurrent Programs

Organizers: Julian Dolby, Orna Grumberg, Peter Müller, and Omer Tripp

Seminar No. 16201

Date: May 16–20, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.56

Participants: Mike Dodds, Julian Dolby, Derek Dreyer, Philippa Gardner, Orna Grumberg, Arie Gurfinkel, Cliff B. Jones, K. Rustan M. Leino, Ben Liblit, Andreas Lochbihler, Peter Müller, Anders Møller, Wytse Oortwijn, Corina Pasareanu, Wolfgang J. Paul, Arnd Poetzsch-Heffter, Murali Krishna Ramanathan, Malavika Samak, Ilya Sergey, Natasha Sharygina, Sharon Shoham Buchbinder, Alexander J. Summers, Michael Tautschnig, Omer Tripp, Caterina Urban, Yakir Vizel, Thomas Wahl

Context and Motivations

Major trends in computing infrastructure, such as multicore processors and data centers, increase the demand for concurrent software that utilizes the available resources. However, concurrent programs are notoriously difficult to develop. They are susceptible to a number of specific errors that do not occur in sequential code, such as data races, deadlock, atomicity violations, starvation, and violations of consistency models. These errors typically manifest themselves only in certain executions (for instance, under certain thread schedules), which makes them extremely difficult to detect, reproduce, localize, and repair. Established techniques for testing, verifying and repairing sequential programs are insufficient to handle concurrent software. In particular, they do not address the following challenges:

- *State space explosion:* The execution of a concurrent program depends not only on the inputs but also on the thread schedule and optimizations, such as memory reordering. This results in an state space that is orders of magnitude larger than for sequential programs. Bug-finding techniques, such as testing and bounded model checking, require effective ways of pruning the state space. Static verification techniques, such as deductive verification and abstract interpretation, require suitable abstractions that allow one to reason about all possible program behaviors. Finally, program repair requires techniques to predict the impact of a program change on the set of possible executions.
- *Modularity:* Modular techniques, such as unit testing or compositional verification, scale to large applications. However, for many properties of concurrent programs there are no modular techniques, or they require a large annotation overhead, for instance to denote the locations protected by a lock or to specify a locking order (or discipline) that ensures deadlock freedom. It is crucial to develop techniques that allow programs to be checked and repaired modularly, for instance to fix an atomicity violation by adding more thread synchronization, but without introducing a deadlock globally.
- *Specifications:* Testing, verification and repair may rely on specifications that express the intended program behavior, for instance in the form of test oracles or program invariants. In addition to functional properties, specifications for concurrent programs also have to express how threads cooperate, for instance via a global locking strategy. While various specification approaches exist for concurrent programs, there is no uniform formalism that handles the full range of concurrency idioms and that supports testing, verification and repair.
- *Error reporting:* Testing, verification and repair techniques need to disambiguate true problems from spurious defects, which is often difficult in concurrent programs. For instance, a data race is not necessarily a bug. If a race occurs within a lock-free data structure, then it may be admissible as part of some higher-level transactional behavior enforced by the data-structure operation. Moreover, it is important to present bugs in an understandable manner, for instance by providing reports with only a small number of threads and by determining whether a bug is inherently concurrent or may also arise in a sequential context.
- *Liveness:* Whereas for most sequential programs, termination is the only relevant liveness property, liveness (such as fairness or the absence of livelocks) is often more prevalent in concurrent programs. It is, therefore, important to develop techniques to check and enforce progress.

Program testing, verification, and repair each offer partial solutions to these challenges. This seminar was conceived with the goal of bringing together these three communities in order to develop a common understanding of the issues as well as to enable collaboration at the level of techniques and tools.

Main Themes

The first step toward exposing, and enabling, synergies between the three main threads of research on correctness and reliability of concurrent programs – verification, testing and repair

– is to analyze the challenges and contributions pertaining to each of these areas in isolation. We survey work that has been done in each of these communities, based on the available literature and presentations given in the seminar, to summarize the current state of the three communities.

Verification. A main challenge in verification of concurrency properties is the prohibitive state space unfolded by thread interleavings. A hybrid solution to this problem is to specialize the static abstraction according to necessary proof conditions, arising during dynamic runs, such that the verification algorithm can scale with fine-grained abstractions (Naik, Yang). Another approach is to retain correlations among local thread states as well as the shared program state (Sagiv, Segalov). In this way, useful invariants can be proved and exploited by the verifier even if an unbounded number of threads is assumed. Refinement techniques are useful when little information is required about the environment to prove a property (Gupta). A useful idea in error reporting is to pinpoint concurrency-specific bugs (differentiating them from sequential bugs) by also running a sequential verifier and performing delta analysis (Joshi). Much like other techniques, verification greatly benefits from user specifications. For example, a parallelizing compiler is more likely to prove disjointness between loop iterations if relevant data structures (or operations) are specified as linearizable (Rinard, Diniz). This also provides a measure of modularity, enabling the separation between library linearizability checking and client verification. Modern program logics (O'Hearn, Parkinson, Gardner) provide a way of constructing correctness proofs for concurrent programs, though in general modular verification of concurrent software remains a hard problem.

Testing. Similarly to verification, testing techniques are also challenged by the state-space problem. Several ideas have been proposed in response to this problem. Open-world testing, whereby data structures or libraries referenced by an application are tested in isolation for concurrency bugs (e.g., atomicity violations), reduces the scope of testing considerably (Shacham). Interestingly, even open-world issues that cannot be recreated within the client application are often fixed by developers, which encourages further research into modular consistency properties (e.g., linearizability) (Shacham). Predictive analysis is a recent form of testing that holds the promise of high coverage at an affordable cost (Smaragdakis). Starting from a concrete trace, predictive analysis applies feasibility-preserving transformations (reordering trace events, typically through constraint solving) to detect concurrency bugs, such that soundness is guaranteed (Dolby, Huang). Another source of state-space reduction is to exploit high-level semantic guarantees, like atomicity, to abstract away intermediate trace transitions (Shacham, Tripp). This also relates to error reporting, where certain read/write conflicts give rise to spurious conflicts that can be eliminated with a higher-level view of conflict as lack of commutativity between atomic operations (Koskinen, Kulkarni). Contrary to memory-level conflict detection, commutativity-based testing requires a specification (Shacham, Tripp). Another form of specification refers to consistency relaxations, e.g. permitting certain types of read/write conflict (Thies) or specifying a computation as nondeterministic (Burnim, Tripp).

Repair. In program repair, error reporting (or localization) plays a key role, deciding the effective scope and nature of the fix. Pinpointing the exact conditions that give rise to a concurrency bug is thus critical, emphasizing the need for better testing and verification tools. Importantly, incorrect fixing may introduce concurrency bugs (e.g., a deadlock resulting from additional synchronization to fix an atomicity violation), which again highlights the need for better synergy between repair and testing/verification (Liu). Incorrect fixing also turns liveness into a concrete concern: Assuming the program previously terminated, does it also terminate after the fix? Existing solutions that ensure termination rely on iterative transformation methods as well as specialized models like Petri nets (Liu, Zhang). A common assumption in the repair community, to hold back the state-space challenge, is that concurrency bugs involve a small number of threads (typically 2) (Liblit, Liu). The hope is that better synergy with testing and verification can work toward relaxing this assumption. Semantic lifting of the concrete code, exploiting e.g. linearizability, has recently been demonstrated as a useful means to apply bottom-up/top-down fixing: First, the code is lifted into an abstract workflow, and then the workflow is concretized into a correct reimplementation (Liu, Tripp). This motivates further exploration of useful specification media for repair of concurrency defects.

Goals of the Seminar

The goal of the seminar was to promote cross fertilization among the verification, testing and repair communities, as they seem to be running into the same challenges, thereby solving increasingly similar problems. At the extreme, verification is about all possible program behaviors, testing is about running the program to see what it does, and repair is about generating new code. However, many techniques in all communities now blur the distinction. Use of dynamic information to guide abstractions in verification is one example; another is how predictive testing looks for bugs in possible executions close to a dynamic one, leading to a form of verification; finally, program repair increasingly uses solvers to synthesize new programs and test them, which overlaps with techniques from the other areas. We intended for the seminar to bring out further areas in which these fields are closely related, and inspire further techniques that fuse these areas, which was fulfilled by some of the discussions throughout the seminar.

Below are concrete examples of connections that we meant to expose, some of which were discussed throughout the seminar:

Benchmarks. Each area has a variety of benchmarks and competitions, and many of them ultimately focus on concurrency-specific challenges like interleavings. It seems likely that the different communities could benefit from sharing. For instance, predictive testing and verification could surely share many benchmarks, and a more standard set of benchmarks could make evaluations easier. At the same time, potential users could help ensure that any benchmarks actually measure what they care about.

Infrastructure. Much progress in both testing and verification has been made possible by progress in solver technology, and a variety of solvers are now common in both areas. There is room to share the infrastructure itself and the common remaining challenges.

Hybrid tools. The path-specific focus of testing and the global focus of verification can aid each other, e.g. current work such as CLAP using a control flow from a specific execution to make model checking more scalable.

Though the seminar touched on techniques and approaches that generalize beyond analysis and repair of concurrent software, we feel that the overall focus on challenges posed by concurrency was justified. With this focus, we were able to stir concrete discussion and tightly connected talks.

6.33 Hardware Security

Organizers: Osnat Keren, Ilia Polian, and Mark M. Tehranipoor

Seminar No. 16202

Date: May 16–20, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.72

Participants: Lejla Batina, Georg T. Becker, Christian Boit, Jan Burchard, Wayne P. Burleson, Jean-Luc Danger, Linus Feiten, Domenic Forte, Fatemeh Ganji, Swaroop Ghosh, Jorge Guajardo Merchan, Tim Erhan Güneysu, Sorin A. Huss, Michael Hutter, Ramesh Karri, Osnat Keren, Tanja Lange, Roel Maes, Nele Mentens, Debdeep Mukhopadhyay, Ruben Niederhagen, Sikhar Patranabis, Elad Peer, Ilia Polian, Wenjing Rao, Francesco Regazzoni, Ulrich Rührmair, Kazuo Sakiyama, Werner Schindler, Georg Sigl, Shahin Tajik, Mark M. Tehranipoor

The convergence of IT systems, data networks (including but not limited to the Internet) and ubiquitous embedded devices within the cyberphysical system paradigm has led to the emergence of new security threats associated with the system hardware. Manipulating the hardware components that implement security functions can compromise system integrity, provide unauthorized access to protected data, and endanger intellectual property. Additionally, secure hardware is required to protect software in a proper manner tampering. Addressing these vulnerabilities is essential in order to prevent the hardware from becoming the Achilles heel of today's systems. Current technology trends point towards massive utilization of hardware circuits in larger cyberphysical systems that are interacting with the physical environment via sensors and actuators. At the same time cyberphysical systems are more and more integrated via open networks, most notably the Internet. Moreover, they interact with each other, forming systems of systems that exhibit highly complex, emergent behavior and constantly change their boundaries, with new sub-systems continuously entering and leaving. As a consequence, hardware-related threats must be addressed by appropriate countermeasures at realistic costs.

The seminar will focus on security threats where hardware components play the main role, and on countermeasures to address these threats. The emphasis is on generic algorithmic advances on the boundary between computer science and other disciplines. While Hardware Security is a very diverse scientific field, the seminar will specifically focus on its three main areas: passive and active side-channel analysis of security-relevant hardware components (cryptographic blocks, true random number generators) which goes beyond classical cryptanalysis; physical unclonable functions (PUFs) and authentication solutions on their basis; and new threats through hardware Trojans and counterfeit ICs as well as techniques for their detection and neutralization.

6

6.34 Algorithms for Optimization Problems in Planar Graphs

Organizers: Jeff Erickson, Philip N. Klein, Dániel Marx, and Claire Mathieu

Seminar No. 16221

Date: May 29 to June 3, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.94

Participants: Anna Adamaszek, Mohammad Hossein Bateni, Hans L. Bodlaender, Sergio Cabello, Vincent Cohen-Addad, Zdenek Dvorak, Jeff Erickson, Thomas Erlebach, Fedor V. Fomin, Kyle Jordan Fox, Eli Fox-Epstein, Michelangelo Grigni, Bart Jansen, Philip N. Klein, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Claire Mathieu, Bojan Mohar, Shay Mozes, Marcin Pilipczuk, Michal Pilipczuk, Peter Rossmanith, Piotr Sankowski, Ignasi Sau Valls, Saket Saurabh, Aaron Schild, Andreas Schmid, Anastasios Sidiropoulos, Dimitrios M. Thilikos, Tom van der Zanden, Erik Jan van Leeuwen, Oren Weimann, Andreas Wiese, Christian Wulff-Nilsen, Hang Zhou, Anna Zych

There is a long tradition of research in algorithms for optimization problems in graphs, including work on many classical problems, both polynomial-time solvable problems and NP-hard problems, e.g. shortest paths, maximum flow and minimum cut, matching, T-joins, disjoint paths, traveling salesman, Steiner tree, graph bisection, vehicle routing, facility location, k-center, and maximum cut. One theme of such research addresses the complexity of these problems when the input graph is required to be a planar graph or a graph embedded on a low-genus surface.

There are three reasons for this theme. First, optimization problems in planar graphs arise in diverse application areas. Second, researchers have discovered that, by exploiting the planarity of the input, much more effective algorithms can be developed – algorithms that are faster or more accurate than those that do not exploit graph structure. Third, the study of algorithms for surface-embedded graphs drives the development of interesting algorithmic techniques. One source of applications for planar-graph algorithms is geographic problems. Road maps are nearly planar, for example, so distances in planar graphs can model, e.g., travel times in road maps. Network design in planar graphs can be used to model scenarios in which cables must be run under roads. Planar graphs can also be used to model metrics on the earth's surface that reflect physical features such as terrain; this aspect of planar graphs has been used in studying wildlife corridors. Another source of applications is image processing. Some algorithms for problems such as image segmentation and stereo involve finding minimum cuts in a grid in which each vertex represents a pixel. Sometimes an aggregation technique (superpixels) coalesces regions into vertices, turning the grid into an arbitrary planar graph. A third example application is VLSI. Algorithmic exploitation of a planar embedding goes back at least to the introduction of maximum flow by Ford and Fulkerson in 1956. Current research can be divided in three parts. For polynomial-time-solvable problems, such as maximum flow, shortest paths, matching, and min-cost circulation, researchers seek planarity-exploiting algorithms whose running times beat those of general-graph algorithms, ideally algorithms whose running times are linear or nearly linear. For NP-hard problems, there are two strategies: fixed-parameter algorithms and approximation algorithms. In all three research subareas, there has recently been significant progress. However, many researchers are expert in only one or two subareas. This Dagstuhl Seminar brought together researchers from the different subareas, to introduce them to techniques from subareas that might be unfamiliar, and to foster collaboration across the subareas. The seminar will thus help to spur further advances in this active and growing area. The scientific program of the seminar consisted of twenty-two talks. Four of these talks were longer (60–90 minute) tutorials overviewing the three main areas of the seminar:

- *Polynomial-time algorithms:* "Tutorial on embedded graph algorithms" (Jeff Erickson) and "Monge property, dense distance graphs and speeding-up max-flow computations in planar graphs" (Piotr Sankowski)
- *Approximation schemes:* "Some techniques for approximation schemes on planar graphs" (Philip Klein)
- *Fixed-parameter tractability:* "The square-root phenomenon in planar graphs" (Dániel Marx)

One of the main goals of the seminar was to encourage collaboration between the three communities, and these well-received tutorials helped by introducing the basics of each of these topics.

The rest of the talks were 25-minute presentations on recent research of the participants. The time between lunch and the afternoon coffee break was left open for individual discussions and collaborations in small groups. An open-problem session was organized on Monday morning. Notes on the presented problems can be found in this report.

6.35 Engineering Moral Agents – from Human Morality to Artificial Morality

Organizers: Michael Fisher, Christian List, Marija Slavkovik, Alan Winfield

Seminar No. 16222

Date: May 29 to June 3, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.5.114

Participants: Michael Anderson, Albert Anglberger, Zohreh Baniasadi, Kevin Baum, Vincent Berenz, Jan M. Broersen, Vicky Charisi, Louise A. Dennis, Sjur K. Dyrkolbotn, Michael Fisher, Joseph Halpern, Holger Hermanns, Johannes Himmelreich, John F. Horty, Susan Leigh Anderson, Robert Lieck, Christian List, Andreas Matthias, James H. Moor, Marcus Pivato, Marek Sergot, Marija Slavkovik, Janina Sombetzki, Kai Spiekermann, Alan FT Winfield, Roman V. Yampolskiy

Artificial morality, also called "machine ethics", is an emerging field in artificial intelligence that explores how artificial agents can be enhanced with sensitivity to and respect for the legal, social, and ethical norms of human society. This field is also concerned with the possibility and necessity of transferring the responsibility for the decisions and actions of the artificial agents from their designers onto the agents themselves. Additional challenging tasks include, but are not limited to: the identification of (un)desired ethical behaviour in artificial agents and its adjustment; the certification and verification of the artificial agents' ethical capacities; the identification of the adequate level of responsibility of an artificial agent; the dependence between the responsibility and the level of autonomy that an artificial agent possesses; and the place of artificial agents within our societal, legal, and ethical normative systems.

Artificial morality has become increasingly salient since the early years of this century, though its origins are older. Isaac Asimov already famously proposed three laws of robotics, requiring that, first, robots must not harm humans or allow them to be harmed; second, robots must obey human orders provided this does not conflict with the first law; and third, robots must protect themselves provided this does not conflict with the first two laws.

Although there has been some discussion and analysis of possible approaches to artificial morality in computer science and related fields, the "algorithmization" and adaptation of the ethical systems developed for human beings is both an open research problem and a difficult engineering challenge. At the same time, formally and mathematically oriented approaches to ethics are attracting the interest of an increasing number of researchers, including in philosophy. As this is still in its infancy, we thought that the area could benefit from an "incubator event" such as an interdisciplinary Dagstuhl seminar.

We conducted a five-day seminar with twenty six participants with diverse academic backgrounds including robotics, automated systems, philosophy, law, security, and political science. The first part of the seminar was dedicated to facilitating the cross-disciplinary communication by giving researchers across the contributing disciplines an integrated overview of current research in machine morality from the artificial intelligence side, and of relevant areas of philosophy from the moral-philosophy, action-theoretic, and social-scientific side. We accomplished this through tutorials and brief self-introductory talks. The second part of the seminar was dedicated to discussions around two key topics: how to formalise ethical theories and reasoning, and how to implement ethical reasoning. This report summarises some of the highlights of those discussions and includes the abstracts of the tutorials and some of the self-introductory talks. We also summarise our conclusions and observations from the seminar.

Although scientists without a philosophical background tend to have a general view of moral philosophy, a formal background and ability to pinpoint key advancements and central work in it cannot be taken for granted. Kevin Baum from the University of Saarland presented a project currently in progress at his university and in which he is involved, of teaching formal ethics to computer-science students. There was great interest in the material of that course from the computer science participants of the seminar. In the first instance, a good catalyst for the computer science–moral philosophy cooperation would be a comprehensive "data base" of moral-dilemma examples from the literature that can be used as benchmarks when formalising and implementing moral reasoning.

The formalisation of moral theories for the purpose of using them as a base for implementing moral reasoning in machines, and artificial autonomous entities in general, was met with great enthusiasm among non-computer scientists. Such work gives a unique opportunity to test the robustness of moral theories.

It is generally recognised that there exist two core approaches to artificial morality: explicitly constraining the potentially

immoral actions of the AI system; and training the AI system to recognise and resolve morally challenging situations and actions. The first, constrained-based approach consists in finding a set of rules and guidelines that the artificial intentional entity has to follow, or that we can use to pre-check and constrain its actions. By contrast, training approaches consist in applying techniques such as machine learning to "teach" an artificial intentional entity to recognise morally problematic situations and to resolve conflicts, much as people are educated by their carers and community to become moral agents. Hybrid approaches combining both methods were also considered.

It emerged that a clear advantage of constraining the potentially immoral actions of the entity, or the "symbolic approach" to ethical reasoning, is the possibility to use formal verification to test that the reasoning works as intended. If the learning approach is used, the learning should happen before the autonomous system is deployed for its moral behaviour to be tested. Unfortunately, the machine-learning community was severely under-represented at the seminar, and more efforts should be devoted to include them in future discussions. The discussions also revealed that implanting moral reasoning into autonomous systems opens up many questions regarding the level of assurance that should be given to users of such systems, as well as the level of transparency into the moral-reasoning software that should be given to users, regulators, governments, and so on.

Machine ethics is a topic that will continue to develop in the coming years, particularly with many industries preparing to launch autonomous systems into our societies in the next five years. It is essential to continue open cross-disciplinary discussions to make sure that the machine reasoning implemented in those machines is designed by experts who have a deep understanding of the topic, rather than by individual companies without the input of such experts. It was our impression as organisers, perhaps immodest, that the seminar advanced the field of machine ethics and opened new communication channels. Therefore we hope to propose a second seminar in 2018 on the same topic, using the experience and lessons we gained here, to continue the discussion and flow of cross-disciplinary collaboration.

6.36 Immersive Analytics

Organizers: Tim Dwyer, Nathalie Henry Riche, Karsten Klein, Wolfgang Stuerzlinger, and Bruce Thomas

Seminar No. 16231

Date: June 5–10, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.1

Participants: Benjamin Bach, Maxime Cordeil, Raimund Dachselt, Pierre Dragicevic, Steven M. Drucker, Tim Dwyer, Niklas Elmqvist, Barrett Ens, Roland Fernandez, Carla Freitas, Carsten Görg, Mark Hancock, Hans-Christian Hege, Nathalie Henry Riche, Christophe Hurter, Pourang P. Irani, Petra Isenberg, Tobias Isenberg, Takayuki Itoh, Yvonne Jansen, Karsten Klein, Bongshin Lee, Todd Margolis, Kim Marriott, Jon McCormack, Chris North, Steffen Oeltze-Jafra, Huamin Qu, Aaron Quigley, Jonathan C. Roberts, Dieter Schmalstieg, Falk Schreiber, Wolfgang Stuerzlinger, Aurélien Tabard, Bruce Thomas, Frank van Ham, Gregory F. Welch, Uwe Wössner

Immersive Analytics is an emerging new field that studies technologies facilitating a deep cognitive, perceptual and/or emotional involvement of humans when understanding and reasoning with data.

Immersive technologies are commonly defined as technologies aiming at blurring the line between physical and virtual worlds, by employing multimodal input and multi-sensory output to create a state of immersion, i.e. a deep mental involvement of a person into an activity and/or an intense concentration or complete absorption into the activity that one does.

The term Immersive Analytics was coined a few years ago, but there is no precise definition of the concept so far, and the corresponding research is scattered across several fields and communities. Hence our goal for this seminar was to discuss and define the field of Immersive Analytics, and to create a community around it. In addition, we planned to develop an outline for a book on the topic.

During the working group and discussion sessions, the participants investigated the potential and the challenges of immersive analytics for research and commercial applications, as well as a variety of aspects like multi-sensory data representation, immersive human-centered data analysis, interaction for immersive analysis, immersion for data-driven narratives, and the use of immersive analytics concepts in application areas like the life sciences and air traffic control.

During the first plenary sessions, major topics for discussion were defined and clustered into working groups, and the participants then joined the proposed working groups based on common interest. Later, the participants could switch between the groups. Each of the working groups was meant to outline a chapter of the book publication. For some of the topics, discussions continued in the evening hours, which were also used to experience new technologies like the Microsoft HoloLens.

6.37 Fair Division

Organizers: Yonatan Aumann, Jérôme Lang, and Ariel D. Procaccia

Seminar No. 16232

Date: June 5–10, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.10

Participants: Yonatan Aumann, Haris Aziz, Sylvain Bouveret, Simina Brânzei, Katarina Cechlarova, Yann Chevaleyre, Marco Dall'Aglio, Edith Elkind, Ulle Endriss, Serge Gaspers, Vasilis Gkatzelis, Laurent Gourves, Avinatan Hassidim, Dorothea Herreiner, Mehmet Ismail, Michael A. Jones, Marc Kilgour, Christian Klamler, David Kurokawa, Jérôme Lang, Simon William Mackenzie, Hervé J. Moulin, Nhan-Tam Nguyen, Trung Thanh Nguyen, Ariel D. Procaccia, Clemens Puppe, Jörg Rothe, Uzi Segal, Erel Segal-Halevi, Walter Stromquist, Taiki Todo, Toby Walsh

Fair division has been an active field of research in economics and mathematics for decades. More recently, the topic has attracted the attention of computer scientists, due to its algorithmic nature and its real-world applications. There had been a first Dagstuhl Seminar on fair division, in 2007, and none since. The aim of the 2016 Dagstuhl seminar on fair division was to bring together top researchers in the field, from among the multiple disparate disciplines where it is studied, both within computer science and from economics and mathematics, to share knowledge and advance the state of the art.

The seminar covered fair division of both divisible and indivisible goods, with a good mix between economics and computer science (with a significant number of talks being about economics *and* computer science). Topics included algorithms, lower bounds, approximations, strategic behavior, tradeoffs between fairness and efficiency, partial divisions, alternative definitions of fairness, and practical applications of fair division. The ratio between the number of participants with a main background in computer science and in economics was about 3–1, with a couple of participants with another main background (mathematics or political science). This ratio is similar to the corresponding ratios for Dagstuhl seminars on computational social choice (2007, 2010, 2012, 2015).

The seminar started by a short presentation of the participants (3 minutes per attendee). The rest of the seminar was composed of technical sessions with regular talks, and discussion sessions distributed over the full week (Tuesday morning, Tuesday afternoon, Wednesday morning, Friday morning). One of these discussion sessions was specifically about *Fair division in the real world*, two were about open problems, and one was about high-level thoughts about the topic and its future. Moreover, there was a significant amount of time left for participants to interact in small groups.

6.38 Graph Polynomials: Towards a Comparative Theory

Organizers: Jo Ellis-Monaghan, Andrew Goodall, Johann A. Makowsky, and Iain Moffatt

Seminar No. 16241

Date: June 12–17, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.26

Participants: Ilia Averbouch, Spencer Backman, Markus Bläser, Béla Bollobás, Robert Brijder, Ada Sze Sze Chan, Carolyn Chun, Bruno Courcelle, Anna De Mier, Holger Dell, Jo Ellis-Monaghan, Graham Farr, Alex Fink, Emeric Gioan, Andrew Goodall, Gary P. Gordon, Krystal Guo, Hendrik Jan Hoogeboom, Thore Husfeldt, Natasa Jonoska, Louis H. Kauffman, Marsha Kleinbauer, Tomer Kotek, Joseph Kung, Martin Loebl, Johann A. Makowsky, Liz McMahon, Iain Moffatt, Kerri Morgan, Ada Morse, Jaroslav Nesetril, Steven Noble, Marc Noy, Seongmin Ok, James Oxley, Thomas Perrett, Vsevolod Rakita, Elena V. Ravve, Guus Regts, Gordon Royle, Benjamin Smith, Peter Tittmann, Martin Trinks, William Whistler

The intent of this 5-day Seminar was to develop a general theory of graph polynomials. Graph polynomials have played a key role in combinatorics and its applications, having effected breakthroughs in conceptual understanding and brought together different strands of scientific thought. The characteristic and matching polynomials advanced graph-theoretical techniques in chemistry; the Tutte polynomial married combinatorics and statistical physics, and helped resolve long-standing problems in knot theory. The area of graph polynomials is incredibly active, with new applications and new graph polynomials being discovered each year. However, the resulting plethora of techniques and results now urgently requires synthesis. Beyond catalogues and classifications we need a comparative theory.

There is a long history in this area of results in one field leading to breakthroughs in another when techniques are transferred, and this Seminar leveraged that paradigm. More critically, experts in the field have recently begun noticing strong resonances in both results and proof techniques among the various polynomials. The species and genera of graph polynomials are diverse, but there are strong interconnections: the Seminar initiated work on creating a general theory that will bring them together under one family. The process of developing such a theory of graph polynomials should expose deeper connections, giving great impetus to both theory and applications. This has immense and exciting potential for all those fields of science where combinatorial information needs to be extracted and interpreted.

The Seminar provided conditions ripe for cross-fertilization of ideas among researchers in graph theory and topological graph theory, in logic and finite model theory, and in current biocomputing and statistical mechanics applications. During the Seminar the participants were offered a conspectus of the broad area of graph polynomials. The view was confirmed that a synthetic approach is needed in order to see the wood for the trees. The discussions and collaborations initiated at the workshop promise well for the development of a unified theory of graph polynomials. This Seminar represented a convincing beginning, and, hopefully, similar meetings in future will further the envisaged project.

In the light of our stated goals, the Seminar provided ample time for discussion groups and tutorials. The participants (44) of the Seminar included some of the leading experts in combinatorics, knot theory, matroid theory and graph polynomials from Europe, the Americas, Asia and Australia. The composition of participants was both age and gender balanced with a quarter of the participants being women. The younger researchers (more than a quarter of the participants) profited from intense contacts and discussions with their more experienced colleagues. An inspiring problem session brought about particular directions for further research.

Fig. 6.11
"@dagstuhl #ImmersiveDagstuhl [...] All participants working hard. Discussing a book outcome. #havingfun" Twitter post by 16231 Dagstuhl Seminar participant Jonathan C. Roberts. https://twitter.com/jcrbrts/status/740483240209842176. Photo courtesy of Jonathan C. Roberts.

6.39 Information-centric Networking and Security

Organizers: Edith Ngai, Börje Ohlman, Gene Tsudik, and Ersin Uzun

Seminar No. 16251

Date: June 19–22, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.49

Participants: Bengt Ahlgren, Tohru Asami, Roland Bless, Randy Bush, Kenneth L. Calvert, Antonio Carzaniga, Mauro Conti, Lars Eggert, Darleen Fisher, Ashish Gehani, Jussi Kangasharju, Ghassan Karame, Dirk Kutscher, John Mattsson, Marc Mosko, Edith Ngai, Börje Ohlman, Jörg Ott, Craig Partridge, Fabio Pianese, Sanjiva Prasad, Thomas C. Schmidt, Sebastian Schönberg, Christoph Schuba, Glenn Scott, Jan Seedorf, Tim Strayer, Christian Tschudin, Gene Tsudik, Ersin Uzun, Matthias Wählisch, Cedric Westphal, Christopher A. Wood

Dagstuhl seminar 16251 "Information-centric Networking and Security" was a short workshop held June 19–21, 2016. The goal was to bring together researchers with different areas of expertise relevant to ICN to discuss security and privacy issues particular to ICN-based architectures. These problems have become increasingly important as ICN technology gradually matures and nears real-world deployment.

Threat models are distinct from IP. Differentiating factors between the two include new application design patterns, trust models and management, as well as a strong emphasis on object-based, instead of channel-based, security. Therefore, it is both timely and important to explore ICN security and privacy issues as well as devise and assess possible mitigation techniques. This was the general purpose of the Dagstuhl seminar. To that end, the attendees focused on the following issues:

- What are the relevant threat models with which ICN must be concerned? How are they different from those in IP-based networks?
- To what extent is trust management a solved problem in ICN? Have we adequately identified the core elements of a trust model, e.g., with NDN trust schemas?
- How practical and realistic is object-based security when framed in the context of accepted privacy measures used in IP-based networks?
- Are there new types of cryptographic schemes or primitives ICN architectures should be using or following that will enable (a) more efficient or secure packet processing or (b) an improved security architecture?

The seminar answered (entirely or partially) some of these questions and fueled discussions for others. To begin, all participants briefly introduced themselves. This was followed by several talks on various topics, ranging from trust management and identity to privacy and anonymity. Subsequently, the attendees split into working groups to focus more intensely on specific topics. Working group topics included routing on encrypted names, ICN and IoT, non-privacy-centric aspects of ICN security, as well as trust and identity in ICN. Once the working group sessions were over, a representative from each presented outcomes to all attendees. (These are documented in the remainder of this report.) The major takeaways from the seminar were as follows.

First, the ICN community still does not have a clear answer for how to handle namespace and identity management. While trust management in ICN can be distributed and function without a global PKI, it seems difficult to break away from this model for namespace management and arbitration. This has strong implications on how names are propagated in the routing fabric. Can any producer application advertise any name, anywhere in the network? If not, how can name prefix advertisements be constrained or limited?

Second, given that ICN focuses on object security, the need for and use of transport protocols that provide forward secrecy should be deferred to higher layers. Attendees found that while most ICN-based architectures do not preclude forward secrecy, it should not be a requirement at the network layer.

Third, there is still deep uncertainty about whether ICN should embrace a content locator and identifier split. Names in architectures such as NDN and CCN serve as both a locator and identifier of data, though there are extensions that permit explicit locators (e.g., through the use of NDN LINK objects). This distinction is necessary under the common understanding that routing should concern itself with topological names. Finding data through non-topological names should not be in the data plane as part of the global routing space. However, if we revert to a distinction between topological locators and identifiers, then features unique to ICN become much more limited. One facet that is certainly unique to ICN is how software is written. Specifically, we have the opportunity to move beyond the mental model of

a fixed address space and re-design existing network stacks and APIs.

Fourth, privacy seems difficult to achieve without major architectural changes to ICN-based systems. In particular, since data names reveal a great deal of information to the passive eavesdropper, privacy demands that names and payloads have no correlation. However, achieving this seems infeasible without the presence of an upper-layer service akin to one that would resolve non-topological identifiers to topological names.

Lastly, there are no compelling reasons to apply esoteric (and often untested) cryptographic techniques in ICN, at least at the network layer. Computationally bounded and "boring" cryptographic primitives, such as digital signatures, hash functions, etc., should be the extent of per-packet cryptographic processing done by routers. Anything more would become fodder for Denial-of-Service attacks that could render the entire infrastructure ineffective. However, architecture designs should not restrict themselves to specific algorithms. In other words, there must be flexibility in accommodating multiple (and evolving) cryptographic primitives. This could be useful if, for example, post-quantum digital signature schemes become necessary for the longevity of content authenticators.

We thank Schloss Dagstuhl for providing a stimulating setting for this seminar. Much progress was made over the course of the seminar and since its completion. This is mainly because of the ease of face-to-face collaboration and interaction at Dagstuhl.

6.40 Engineering Academic Software

Organizers: Carole Goble, James Howison, Claude Kirchner, Oscar Nierstrasz, and Jurgen J. Vinju

Seminar No. 16252

Date: June 20–24, 2016 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.5362/DagRep.6.6.62

Participants: Alice Allen, Cecilia Aragon, Christoph Becker, Jeffrey Carver, Andrei Chis, Benoit Combemale, Mike Croucher, Kevin Crowston, Daniel Garijo, Ashish Gehani, Carole Goble, Robert Haines, Robert Hirschfeld, James Howison, Katy Huff, Caroline Jay, Daniel S. Katz, Claude Kirchner, Katie Kuksenok, Ralf Lämmel, Oscar M. Nierstrasz, Matthew J. Turk, Rob van Nieuwpoort, Matthew Vaughn, Jurgen J. Vinju

This Dagstuhl Perspectives Workshop brought together activists, experts and stakeholders on the subject of high quality software produced in an academic context.[38] Our current dependence on software across the sciences is already significant, yet there are still more opportunities to be explored and risks to be overcome. The academic context is unique in terms of its personnel, its goals of exploring the unknown and its demands on quality assurance and reproducibility.

We refer to the IEEE Internet Computing article "Better Software, Better Research" [1] which motivated the topic. In this workshop we took the following perspective of a research team which is in either or both of the following situations:

- consuming or producing software as *an output* of the academic process;
- consuming or producing software as *a component* of the research methods.

Society is now in the tricky situation where several deeply established academic fields (e.g. physics, biology, mathematics) are shifting towards dependence on software, programming technology and software engineering methodology which are backed only by young and rapidly evolving fields of research (computer science and software engineering). Full accountability and even validity of software-based research results are now duly being challenged.

With the outputs of this interactive and productive perspectives workshop, we strive to contribute in a positive manner to the above challenges. We formulated taxonomies with definitions to clarify the domain, we co-authored concrete policy and process documents to improve the status and recognition of academic software development and academic software engineers, and finally we formulated a list of 18 concrete declarations of intent ("I will" pledges). This list was presented to the WSSSPE community [2] in September 2016 to acquire feedback and it will be the backbone of the Dagstuhl Manifesto document we are editing. It serves to motivate change by proposing policy changes with concrete actions and instilling positive attitudes towards academic software.

Participants. The participants of the workshop came from three major groups. The first group consists of active and visible members of the *global academic software engineering community*. They represent (formal) institutions such as the Software Sustainability Institute, the Software Carpentry Foundation, and eScience and data science centers from across the globe. The second group contributed researchers in *empirical software engineering*, with a specific eye on studying the principles and practices of academic software engineering. The final group contributed *researchers as an audience*: software engineering researchers with a long experience in engineering software for software itself or software for specific academic research fields.

We found that without exception the participants were strongly motivated and able to actively contribute to the proceedings of the workshop; the mix of people proved to be well-balanced. This balance is an accomplishment, given that invitees from computer science were far more likely to know of Dagstuhl workshops than other groups. To attest to our outcomes we've selectively listed three (paraphrased) verbal statements here:

- "The workshop was a transformational experience for me; I've learned an entire new perspective on my field and I intend to apply the insights in my daily practice."

[38] We include any software which is part of either research processes and/or output, while excluding more generic administrative software for research and education management.

- "I had an epiphany yesterday after dinner; now I understand how to connect the data science research at my university to the computer science department."
- "Before the workshop I had no idea so many initiatives were already underway in [improving] academic software engineering; this has given my understanding of the challenges a real boost and I know what the some of the next steps to take are."

Schedule. The schedule of the workshop was designed to maximize both interactive discussion and work towards tangible outputs. Key points were: to start the day with inspiring presentations to set the stage, then to have at least 40% of the day time allocated to free discussion time, and to explicitly share successes (output) of each day's breakout groups in a plenary session.

The workshop started on Monday with a quick and tightly timed round of 2 minute personal introductions. Otherwise on Monday, Tuesday and Thursday the program was structured equally: in the morning we would have plenary presentations which included exploratory discussions. These sessions were meant to bring everybody up-to-speed with ongoing and past initiatives. During and after lunch we used a board with sticky notes to define break-out groups. Each break-out group was centred around a specific discussion topic and (usually) a specific idea for an output document was associated with it. After coffee we would go back to the same break-out group to collaboratively record the notes and lessons from each group (stored in a shared online document). Between 17:00 and 18:00 we reconvened and harvested the results of each breakout group with the others. People could and did freely switch between breakout groups but this was not a common thing.

On Wednesday we had an "open-mic" session with 8 presentations of around 10 minutes, sharing experiences and results, before we had a long walk in the surroundings. The organizers also designed an initial skeleton structure and ideas for the manifesto that day.

On Thursday afternoon and Friday morning we all worked together on our Dagstuhl Manifesto by first reworking our notes into the ideas around the manifesto, specifically a list of "I will" pledges with references and motivation. Finally, Friday afternoon a small remaining group re-ordered the group's manifesto notes into a well-structured list of 18 pledges. Two of the organizers remained to continue to edit the current report and the manifesto document.

Output. Output documents of the workshop are organized under the "DagstuhlEAS" organisation on GitHub.[39] This currently features 6 draft documents, including the current report and (a) the manifesto, (b) the Research Software Engineering Handbook, (c) a Literature Survey, (d) a Taxonomy on Software Credit Roles, and (e) a Software Award Proposal. Next to these documents, an R&D project proposal was produced on measuring the impact of academic software.

The remainder of this document summarizes the morning sessions by listing the abstracts of each talk, the afternoon breakouts by describing each topic and its results, and finally the research questions on the topic of engineering academic software we have collected.

References

1 Carole Goble. Better software, better research. *IEEE Internet Computing*, 18(5):4–8, Sep 2014.

2 Alice Allen, Cecilia Aragon, Christophe Becker, Jeffrey C. Carver, Andrei Chis, Benoit Combemale, Mike Croucher, Kevin Crowston, Daniel Garijo, Ashish Gehani, Carole Goble, Robert Haines, Robert Hirschfeld, James Howison, Kathryn Huff, Caroline Jay, Daniel S. Katz, Claude Kirchner, Kateryna Kuksenok, Ralf Lämmel, Oscar Nierstrasz, Matthew Turk, Rob van Nieuwpoort, Matthew Vaughn, and Jurgen Vinju. Lightning talk: "I solemnly pledge" – a manifesto for personal responsibility in the engineering of academic software. In *Proceedings of the Fourth Workshop on Sustainable Software for Science: Practice and Experiences (WSSSPE4)*.

[39] https://github.com/DagstuhlEAS

6.41 Integration of Expert Knowledge for Interpretable Models in Biomedical Data Analysis

Organizers: Gyan Bhanot, Michael Biehl, Thomas Villmann, and Dietlind Zühlke

Seminar No. 16261

Date: June 26 to July 1, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.88

Participants: Gyan Bhanot, Michael Biehl, Kerstin Bunte, Aguirre de Cubas, Gert-Jan de Vries, Sebastian Doniach, Shridar Ganesan, Tina Geweniger, Gernoth Grunst, Barbara Hammer, Marika Kaden, Hossein Khiabanian, Saurabh V. Laddha, John A. Lee, Pietro Lio', Paulo J. Lisboa, Markus Lux, Elke K. Markert, John Martens, Thomas Martinetz, Friedrich Melchert, Erzsébet Merényi, Klaus Mueller, David Nebel, Jeffrey Rathmell, W. Kimryn Rathmell, Anupama Reddy, Timo Ropinski, Joshua T. Taylor, Peter Tino, Alexei Vazquez, Thomas Villmann, Gunther H. Weber, Ole Winther, Thomas Wischgoll, Röbbe Wünschiers, Dietlind Zühlke

The participants were drawn from three distinct disciplines: Biomedical Research, Machine Learning and Visualizations. On the first day, three overview talks on different aspects of bio-medical research were presented, including an overview of omics and clinical data and databases, a summary of current problems in cancer prognosis and metastasis, and steroid metabolomics and its relevance to disease. On the next two days, there were four overview talks on computer science topics, including machine learning, modeling and visualization. Participants also had the opportunity to give shorter presentations of their current research areas and describe open problems, as well as introduce new and relevant datasets and methods. In total, 16 such short talks were presented, covering various areas of biomedical research and computer science. All talks served as starting points for extensive plenary and individual evening and after dinner discussions about the integration of expert knowledge into data analysis and modeling, specifically targeted to cancer informatics. From these discussions, it was clear that there was an urgent need for interactive collaboration to foster successful analysis and interpretation of biomedical data and the success of such collaboration would hinge on active participation from domain experts from biomedical research, data mining and visualization.

Motivated by this conclusion, we identified a joint project in cancer genomics, which would exploit the expertise represented by the seminar participants. On the fourth day, participants discussed the interactive methodology we will follow in the project. Following this, first results obtained by analysis of cancer data from The Cancer Genome Atlas was presented in a joint talk by representatives from all three disciplines (biology, machine learning, visualization). We will extend this project further in the coming months with active participation from the clinicians and computer scientists. The goal of this effort is not just to solve a relevant and outstanding problem in cancer biology but also to work towards publication of our findings in a high-impact journal authored by all participants. To foster this project, we will establish a Wiki, which will serve as a platform for collaboration and communication.

The participants gave feedback on Friday on the organization and content of the seminar. All participants were appreciative of the open, friendly and constructive atmosphere that made learning and insight possible for experts from very diverse disciplines. Getting to know the basic methods used in each field was seen as the perfect starting point for future collaborations. The idea of a joined wiki page as a collaboration platform as well as the already started joined project were highlighted as especially important. Follow-up-meetings of newly formed interdisciplinary teams were initiated and planned e.g. one in Copenhagen. The participants were very enthusiastic about having a further meeting after about a year to discuss results and new directions resulting from the joint project initiated here. Apart from working on a specific project in cancer biology, the goal of the collaboration is to establish a methodology for interactions, disseminate ideas and protocols among the disciplines and establish a common language to foster understanding.

In summary, biologists, both medical and computational experts in the seminar are enthusiastic about joining forces to solve outstanding problems in understanding biological processes. Many of the machine learning methods presented by participants are ready to be applied in real environments such as in clinical use or in research laboratories, after proper technology transfer. Such technology transfer requires targeted funding and agreed upon protocols to ensure adequate resources and necessary quality control, for subsequent release to the community.

The participants felt that influential members in each community should seek opportunities and avenues to urge the appropriate agencies (NIH, NFS, EU Scientific bodies) to establish a targeted program for technology transfer of computational solutions to challenges in the interpretation of biomedical data. Such a program would solicit competitive funding proposals from groups consisting of both biomedical and computational experts, and require products that are rigorously demonstrated

on real problems, as well as satisfy appropriate coding and user interface standards, and where appropriate, satisfy requirements of interfacing or integration with existing established systems currently in use by the community.

In medicine the data is treasure
Whose value's beyond any measure
But it is not surprising
That without analysing
Acquisition is meaningless pleasure

(Michael Biehl and Gyan Bhanot)

6.42 Automotive User Interfaces in the Age of Automation

Organizers: Andreas Riener, Susanne Boll, and Andrew L. Kun

Seminar No. 16262

Date: June 26 to July 1, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.111

Participants: Ignacio Alvarez, Martin Baumann, Susanne Boll, Linda Boyle, Nora Broy, Duncan Brumby, Gary Burnett, Lewis Chuang, Frank Flemisch, Anna-Katharina Frison, Christian P. Janssen, Jeon Myounghoon, Wendy Ju, Andrew Kun, Andreas Löcken, Rod McCall, Alexander Mirnig, Sebastian Osswald, Ingrid Pettersson, Bastian Pfleging, Ioannis Politis, Benjamin Poppinga, Andreas Riener, Maria Rimini-Doering, Shadan Sadeghian Borojeni, Albrecht Schmidt, Steven Shladover, Christine Sutter, Jacques Terken, Mohan Trivedi, Philipp Wintersberger, Jürgen Ziegler

The next big change in the automotive domain will be the move towards semi-automated and automated driving. The pathway to autonomous driving supported by rapid advance of a wide range of novel vehicle-related technology presents industry, academia, and regulatory agencies with new opportunities and challenges in re-imagining human interactions in the vehicle. While expectations are high towards automated driving the revolution will proceed in incremental steps; with the progress of technology new tasks and driving phases will be supported by automation. All of this will unfold in traffic scenarios in which different levels of automation will coexist for many years in which user interfaces play a key role.

We see three core challenges for automotive user interfaces in the age of automation, which we have addressed during the seminar.

- **Transforming vehicles into places of productivity and play.** People in automated vehicles will be able to turn their attention to non-driving tasks some of the time, or even much of the time. This will allow user interface designers to explore a range of possible interactions, which are might be too distracting in manually driven vehicles. For highly automated vehicles our constraints will have to do less with the driver's attention to the road, and more with the characteristics of the vehicle, such as the area available for interaction, the motion of the vehicle, as well as its computational power and the sensors that are available in the cockpit. User interactions will include other people in the vehicle, but might also include people in other vehicles. Novel user interfaces may turn the car into an infotainment and entertainment platform in which the automation allows for new secondary tasks in the car with driver and passengers that were not possible before.
- **Re-engagement of drivers into the driving task.** As automated driving makes advances, drivers will often be able to disengage from driving, and safely turn their attention to a secondary task. But until our vehicles are fully automated, drivers will eventually have to re-engage in driving. As the non-driving tasks may vary in time but also in the engagement of the user, it will be a challenge to safely and timely return to the primary task. For handling a critical situation the driver must perceive, and act upon, a sequence of information and entities. This can be a complex maneuver in a traffic scenario but also a time critical course of actions in the treatment of an emergency case. Much work needs to be done on user interface design in order to make re-engagement in different kinds of situations and different kinds of complexity safe.
- **Collaboration in mixed traffic scenarios.** Traffic automation will come to the streets peu-a-peu. Thereby and for many years, mixed scenarios in which vehicles with no-, partial-, and full automation will coexist and cooperate in daily traffic. This road sharing involves communicating autonomous operations to the driver of the autonomous car and also a communication strategy to keep non-autonomous vehicles and their drivers in the loop. Road sharing means avoiding collisions, but automated vehicles will also cooperate, for example by traveling in platoons in order to save energy and improve the utilization of the road infrastructure. Research is needed to create user interfaces that allow for safe operation of the vehicle in all of these mixed traffic scenarios.

Along with these topics, we also discussed the role of trust, e. g., how user interfaces will support the communication of trust in typical situations with mixed levels of automation. We further discussed about future technologies in and around the car (e. g., novel sensors, interaction concepts, and feedback systems) and about the recent strategy change of automakers to fund apps and invest a lot in app development to make car dashboards/instrument clusters more sustainable.

This Dagstuhl Seminar brought together researchers from human computer interaction, cognitive psychology, human factors, psychology, and also from automotive industry and OEMs to discuss the new interface paradigms for (semi-)automated driving.

Fig. 6.12
"A beautiful day at Schloss Dagstuhl for #dagstuhleas" Twitter post by 16252 Dagstuhl Seminar participant Daniel S. Katz. https://twitter.com/danielskatz/status/745909401974882304. Photo courtesy of Daniel S. Katz.

6.43 Algorithmic Foundations of Programmable Matter

Organizers: Sándor Fekete, Andréa W. Richa, Kay Römer, and Christian Scheideler

Seminar No. 16271

Date: July 3–8, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.7.1

Participants: Luca Becchetti, Julien Bourgeois, Sarah Cannon, Ioannis Chatzigiannakis, Nikolaus Correll, David Doty, Yuval Emek, Sándor Fekete, Robert Gmyr, Heiko Hamann, Jacob Hendricks, Stephan Holzer, Irina Kostitsyna, Dominik Krupke, Fabian Daniel Kuhn, Matteo Lasagni, Othon Michail, Venkateswarlu Muni, André Naz, Pekka Orponen, Matthew J. Patitz, Theodore P. Pavlic, Benoît Piranda, Andréa W. Richa, Kay Römer, Trent Rogers, Nicola Santoro, Christian Scheideler, Arne Schmidt, Robert Schweller, Thim Frederik Strothmann, Sebastian von Mammen, Jennifer L. Welch, Andrew Winslow, Damien Woods, Yukiko Yamauchi

Programmable matter refers to a substance that has the ability to change its physical properties (shape, density, moduli, conductivity, optical properties, etc.) in a programmable fashion, based upon user input or autonomous sensing. The potential applications are endless, e.g., smart materials, autonomous monitoring and repair, or minimal invasive surgery. Thus, there is a high relevance of this topic to industry and society in general, and much research has been invested in the past decade to fabricate programmable matter. However, fabrication is only part of the story: without a proper understanding of how to program that matter, complex tasks such as minimal invasive surgery will be out of reach. Unfortunately, only very few people in the algorithms community have worked on programmable matter so far, so programmable matter has not received the attention it deserves given the importance of that topic.

The Dagstuhl seminar "Algorithmic Foundations of Programmable Matter" aimed at resolving that problem by getting together a critical mass of people from algorithms with a selection of experts from distributed systems and robotics in order to discuss and develop models, algorithms, and technical solutions for programmable matter.

The aim of the proposed seminar was to bring together researchers from the algorithms community with selected experts from robotics and distributed systems in order to set a solid base for the development of models, technical solutions, and algorithms that can control programmable matter. The overall mix worked quite well: researchers from the more practical side (such as Julien Bourgeois, Nikolaus Correll, Ted Pavlic, Kay Römer, among others) interacted well with participants from the theoretical side (e.g., Jennifer Welch, Andrea Richa, Christian Scheideler, Sándor Fekete, and many others). Particularly interesting to see were well-developed but still expanding areas, such as tile self-assembly that already combines theory and practice (with visible and well-connected scientists such as Damien Woods, Matt Patitz, David Doty, Andrew Winslow, Robert Schweller) or multi-robot systems (Julien Bourgeois, Nikolaus Correll, Matteo Lasagni, André Naz, Benoît Piranda, Kay Römer).

The seminar program started with a set of four tutorial talks given by representatives from the different sets of participants to establish a common ground for discussion. From the robotics and distributed system side, Nikolaus Correll and Julien Bourgeois gave tutorials on smart programmable materials and on the claytronics programmable matter framework respectively. From the bioengineering side, Ted Pavlic gave a tutorial on natural systems that may inspire programmable matter. From the algorithmic side, Jacob Hendricks gave a tutorial on algorithmic self-assembly. In the mornings of the remaining four days, selected participants offered shorter presentations with a special focus on experience from the past work and especially also open problems and challenges. Two of the afternoons were devoted to discussions in breakout groups. Four breakout groups were formed, each with less than 10 participants to allow for intense interaction. Inspired by a classification of research questions in biology into "why?" and "how?" questions presented in Ted Pavlic's tutorial, the first breakout session was devoted to the "why?" questions underpinning programmable matter, especially also appropriate models of programmable matter systems (both biological or engineered) suitable for algorithmic research. The second breakout sessions towards the end of the seminar was devoted to a set of specific questions given by the organizers that resulted from the discussions among the participants, they included both research questions and organizational questions (e.g., how to proceed after the Dagstuhl seminar). After each of the two breakout sessions, one participant of each of the four breakout groups reported back the main findings of the discussions to the plenum, leading to further discussion among all participants. One of the afternoons was devoted to a hike to a

nearby village, where the participants also visited a small museum devoted to programmable mechanical musical devices.

The seminar was an overwhelming success. In particular, bringing together participants from a number of different but partially overlapping areas, in order to exchange problems and challenges on a newly developing field turned out to be excellent for the setting of Dagstuhl – and the opportunities provided at Dagstuhl are perfect for starting a new community.

Participants were enthusiastic on a number of different levels:

- Meeting experts from other fields provided additional insights, challenges and focus when considering work on programmable matter.
- Interacting with colleagues in a close and social manner gave many starting points for continuing collaboration.
- Getting together in a strong, large and enthusiastic group provided the opportunity to plan a number of followup activities.

The latter include connecting participants via a mailing list, the planning and writing of survey articles in highly visible publication outlets, and a starting point for specific scientific workshops and conferences.

Participants were highly enthusiastic about the possibility of another Dagstuhl workshop in the future; organizers will keep the ball rolling on this – most likely, for an application in the coming spring, so that some more details can be worked out in the meantime.

6.44 Network Latency Control in Data Centres

Organizers: Mohammad Alizadeh Attar, Jon Crowcroft, Lars Eggert, and Klaus Wehrle

Seminar No. 16281

Date: July 10–13, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.7.15

Participants: Alexandru Agache, Mohammad Alizadeh Attar, Jari Arkko, Marinho Barcellos, Olivier Bonaventure, Bob Briscoe, Edouard Bugnion, Marco Canini, Julian Chesterfield, Mosharaf Chowdhury, Jon Crowcroft, Koen De Schepper, Aaron Yi Ding, Nandita Dukkipati, Lars Eggert, Patrick Thomas Eugster, Matthew P. Grosvenor, Dongsu Han, Oliver Hohlfeld, Michio Honda, Patrick Jahnke, Lavanya Jose, Wolfgang Kellerer, Kirill Kogan, Mirja Kühlewind, Mike Marty, Andrew W. Moore, John Ousterhout, Jitendra Padhye, Rong Pan, Max Plauth, Ihsan Ayyub Qazi, Costin Raiciu, Michael Scharf, Hakim Weatherspoon, Klaus Wehrle, Michael Welzl, Noa Zilberman

Data centres are at the heart of the modern Internet. They host web services, social networking, cloud computing and are increasingly used by operators to host virtual network functions. All these services have one thing in common: they require extremely low latency communication in the data centre. Consequently we have seen the birth of a new field in networking research – data centre latency control.

Unlike the earlier generation of high-performance computing clusters, data centres have tended to use commodity off-the-shelf servers and switches, and run standard operating systems. However, traditional networking equipment and TCP-IP stacks were designed for wide-area networks, where the goal is to maximize throughput, and the control loop between end systems is measured in 10 s of milliseconds. By contrast, data centres operate on timescales that are several orders of magnitude lower. And while throughput is important, the plentiful bandwidth of data centre networks makes throughput a secondary concern to latency.

This seminar explored existing and future techniques for controlling data centre latency across the entire software and hardware stack, including in-network solutions, end-host solutions, and others. The aims of the seminar are to foster closer collaboration between academic researchers, industry, and operators. 38 researchers attended the multidisciplinary seminar. Over the course of the 3-day seminar, seven presentations were given on various aspects of data center networking. Taking the presentations as input, the workshop then broke into six working groups to discuss research aspects of latency control. The seminar was concluded by voting and discussing on possible conclusions from our discussions. Each conclusion was discussed briefly, then voted on. The outcome of the breakout session as well as the concluding statements are summarized in the full report.

Fig. 6.13
Impressions of the creation of a second large lecture hall at Schloss Dagstuhl.

6.45 Topological Methods in Distributed Computing

Organizers: Dmitry Feichtner-Kozlov and Damien Imbs

Seminar No. 16282

Date: July 10–15, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.7.31

Participants: Peva Blanchard, Jeremie Chalopin, Emanuele Delucchi, Dmitry Feichtner-Kozlov, Etienne Fieux, Eli Gafni, Emmanuel Godard, Hans Hagen, Damien Imbs, John Frederick Jardine, Petr Kuznetsov, Jan-Philipp Litza, Hammurabi Mendes, Roy Meshulam, Susumu Nishimura, Ami Paz, Eloi Perdereau, Thibault Rieutord, Matthieu Roy, Ulrich Schmid, Jan Felix Senge, Julien Stainer

This seminar brought together 22 researchers in combinatorial topology and in theoretical distributed computing. Participants came from Germany, France, Israel, Switzerland, United States, Canada, Japan and Austria. The seminar featured a combination of 1-hour talks, group sessions and an open problems session.

Scientific background and topics of the seminar

In the classical sequential computational model, computability is viewed through the Church-Turing thesis, where computations are reduced to those done by Turing machines, and complexity issues are of central importance. In the distributed setting, the situation is quite different. Since the threads of executions may intertwine in various ways (depending on the model), one of the central issues becomes dealing with execution ambiguity, and deciding whether certain standard tasks (Consensus, Renaming, etc.) are computable at all.

In this sense, the distributed setting is harder to analyze rigorously than the sequential one, or at least the difficulties are of quite different type. At the same time very many real-life situations need to be modeled by the distributed setting. These include networks of banking machines, or networks of flight controllers and airplanes, who need to reach a common decision in a decentralized setting. Another example is the parallel chip design, where we need to understand what type of elementary operations – so-called computational primitives, have to be implemented on the hardware level, so that the resulting computational system is powerful enough for our needs.

In the 80s it was realized (due in particular to the work of Fischer, Lynch and Paterson) that certain standard tasks (Consensus) cannot be solved in standard computational models (such as Message-Passing) in the presence of even simple processor crash failures. As spectacular as it is, it has become one of the steps in the development of a sophisticated and beautiful subject of theoretical distributed computing; we refer here to the classical books of Lynch and Attiya & Welch.

In the late 90s and in the early years of our millenium, it was realized by at least 3 independent groups of researchers that topological methods are applicable in proving impossibility results in theoretical distributed computing. There followed a process of further penetration of topological methods, which by now have gained a definite foothold in distributed computing. Additionally, there has also been some work on mathematical foundations, though much remains to be done when it comes to precise definitions and rigorous proofs. Independently, we feel that it is of great interest to develop the mathematics which is inspired by these methods.

The state-of-the-art of the subject was recently summarized in a book by Herlihy, Kozlov and Rajsbaum. One of the paradigms introduced there is to replace the computational task specification by the triple: input complex, output complex, and task specification map, there the input and the output complexes are simplicial complexes with additional structure, and the task specification map is what we call a carrier map, whose definition reflects our desire to restrict ourselves to the wait-free protocols. All the wait-free tasks can be encoded this way, and as a result one obtains both well-known as well as new structures from combinatorial topology.

Furthermore, one can consider the simplicial model for the totality of all executions of a given protocol – the so-called protocol complex. In the full formal setting one actually considers a triple of two simplicial complexes and a carrier map, each one equipped with an additional structure. The intuition here is that the second simplicial complex, as well as the carrier map depend heavily on the model of computation that we choose. One standard example is to take the so-called Immediate Snapshot model. On Figure 6.14 we show the protocol complex for the

one-round execution of the standard immediate snapshot protocol for 3 processors. As already this example shows, frequently there is a purely combinatorial description of the arising simplicial structure. The question of wait-free computability of a given task in a given computational model reduces then to the question of existence of a simplicial map from the protocol complex to the output complex, the so-called decision map, which satisfies certain conditions, which in essence mean that the outputs obtained by the protocol are valid under the task specification. Furthermore, we also have mathematical models for anonymous tasks, and anonymous protocols, as well as for colorless tasks.

As one can see, the mathematics needed for the current model is essentially that of simplicial complexes and carrier maps between them. With subsequent deepening of the theory and diversification of the considered questions, many further mathematical fields are coming in: for example, one needs to consider group actions and equivariant maps, as well as simplicial and carrier maps which satisfy other, less standard conditions. Many of the questions which arise in this setup are somewhat different from the questions classically studied in the simplicial context.

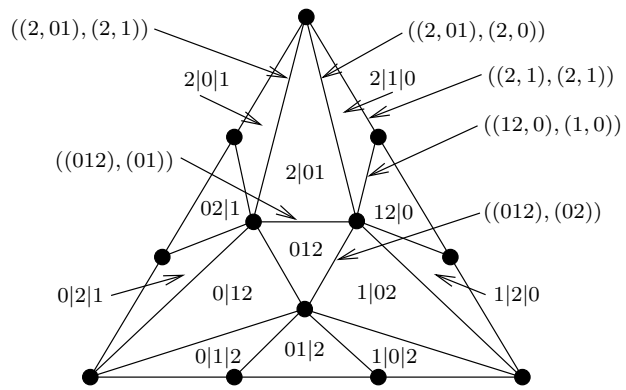

Fig. 6.14
The protocol complex for the one-round execution of the standard immediate snapshot protocol for 3 processors.

6.46 Data, Responsibly

Organizers: Serge Abiteboul, Gerome Miklau, Julia Stoyanovich, and Gerhard Weikum

Seminar No. 16291

Date: July 17–22, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.7.42

Participants: Serge Abiteboul, Marcelo Arenas, Solon Barocas, Chaitanya Baru, Claudia Bauzer Medeiros, Bettina Berendt, Claude Castelluccia, Nicholas Diakopoulos, Marina Drosou, Gerald Friedland, Sorelle Friedler, Irini Fundulaki, Krishna P. Gummadi, Michael Hay, Bill Howe, H. V. Jagadish, Benny Kimelfeld, Amélie Marian, Pauli Miettinen, Gerome Miklau, Wolfgang Nejdl, Benjamin Nguyen, Evaggelia Pitoura, Salvatore Ruggieri, Rishiraj Saha Roy, Arnaud Sahuguet, Eric Simon, Julia Stoyanovich, Jannik Strötgen, Fabian Suchanek, Kristene Unsworth, Jan Van den Bussche, Suresh Venkatasubramanian, Agnès Voisard, Nicholas Weaver, Gerhard Weikum, Christo Wilson, Cong Yu, Ben Zevenbergen

Our society is data-driven. Large scale data analysis, known as Big data, is distinctly present in the private lives of individuals, is a dominant force in commercial domains as varied as automatic manufacturing, e-commerce and personalized medicine, and assists in – or fully automates – decision making in the public and private sectors. Data-driven algorithms are used in criminal sentencing – ruling who goes free and who remains behind bars, in college admissions – granting or denying access to education, and in employment and credit decisions – offering or withholding economic opportunities.

The promise of Big data is to improve people's lives, accelerate scientific discovery and innovation, and enable broader participation. Yet, if not used responsibly, Big data can increase economic inequality and affirm systemic bias, polarize rather than democratize, and deny opportunities rather than improve access. Worse yet, all this can be done in a way that is non-transparent and defies public scrutiny.

Big data impacts individuals, groups and society as a whole. Because of the central role played by this technology, it must be used *responsibly* – in accordance with the ethical and moral norms that govern our society, and adhering to the appropriate legal and policy frameworks. And as journalists [3], legal and policy scholars [1, 2] and governments [4, 5] are calling for algorithmic fairness and greater insight into data-driven algorithmic processes, there is an urgent need to define a broad and coordinated computer science research agenda in this area. The primary goal of the Dagstuhl Seminar "Data, Responsibly" was to make progress towards such an agenda.

The seminar brought together academic and industry researchers from several areas of computer science, including a broad representation of data management, but also data mining, security/privacy, and computer networks, as well as social sciences researchers, data journalists, and those active in government think-tanks and policy initiatives. The problem we aim to address is inherently transdisciplinary. For this reason, it was important to have input from policy and legal scholars, and to have representation from multiple areas within computer science. We were able to attract a mix of European, North American, and South American participants. Out of 39 participants, 10 were women.

Specific goals of the seminar were to:

- assess the state of data analysis in terms of fairness, transparency and diversity;
- identify new research challenges;
- develop an agenda for computer science research in responsible data analysis and use, with a particular focus on potential high-impact contributions from the data management community;
- solicit perspectives on the necessary education efforts, and on responsible research and innovation practices.

The seminar included technical talks and break-out sessions. Technical talks were organized into themes, which included fairness and diversity, transparency and accountability, tracking and transparency, personal information management, education, and responsible research and innovation. Participants suggested topics for seven working groups, which met over one or multiple days.

The organizers felt that the seminar was very successful – ideas were exchanged, discussions were lively and insightful, and we are aware of several collaborations that were started as a result of the seminar. The participants and the organizers all felt that the topic of the seminar is broad, fast moving and extremely important, and that it would be beneficial to hold another seminar on this topic in the near future.

Details about the program are contained in the full report.

References

1 Kate Crawford. Artificial Intelligence's White Guy Problem. The New York Times, June 25, 2016.

2 Kate Crawford and Ryan Calo. There is a blind spot in AI research. Nature / Comment 538(7625), October 13, 2016.

3 Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. ProPublica, May 23, 2016.

4 Executive Office of the President, The White House. Big Data: Seizing Opportunities, Preserving Values. May 2014

5 Parliament and Council of the European Union. General Data Protection Regulation. 2016

6.47 Coding Theory in the Time of Big Data

Organizers: Martin Bossert, Eimear Byrne, and Emina Soljanin

Seminar No. 16321

Date: August 7–12, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.8.1

Participants: Iryna Andriyanova, Allison Beemer, Jessalyn Bolkema, Martin Bossert, Michael Braun, Eimear Byrne, Viveck Cadambe, Gerard Cohen, Alex Dimakis, Iwan M. Duursma, Michelle Effros, Rafah El-Khatib, Tuvi Etzion, Christina Fragouli, Heide Gluesing-Luerssen, Elisa Gorla, Marcus Greferath, Venkatesan Guruswami, Daniel Heinlein, Camilla Hollanti, Anna-Lena Horlemann-Trautmann, Christine A. Kelley, P. Vijay Kumar, Julia Lieb, Hans-Andrea Loeliger, Michael Luby, Felice Manganiello, Carolyn Mayer, Sihem Mesnager, Olgica Milenkovic, Elie Najm, Alessandro Neri, Mario Osvin Pavcevic, Sven Puchinger, Tovohery Randrianarisoa, Alberto Ravagnani, Johan Rosenkilde, Joachim Rosenthal, John Sheekey, M. Amin Shokrollahi, Vitaly Skachek, Emina Soljanin, Pascal Vontobel, Jing Zhong

The Dagstuhl Seminar 16321 *Coding Theory in the Time of Big Data*, held in August 7–12, 2016, was the third of a series of Dagstuhl seminars relating modern aspects of coding theory and its applications in computer science. The overarching technical theme was on how fundamentals of coding theory could be applied to data storage and transmission in the context of big data and conversely, on emerging topics in coding theory arising from such applications. In Dagstuhl Seminar 11461 the main topics discussed were list decoding, codes on graphs, network coding and the relations between them. The themes of distributed storage, network coding and polar codes were central to Dagstuhl Seminar 13351.

The conference was organised into six main working groups, as listed below:

1. Distributed Storage & Index Coding,
2. Private Information Retrieval for Storage Codes,
3. DNA-Based Storage,
4. Age & Delay of Information.
5. Code-Based Cryptography,
6. Rank-Metric Codes.

The amount of data that is being stored is scaling at a rapid pace making efficient data storage an important problem that inspires several lines of scientific research. During the seminar, several discussions were conducted on the theme of using classical and new techniques from coding theory to store/compute data efficiently in distributed storage systems. A number of open problems were identified, such as the design of codes with optimal repair bandwidth, fundamental trade-offs between storage & communication cost, applications to content distribution networks, connections between fundamental limits of storage/caching and the index coding problem and applications of coding theory for parallel computing. A theoretical framework and numerical simulation for the long term reliability of a distributed storage system were presented by Luby.

DNA-based storage was recently proposed to address new challenges to handle extremely high volume recording media to propose new compression methods for non-traditional data formats. Since DNA may be easily replicated and a massive amount of information stored reliably with minimal space requirements, it has enormous potential as a method of big data storage. This was the focus of the DNA working group. Problems such read and write cost, insertion and deletion errors arising in sequences, error reduction were discussed. Milenkovic gave an introductory talk describing several problems associated with whole genome, sequencing read, RNA-seq and ChiP-seq data compression, and outlined the first portable DNA-based rewritable and random access storage system.

Private information retrieval (PIR) enables a user to retrieve a data item from a database without disclosing the identity of the item retrieved, while the data itself may be public. The PIR working group considered this problem in the context of storage codes, in particular for dynamic coded storage and adversarial PIR, with some extensions to asynchronized systems, batch codes and private keyword search. Hollanti gave a tutorial overview of recent results in the area.

Age of information is a metric for status updating systems, where a monitor is interested in staying timely about the status of a source. The optimal updating strategy that minimizes the average age exists when the updating rate is constrained by limited network resources. Streaming source coding problems can be applied to the problem of age analysis. The main focus the Age & Delay working group was to introduce the age of information concept to participating coding theorists and explore potential age and delay problems in coding and storage. An adaptive arithmetic coding scheme was proposed as a potential solution to avoid huge decoding delay. Several possible delay problems in

file downloading from multiple servers were discussed. Two PhD students, Zhong and Najm gave a tutorial overview of the topic.

Code-based crypto-systems are some of the very few that resist quantum-based attacks. In the case subfield subcodes such as the Goppa or Srivastava codes no successful attack is known yet. Moderate-density parity-check (MDPC) codes have been proposed for key size reduction in such crytposystems. The group identified open problems such as investigating other subfield subcodes and attacks on MDPC structured codes. An overview was presented by Bossert.

Rank-metric codes have applications in random network coding, coded-caching and in code-based cryptography. The working group focussed on maximum rank distance (MRD) codes, specifically their classifications and on algebraic methods for constructing and decoding families of them. New nontrivial classifications were obtained. Further research directions on the classification problem were identified such as adapting semi-field theory techniques and searches for codes with high symmetry. Given the known limitation of list decoders for Gabidulin codes, the group worked on adapting decoders for Gabidulin codes to recent families of MRD codes. Sheekey presented recent results on MRD codes and described links to semifields.

A total of 44 researchers participated in the seminar across these working groups. In addition, several participants took the opportunity to collaborate with others on specific related projects. There were 16 talks in total, several related to storage of big data and others on topics such as maximum rank distance codes, chip-to-chip communication, the MDS conjecture, the SAGE computer algebra system, age of information, the edge removal problem, convolutional codes and network coding. Among the talks given were some tutorial presentations, aimed at introducing researchers to fundamentals of a related working group. The working groups focussed on identifying and addressing new and/or important open problems in the area. Age & Delay, PIR for storage codes and DNA-based storage were new topics to many participants and generated considerable interest.

6.48 Integrating Process-Oriented and Event-Based Systems

Organizers: David Eyers, Avigdor Gal, Hans-Arno Jacobsen, and Matthias Weidlich

Seminar No. 16341

Date: August 21–26, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.8.21

Participants: Alexander Artikis, Jean Bacon, Anne Baumgraß, Alejandro P. Buchmann, Claudio Di Ciccio, David Eyers, Avigdor Gal, Annika M. Hinze, Hans-Arno Jacobsen, Martin Jergler, Boris Koldehofe, Oliver Kopp, Agnes Koschmider, Sankalita Mandal, Alessandro Margara, Ken Moody, Cesare Pautasso, Wolfgang Reisig, Stefanie Rinderle-Ma, Mohammad Sadoghi Hamedani, Stefan Schulte, Arik Senderovich, Vinay Setty, Jatinder Singh, Sergey Smirnov, Pnina Soffer, Wei Song, Jan Sürmeli, Lucinéia Heloisa Thom, Martin Ugarte, Stijn Vansummeren, Roman Vitenberg, Matthias Weidlich, Lijie Wen, Mathias Weske, Kaiwen Zhang, Holger Ziekow

Background and Motivation Process-oriented information systems are software systems that execute and manage a process, broadly defined as a coordinated execution of actions to achieve a certain goal. As such, they support Business Process Management (BPM) initiatives. Process-oriented systems have been traditionally used in domains such as business process automation, enterprise application integration, and collaborative work. Recently, there has also been a significant uptake of process-oriented information systems in transportation, logistics, and medical infrastructures – domains that impose new challenges in terms of system reactivity and adaptability. Here, trends such as sensing of data (e.g., based on RFID technology) and advancing system integration (driven by technical standards such as EPCglobal) represent opportunities to strengthen the event-perspective in process-oriented systems in order to achieve more flexible and comprehensive process control.

Event-based systems, in turn, have been put forward to integrate heterogeneous systems in a flexible and scalable manner by separating communication from application logic. These systems provide interaction models, mechanisms for routing events between components, and techniques for the detection of composite events, i.e., for Complex Event Processing (CEP). Although event-based systems are typically positioned as general-purpose technology, they have found their way into many applications where event generation is comparatively deterministic and follows structured behaviour. In domains such as transportation, logistics, and the medical sector, events handled by event-based systems stem from the execution of processes, which are partially supported by process-oriented information systems. Exploiting the process-perspective, therefore, promises to lead to advancements in the design, analysis, and optimisation of event-based systems.

The increasing overlap of application scenarios that involve concepts and techniques of process-oriented as well as event-based systems, however, is only marginally supported by exchange and convergence of the related research fields. Strong communities have been established for research on either type of system. Yet, due to the missing link between these communities, manifold opportunities for ground-breaking research and broad impact in industry are missed out. Research efforts related to the underlying theory as well as specific platforms are duplicated and similar approaches are developed in both communities.

Breaking this disconnect had been the goal that the seminar aimed to achieve by identifying the links between conceptual models, formal analysis methods, and engineering techniques developed for either type of system.

Seminar Structure Given that seminar attendees came from two rather disconnected communities, the first day of the seminar featured four tutorials to establish a joint understanding of essential concepts and terminology. First, Alessandro Margara presented an overview of the basic techniques to manage streams of events. Mathias Weske then gave a primer on BPM, elaborating on the main concepts, models, and the role of events for process management. An advanced view on techniques for event processing was given by Alejandro Buchmann. Stefanie Rinderle-Ma closed this part of the seminar with a tutorial on management, utilisation, and analysis of instance data in distributed process management.

The remainder of the seminar week was centred on break-out sessions, in which participants worked on particular topics on the intersection of process-oriented and event-based systems. In these working groups, participants discussed the relevant state-of-the-art and identified the research challenges under a near-, mid-, or long-term perspective. In addition, there were two sessions in which seminar participants gave a very short overview of their recent research work.

Topics and Key Challenges The working groups focussed on a diverse set of topics, highlighting the key challenges that need to be addressed:

Event Models for BPM: Semantics of Events and Patterns. Starting from the observation that event models are well-established in both BPM and CEP and that their coupling has obvious benefits, the challenge relates to the question of how events can guide the evolution or adaptation of process instances.

Towards Automatic Event-Based Monitoring of Processes. Event-based monitoring of processes is influenced by the availability of patterns, the consequences of monitoring results, and the integration of contextual information. These dimensions render it particularly challenging to comprehensively discover and utilise patterns for process monitoring.

Patterns and Models for Communication. The communication models underlying an event-based middleware have diverse implications for the interplay of processes and event patterns – and a major challenge is the identification of requirements that are imposed by process scenarios on communication models.

Choreographies and Inter-Process Correlation. Common languages for the description of interacting processes lack capabilities for the specification of event-based processing. The challenge is to develop a better grounding of choreography languages and enable analysis of the information flow between processes.

Abstraction Levels: Processes versus Events. Observing that methods in BPM mainly proceed top-down, whereas event processing is often approached bottom-up, a key challenge is the identification of the right abstraction level on which concepts and methods shall be integrated.

Context in Events and Processes. The context of a process may influence event processing, and the context as materialised in complex events impacts the execution of a process. Yet, a suitable representation and dynamic evolution of context information is an open research challenge.

Integrated Platforms for BPM & CEP. The integration of traditional BPM or CEP engines promises accelerated application development and lower maintenance cost. To attain this end, the challenge of developing a unified model for events and processes, enabling well-grounded architectural decisions, needs to be addressed.

(Highly) Distributed Processes & The Role of Events. Events and processes can both be handled in a centralised or distributed infrastructure and open challenges relate to the tradeoffs regarding trustworthiness, reliability, and scalability.

Event Data Quality. Event data may be uncertain, which needs to be reflected in processes that are influenced by these events. The challenge is how to capture such uncertainty and make explicit how it influences decision making on the level of the process.

From Event Streams to Process Models and Back. Event patterns and processes are typically concerned with events on different levels of abstractions, which can be bridged only on the basis of a unifying formal model. Further challenges arise from the imprecision of event definitions in processes and the expressiveness of CEP languages when capturing procedural behaviour.

Compliance, Audit, Privacy and Security. Compliance checking of business processes may benefit from CEP systems and BPM tools may be useful to express service level agreements in event-based systems. Challenges, however, are methods for a structured integration of BPM and CEP technology and their alignment with informal compliance requirements.

Main Recommendations From the discussions and the exchange of ideas during the workshop, a set of recommendations was able to be distilled in order to materialise the benefits of integrating process-oriented and event-based systems.

Build a community around BPM and CEP. The topics on the intersection of process-oriented and event-based systems provide a rich field for high-impact research. The number and diversity of open research questions call for a long-term research initiative, so that a respective community needs to be built up. To achieve this, it is recommended that joint workshops be initiated at the flagship conferences in either field, the BPM conference and the DEBS conference, and to evaluate potential co-location of the conferences in future.

Start research on integrated models. For many of the aforementioned challenges, the lack of integrated models, in which processes and events are first-class citizens, turns out to be a major issue. Research shall be devoted to creating such models, clarifying which basic notions of events exist, and considering the semantics of distributed event generation.

Facilitate joint research. Joint research is currently hindered not only by the disconnect of the research communities, but also by a lack of a common set of standard concepts in either community. There is a need for concise overviews of the most important concepts and methods in either field, e.g., by means of standard textbooks. Researchers from one field need to be able to quickly gather the level of understanding of the other field that is required for joint research initiatives.

Engage industry. The integration of process-oriented and event-based systems is driven by particular domains, such as logistics, health, and mobility. The prioritisation of challenges and the evaluation of developed solutions critically depends on the involvement of industrial partners from these domains. As such, it is recommended to reach out to industry to develop evaluation scenarios and benchmark datasets. One viable means for this are the research proposals on the EU and national levels that involve BPM and CEP experts from both academia and industry.

6.49 Foundations of Secure Scaling

Organizers: Lejla Batina, Swarup Bhunia, and Patrick Schaumont

Seminar No. 16342

Date: August 21–26, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.8.65

Participants: Debapriya Basu Roy, Lejla Batina, Guido Bertoni, Swarup Bhunia, Christian Boit, Chen-Mou Cheng, Joan Daemen, Jia Di, Thomas Eisenbarth, Naofumi Homma, Yier Jin, Nele Mentens, Debdeep Mukhopadhyay, Ventzislav Nikov, Svetla Petkova-Nikova, Bart Preneel, Sandip Ray, Francesco Regazzoni, Kazuo Sakiyama, Patrick Schaumont, Peter Schwabe, Georg Sigl, Shahin Tajik, Ingrid Verbauwhede, Hirotaka Yoshida, Bilgiday Yuce

In electronic system design, scaling is a fundamental force present at every abstraction level. Over time, chip feature sizes shrink; the length of cryptographic keys and the complexity of cryptographic algorithms grows; and the number of components integrated in a chip increases. While scaling is generally thought of as beneficial to the resulting implementations, this does not hold for secure electronic design. Larger and faster chips, for example, are not necessarily more secure. Indeed, the relations between scaling and the resulting security are poorly understood. This Dagstuhl Seminar hosted researchers in secure electronic system design, spanning all abstraction levels from cryptographic engineering over chip design to system integration.

Discussion Topics

The mechanisms of secure scaling require investigation of the links between Cryptography, Technology, and Digital Integration. Cryptographers are concerned with novel and secure algorithms that remain secure even as cryptanalytic capabilities improve. Technologists are concerned with the next generation of transistors and their implementation into a reliable and stable process technology. Integrators are concerned with electronic design automation tools that can manage the rapidly increasing complexity of electronic design, and the are concerned with the integration of components on a complex system-on-chip.

Through its participants, the seminar offered a unique opportunity to discuss cross-cutting topics in Secure Scaling. The following list are examples of such cross-cutting topics.

- Scaling effects in Privacy and Security. The massive amount of connected devices will create significant challenges towards security and privacy. Major questions involve data ownership and key ownership and management.
- Power/Energy Efficient Crypto: Secure wireless devices and Secure RFID are two well known examples of applications that require security under severe power and/or energy constraints. Optimizing a cryptographic algorithm for power/energy efficiency needs to consider all abstraction levels of design.
- High-Performance Crypto: Information Technology is increasingly asymmetric, with larger, high-performance servers at one end, and a large population of tiny devices at the other side. Cryptographic designs must scale towards high-performance, high-throughput implementations while it must also accommodate small-footprint, low-latency designs.
- Secure Test: Complex chips utilize a number of testing strategies such as BIST and JTAG. When a chip includes a secure part, the test infrastructure carries a potential risk of abuse. Secure Test is a test strategy for complex chips that takes this risk fully into account.
- Complexity Management in Secure SoC: Managing and integrating a secure module into system-on-chip context is challenging and creates a hard verification problem that cuts through multiple traditional layers of design. Furthermore, managing multiple stakeholders in a single chip design is extremely challenging and may result in conflicting design requirements.
- Implementation Attacks: In modern cryptographic designs, side-channel analysis, fault-analysis and physical tampering are an integral part of the threat model. This requires design techniques that fully integrate countermeasures as part of the design process. In addition, the design of a countermeasure effective against most forms of tampering is an open research issue.
- Technology effects on implementation attacks. Better insight the internal operation of secure implementations at all abstraction levels leads to novel implementation attacks, that work at finer granularity, and that use novel source of leakage such as optical leakage.

The seminar supported participants in learning about the state-of-the-art developments in the three different domains covered in the workshop (Cryptography, Integration, and Technology). The seminar also supported the presentation of specific cross-cutting topics, as well as round-table (panel-style) discussions.

6.50 Next Generation Sequencing – Algorithms, and Software For Biomedical Applications

Organizers: Gene Myers, Mihai Pop, Knut Reinert, and Tandy Warnow

Seminar No. 16351

Date: August 28 to September 2, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.8.91

Participants: Niko Beerenwinkel, Ewan Birney, Christina Boucher, Jason Chin, Pascal Costanza, Anthony J. Cox, Fabio Cunial, Richard Durbin, Mohammed El-Kebir, Anne-Katrin Emde, Simon Gog, Hannes Hauswedell, Daniel H. Huson, André Kahles, Birte Kehr, Gunnar W. Klau, Oliver Kohlbacher, Ben Langmead, Pietro Lio', Veli Mäkinen, Tobias Marschall, Alice Carolyn McHardy, Siavash Mirarab, Laurent Mouchard, Gene Myers, Luay Nakhleh, Kay Nieselt, Enno Ohlebusch, Adam M. Phillippy, Mihai Pop, Simon J. Puglisi, Gunnar Rätsch, Tobias Rausch, Knut Reinert, Karin Remington, Bernhard Renard, S. Cenk Sahinalp, Enrico Siragusa, Peter F. Stadler, Granger Sutton, German Tischler, Esko Ukkonen, Tandy Warnow, David Weese, Shibu Yooseph

Motivation

In recent years, Next Generation Sequencing (NGS) data have begun to appear in many applications that are clinically relevant, such as resequencing of cancer patients, disease-gene discovery and diagnostics for rare diseases, microbiome analyses, and gene expression profiling, to name but a few. Other fields of biological research, such as phylogenomics, functional genomics, and metagenomics, are also making increasing use of the new sequencing technologies.

The analysis of sequencing data is demanding because of the enormous data volume and the need for fast turnaround time, accuracy, reproducibility, and data security. Addressing these issues requires expertise in a large variety of areas: algorithm design, high performance computing on big data (and hardware acceleration), statistical modeling and estimation, and specific domain knowledge for each medical problem. In this Dagstuhl Seminar we aimed at bringing together leading experts from both sides – computer scientists including theoreticians, algorithmicists and tool developers, as well as leading researchers who work primarily on the application side in the biomedical sector – to discuss the state-of-the art and to identify areas of research that might benefit from a joint effort of all the groups involved.

Goal of the seminar

The key goal of this seminar was a free and deep exchange of ideas and needs between the communities of algorithmicists and theoreticians and practitioners from the biomedical field. This exchange should have triggered discussions about the implications that new types of data or experimental protocols have on the needed algorithms or data structures.

Results

We started the seminar with a number of *challenge talks* to encourage discussion about the various topics introduced in the proposal. Before the seminar started we identified three areas the participants were most interested in, namely:

1. Data structures and algorithms for large data sets, hardware acceleration
2. New problems in the upcoming age of genomes
3. Challenges arising from new experimental frontiers and validation

For the first area Laurent Mouchard, Gene Myers, and Simon Gog presented results and challenges; for the second area Siavash Mirarab, Niko Beerenwinkel, Shibu Yooseph, and Kay Nieselt introduced some thoughts; and finally, for the last area, Jason Chin, Ewan Birney, Alice McHardy, and Pascal Costanza talked about challenges. For most of those talks the abstracts can be in the full report. Following this introductionary phase, the participants organized themselves into various working groups the topics of which were relatively broad. Those first breakout groups were about

- Haplotype phasing
- Big data
- Pangenomics data representation
- Cancer genomics
- Metagenomics
- Assembly

The results of the groups were discussed in plenary sessions interleaved with some impromptu talks. As a result the participants split up into smaller, more focused breakout groups that were received very well. Indeed, some participants did already extend data formats for assembly or improved recent results on full text string indices.

Based on the initial feedback from the participants we think that the topic of the seminar was interesting and led to a lively exchange of ideas. We thus intend to revisit the field in the coming years in a Dagstuhl seminar again, most likely organized by different leaders of the field in order to account for these upcoming changes. In such a seminar we intend to encourage more people from clinical bioinformatics to join into the discussions.

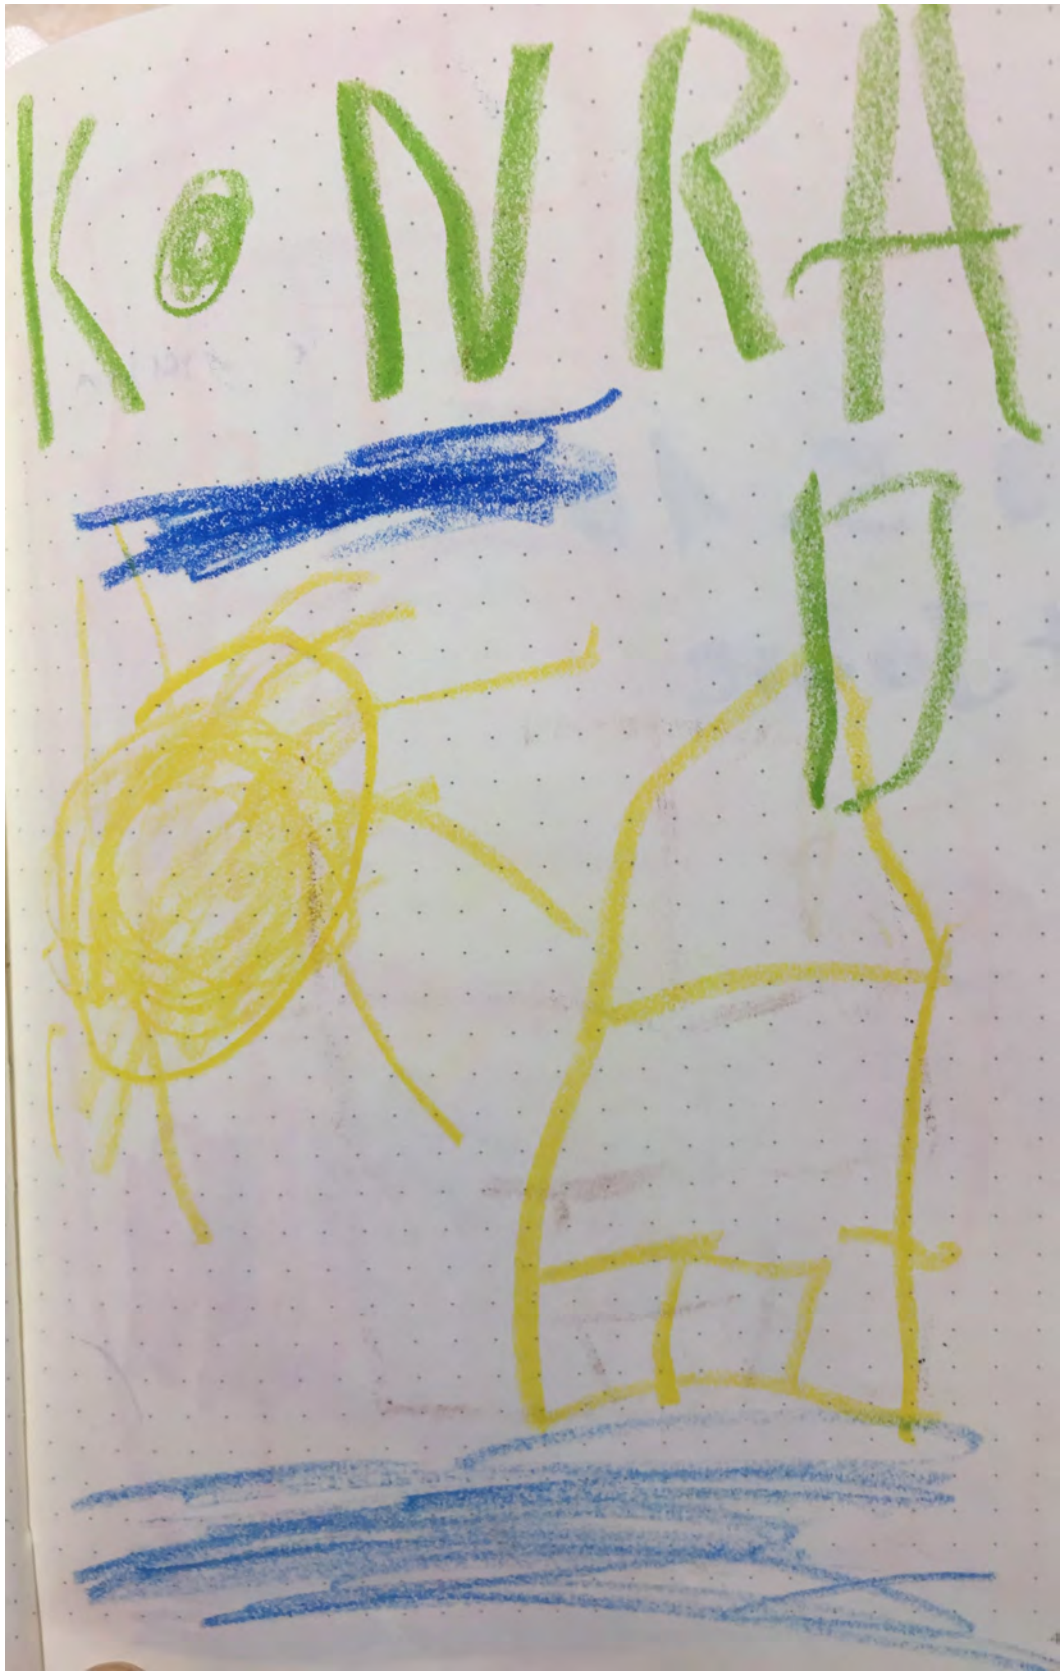

Fig. 6.15
Drawing for the Dagstuhl children's guest book by Konrad.

6.51 Network Attack Detection and Defense – Security Challenges and Opportunities of Software-Defined Networking

Organizers: Marc C. Dacier, Sven Dietrich, Frank Kargl, and Hartmut König

Seminar No. 16361

Date: September 4–9, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.1

Participants: Johanna Amann, Kpatcha Mazabalo Bayarou, José Jair C. de Santanna, L. Jean Camp, Georg Carle, Radoslaw Cwalinski, Marc C. Dacier, Hervé Debar, Sven Dietrich, Falko Dressler, Marc Eisenbarth, Felix Erlacher, Paulo Jorge Esteves-Veríssimo, Dieter Gollmann, Peter Herrmann, Marko Jahnke, Mattijs Jonker, Frank Kargl, Thomas Kemmerich, Issa Khalil, Hartmut König, Jan Kohlrausch, Boris Koldehofe, Tobias Limmer, Claas Lorenz, Thomas Lukaseder, Evangelos Markatos, Michael Meier, Michael Menth, Simin Nadjm-Tehrani, Rene Rietz, Christian Röpke, Christian Rossow, Ramin Sadre, Thomas Scheffler, Björn Scheuermann, Sebastian Schmerl, Bettina Schnor, Robin Sommer, Radu State, Jens Tölle, Alexander von Gernler, Han Xu, Emmanuele Zambon

From September 4 through 9, 2016, more than 40 researchers from the domains of computer networks and cyber security met at Schloss Dagstuhl to discuss security challenges and opportunities of software-defined networking (SDN).

Software-defined networking has attracted a great attention both in industry and academia since the beginning of the decade. This attention keeps undiminished. In 2014, IDC predicted that the market for SDN network applications would reach $1.1bn. Especially in industry, the vision of "programming computer networks" has electrified many IT managers and decision makers. There are great expectations regarding the promises of SDN. Leading IT companies, such as Alcatel-Lucent, Cisco systems, Dell, Juniper Networks, IBM, and VMware, have developed their own SDN strategies. Major switch vendors already offer SDN-enabled switches.

Software-defined networking provides a way to virtualize the network infrastructure to make it simpler to configure and manage. It separates the control plane in routers and switches, which decides where packets are sent, from the data plane, which forwards traffic to its destination, with the aim to control network flows from a centralized control application, running on a physical or virtual machine. From this controller, admins can write and rewrite rules for how network traffic, data packets, and frames are handled and routed by the network infrastructure. Routers and switches in a sense become "slaves" of this application-driven central server. SDN-enabled networks are capable of supporting user requirements from various business applications (SLAs, QoS, Policy Management, etc.). This is not limited to the network devices of a certain vendor. It can be applied to devices from various vendors if the same protocol is used. Most SDN infrastructure utilizes the widely-used OpenFlow protocol and architecture to provide communication between controllers and networking equipment.

Security-related aspects of software-defined networking have only been considered more recently. Opinions differ widely. Some believe that the security problems introduced by SDN are manageable – that SDN can even bring security benefits; others think that Pandora's Box has been opened where SDN and SDN-enabled networks can never be secured properly.

No doubt, there are a number of serious security problems as the following examples show. SDN controllers represent single points of failures. The controllers as well as the connections between controllers and network devices might be subject to distributed denial of service attacks. Compromising the central control could give an attacker command of the entire network. The SDN controllers are configured by network operators. Configuration errors can have more complex consequences than in traditional settings because they may unpredictably influence the physical network infrastructure. Furthermore, the idea of introducing 'network applications' that interact with the controller to modify network behavior seems like a complexity nightmare in terms of required authentication and authorization schemes. Finally, the SDN paradigm is a major turn around with respect to the basic design rules that have made the Internet successful so far, namely a well-defined layered approach. Whereas in today's world, applications have no say in routing decisions, SDN's promise for highly flexible and application-tailored networking requires a way for applications to optimize networking decisions for their own benefits. However, it is unclear to what extent fairness can be ensured, how conflicting decisions can be resolved, etc. Along the same line, members of the security community worry about the possibility to intentionally design SDN applications that could eventually be turned into attack weapons or simply be misused by malicious attackers. Whether these fears are substantiated or not is something which has not received any scrutiny so far.

On the other hand, SDN is also considered by many researchers as an effective means to improve the security of

networks. SDN controllers can be used, for instance, to store rules about the permission of certain requests which cannot be decided at the level of a single switch or router because this requires full overview over network status or additional information and interactions which are not contained in the current protocol versions. Attacks that can be detected this way are ARP spoofing, MAC flooding, rogue DHCP server, and spanning tree attacks. Also, by enabling the creation of virtual networks per application, people speculate that intrusion detection techniques relying on the modeling of the normal behavior of network traffic will become much easier to implement and more reliable in terms of false positive and negatives. Similarly, SDN apps could offer a very simple and effective way to implement quarantine zones for infected machines without cutting them off completely from the network since the quarantine could be customized at the application level (letting DNS and HTTP traffic for a given machine go through but not SMTP, for instance).

These two contrary facets of SDN security were the key ingredients for an extremely lively and very fruitful seminar. The seminar brought together junior and senior experts from both industry and academia, covering different areas of computer networking and IT security. The seminar started with two invited talks by Boris Koldehofe (TU Darmstadt, DE) and Paulo Jorge Esteves-Veríssimo (University of Luxembourg, LU) on the basics and security aspects of software-defined networking. After that we organized six working groups to discuss in two rounds the Good and the Bad of using SDN from the security point of view. Based on the outcome of the working groups and a plenary discussion, we formed another four working groups to discuss required research directions. The first six working groups focus on the following issues: (1) centralization in SDN, (2) standardization and transparency, (3) flexibility and adaptability for attackers and defenders, (4) complexity of SDN, (5) attack surface and defense, and (6) novelty and practicability. The research direction working groups dealt with (1) improving SDN network security, (2) a secure architecture for SDN, (3) secure operation in SDN-based environments, and (4) SDN-based security. The discussion in the working groups was supplemented by short talks of participants to express their positions on the topic or to report about ongoing research activities. Based on the talks, discussions, and working groups, the Dagstuhl seminar was closed with a final plenary discussion which summarized again the results from the working groups and led to a compilation of a list of statements regarding the security challenges and opportunities of software-defined networking. The participants agreed that SDN provides new possibilities to better secure networks, but also offers a number of serious security problems which have to be solved for being SDN a successful technology. The outcome of these discussions and the proposed research directions are presented in the following.

6.52 Robustness in Cyber-Physical Systems

Organizers: Martin Fränzle, James Kapinski, and Pavithra Prabhakar

Seminar No. 16362

Date: September 4–9, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.29

Participants: Houssam Abbas, Paul Bogdan, Alexandre Donzé, Rüdiger Ehlers, Georgios Fainekos, Martin Fränzle, Nathan Fulton, Miriam García Soto, Khalil Ghorbal, James Kapinski, Scott C. Livingston, Sarah M. Loos, Rupak Majumdar, Jens Oehlerking, Jan Otop, Necmiye Ozay, Pavithra Prabhakar, Sylvie Putot, Stefan Ratschan, Matthias Rungger, Paulo Tabuada, Ufuk Topcu, Eric M. Wolff, Bai Xue, Paolo Zuliani

Overview and Goals of the Seminar

Engineering robustness into systems under development has always been at the heart of good engineering practice, be it robustness against manufacturing tolerances and against variations in purity of construction materials in mechanical engineering, robustness against concentrations of educts in chemical engineering, against parameter variations in the plant model within control engineering, against quantization and measurement noise in signal processing, against faults in computer architecture, against attacks in security engineering, or against unexpected inputs or results in programming. In cyber-physical systems (CPS), all the aforementioned engineering disciplines meet, as the digital networking and embedded control involved in CPS brings many kinds of physical processes into the sphere of human and computer control. This convergence of disciplines has proven extremely fruitful in the past, inspiring profound research on hybrid and distributed control, transferring notions and methods for safety verification from computer science to control theory, transferring proof methods for stability from control theory to computer science, and shedding light on the complex interplay of control objectives and security threats, to name just a few of the many interdisciplinary breakthroughs achieved over the past two decades. Unfortunately, a joint, interdisciplinary approach to robustness remains evasive. While most researchers in the field of CPS concede that unifying notions across the disciplinary borders to reflect the close functional dependencies between heterogeneous components would be of utmost importance, the current state of affairs is a fragmentary coverage by the aforementioned disciplinary notions.

Synergies and research questions. The seminar set out to close the gap in the robustness investigations across the overlapping disciplines under the umbrella of CPS by gathering scientists from the entire spectrum of fields involved in the development of cyber-physical systems and their pertinent design theories. The seminar fostered interdisciplinary research answering the following central questions:

1. What is the rationale behind the plethora of existing notions of robustness and how are they related?
2. What measures have to be taken in a particular design domain (e.g., embedded software design) to be faithful to notions of robustness central to another domain it has functional impact on (e.g., feedback control)?
3. What forms of correctness guarantees are provided by the different notions of robustness and would there be potential for unification or synergy?
4. What design measures have been established by different disciplines for achieving robustness by construction, and how can they be lifted to other disciplines?
5. Where do current notions of robustness or current techniques of system design fall short and can this be alleviated by adopting ideas from related disciplines?

The overarching objective of such research would be to establish trusted engineering approaches incorporating methods for producing cyber-physical system designs

1. that sustain their correctness and performance guarantees even when used in a well-defined vicinity of their nominal operational regimes, and
2. that can be trusted to degrade gracefully even when some of the underlying modeling and analysis assumptions turn out to be false.

To satisfy these design objectives, we require notions of robustness that go well beyond the classical impurities of embedded systems, like sampling, measurement noise, jitter, and machine tolerances, and must draw on concepts of robustness from disparate fields. This seminar identified parallels between related notions of robustness from the many varied domains related to

CPS design and bridged the divide between disciplines, with the goal of achieving the above objectives.

Topics of the Seminar

This seminar aimed to identify fundamental similarities and distinctions between various notions of robustness and accompanying design and analysis methods, with the goal of bringing together disparate notions of robustness from multiple academic disciplines and application domains. The following is a brief compendium of the robustness notions and application domains that were addressed in this seminar.

Robustness Notions and Design/Analysis Methods. One goal of this seminar was to identify crosscutting frameworks and design methodologies among the different approaches used to study robustness in the domains of control theory, computer science, and mechanical engineering. We considered the following broad classifications of robustness with the ultimate goal of synergizing the notions and techniques from the various disciplines.

- Input/Output Robustness
- Robustness with respect to system parameters
- Robustness in real-time system implementation
- Robustness due to unpredictable environments
- Robustness to Faults

Application Domains. The applications for the topics addressed in this seminar include cyber-physical systems for which robustness is a vital concern. The following is a partial list of these application domains.

- Automotive
- Aeronautics
- Medical devices
- Robotics
- Smart buildings
- Smart infrastructure

Outcome

We summarize the outcomes of the discussions in the break-out sessions that were conducted by forming subgroups among the participants. The topics referred to different approaches and/or applications in the framework of robustness. One of the topics was about robustness for discrete systems. In this session, the need for defining robustness for these systems was extensively discussed, and one of the most relevant challenges identified was to define appropriate metrics on the state-space relevant to the application. Also some specific robustness issues in the domain of medical devices and automotive systems were identified.

Another discussion was about guaranteeing robust performance from systems based on machine learning. This issue is a difficult task and it is growing in importance as many new safety critical applications, such as self-driving cars, are being designed using machine learning techniques. A challenge is to develop reliable methodologies for certifying or designing for robust performance for systems based on machine learning.

Discussions in a third break-out group were centered around the issue of established engineering means for obtaining robustness by design and how to accommodate these in rigorous safety cases or formal proofs of correctness. A finding was that most formal models would currently require rather low-level coding of the dynamic behavior of such mechanisms, thereby requiring them to be re-evaluated on each new design rather than exploiting their guaranteed properties to simplify system analysis, which would be in line with their actual impact on engineering processes.

6.53 Public-Key Cryptography

Organizers: Marc Fischlin, Alexander May, David Pointcheval, and Tal Rabin

Seminar No. 16371

Date: September 11–16, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.46

Participants: Adekunle Oluseyi Afolabi, Fabrice Benhamouda, Johannes A. Buchmann, David Cash, Pooya Farshim, Marc Fischlin, Pierre-Alain Fouque, Vipul Goyal, Iftach Haitner, Dennis Hofheinz, Antoine Joux, Eike Kiltz, Alexander Koch, Tal Malkin, Alexander May, Jörn Müller-Quade, Phong Q. Nguyen, Kenneth G. Paterson, Krzysztof Pietrzak, David Pointcheval, Tal Rabin, Sven Schäge, Suzanna Schmeelk, Dominique Schröder, Jacob Schuldt, Vinod Vaikuntanathan, Hoeteck Wee

Cryptography has turned out to be an invaluable tool for protecting the confidentiality and integrity of digital data. At the same time, cryptography does not yet provide satisfying solutions to all practical scenarios and threats. To accomplish appropriate protection of the data, cryptography needs to address several challenges.

Cryptography has always been a prominent theme within the Dagstuhl Seminar series, with the first meeting about cryptography held in 1993, and subsequent seminars on this topic about every 5 years. In 2007 and 2012 a seminar for the subarea of "Symmetric Cryptography" has been added, inciting us to coin the seminar here "Public-Key Cryptography" for sake of distinction. The public-key branch has been held for the second time, after the first event in 2011.

The seminar brought together 27 scientists in the area of public-key cryptography, including three student researchers who were invited by Dagstuhl to pick a seminar to participate in. The participants came from all over the world, including countries like the US, Great Britain, Israel, France, or Japan. Among the affiliations, Germany lead the number with 9 participants, followed by the US and France with 6 each. The program contained 21 talks, each of 25 to 60 minutes, and a panel discussion about the uneasiness with the current state of our reviewing system, with a free afternoon on Wednesday for social activities and the afternoon on Thursday for collaborations. Before the seminar, we asked the participants to present very recent and ongoing work which, ideally, should not have been published or accepted to publication yet. Most of the participants followed our suggestion and to a large extend the presentations covered topics which have not even been submitted at the time.

The topics of the talks represented the diversity of public-key cryptography. The goal of the seminar was to bring together three challenge areas in cryptography, namely, cryptanalysis and foundations (investigating and evaluating new primitives), optimization (making solutions more efficient), and deployment (designing real-world protocols). As envisioned, the seminar thus has a good mixture of talks from these areas. There were also suggestions to try to co-locate future events of the seminar with other security-related events at Dagstuhl to foster even broader interdisciplinary research. Discussions during and after the talks were lively. It seems as if the goal of stimulating collaborations among these areas has been met. The discussion about the reviewing system has led to some hands-on practices which could be deployed to improve the quality of reviews. This includes incentives such as"Best Reviewer Awards" and teaching students about proper reviewing.

Fig. 6.16
Impressions of a fire drill at Schloss Dagstuhl in summer 2016. Photos courtesy of Michael Wagner.

6.54 Uncertainty Quantification and High Performance Computing

Organizers: Vincent Heuveline, Michael Schick, Clayton Webster, and Peter Zaspel

Seminar No. 16372

Date: September 11–16, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.59

Participants: Peter Benner, Steffen Börm, Nick Dexter, Tim Dodwell, Omar Ghattas, Helmut Harbrecht, Vincent Heuveline, Olivier Le Maitre, Arnaud Legrand, Youssef M. Marzouk, Habib Najm, Dirk Nuyens, Ivan Oseledets, Michael Peters, Dirk Pflüger, Christian Rieger, Michael Schick, Miroslav Stoyanov, Aretha Teckentrup, Clayton Webster, Peter Zaspel, Guannan Zhang

Topics

Uncertainty quantification (UQ) aims at approximating measures for the impact of uncertainties in e.g. simulation parameters or simulation domains. By this way, it is of great importance for both academic research and industrial development. In uncertainty quantification, one distinguishes between classical forward uncertainty propagation and more involved inference, optimization or control problems under uncertainties. Forward uncertainty propagation is concerned with deterministic numerical models for e.g. engineering problems, in which parts of the input data (domain, parameters, ...) might be affected by uncertainties, i.e. they have a random nature. Randomness is usually characterized by random fields that replace the originally deterministic inputs. In Bayesian inference, parameters of a system shall be derived for given measurements. Since the measurements are assumed to be affected by some (stochastic) error, this inference approach tries to derive probabilities under which a given parameter leads to the observed measurements. In some sense, Bayesian inference complements classical inverse problems in a stochastic sense. Other fields of interest for a similar uncertainty analysis are optimization and control.

High performance computing (HPC) is an interdisciplinary research field in computer science, mathematics and engineering. Its aim is to develop hardware, algorithmic approaches and software to solve (usually) mathematically formulated problems on large clusters of interconnected computers. The dominant part of the involved research is done in parallel computing. From a hardware perspective, HPC or parallel computing requires to develop computing technologies that can e.g. solve several problems at the same time at high performance and low power. Moreover, hardware developments in HPC often aim at improving network communication technologies, which are necessary to let a (potentially) large set of computers solve a single problem in a distributed way. From an algorithmic perspective, methods known from numerical mathematics and data processing are adapted such that they can run in a distributed way on different computers. Here, a key notion is (parallel) scalability which describes the ability to improve the performance or throughput of a given method by increasing the number of used computers. Most algorithmic developments shall improve this scalability for numerical methods. Research in software aims at defining appropriate programming models for parallel algorithms, providing efficient management layers for the underlying hardware and implementing the proposed parallel algorithms in real software.

Challenges

In UQ, (partial) differential equations with random data are approximately solved by either intrusive or non-intrusive methods. An intrusive technique simultaneously discretizes stochastic and physical space with the classical example of stochastic Galerkin approaches. This method delivers favorable properties such as small errors with fewer number of equations and potentially small overall run-time. To achieve that, it requires to re-discretize and re-implement existing deterministic PDE solvers. On the other hand, non-intrusive techniques (e.g. (quasi-)Monte Carlo, multi-level Monte Carlo, stochastic collocation, ...) reuse existing solvers / simulation tools and generate a series of deterministic solutions which are used to approximate stochastic moments. It is thereby possible to perform uncertainty quantification analysis even for very complex large-scale applications for which a re-implementation of existing solvers is no option. The non-intrusive approach is connected to a rather extreme computational effort, with at least hundreds, thousands or even more deterministic problems that have to be solved. While a single real-world forward uncertainty propagation problem is already extremely computational intensive, even on a larger parallel computer, inference, optimization and control

under uncertainties often go beyond the limits of currently available parallel computers.

In HPC, we have to distinguish methods that are intrinsically (often also called embarrassingly) parallel and those that have to exchange data to compute a result. That is, embarrassingly parallel algorithms are able to independently compute on completely decoupled parts of a given problem. A prominent example in UQ are Monte-Carlo-type methods. The other extreme are approaches that require to exchange a lot of data in order to solve a given problem. Here, prominent examples are adaptive and multi-level methods in general and stochastic Galerkin methods. Both method types tend to have excellent approximation properties, but require a considerable effort in parallel algorithms to be scalable on parallel computers. Scalability considerations might become even more important on the next generation of the largest parallel computers, which are expected to be available at the beginning of the next decade. These parallel Exascale computers will be able to process on the exaFLOP level, thus they will be able to issue 10 18 floating-point instructions within a second. Technological limitations in chip production will force computing centers to install systems with a parallel processor count which is by orders of magnitude higher than in current systems. Current parallel algorithms might not be prepared for this next step.

The Dagstuhl Seminar on "Uncertainty Quantification and High Performance Computing", brought together experts from UQ and HPC to discuss some of the following challenging questions:

- How can real-world forward uncertainty problems or even inference, control and optimization under uncertainties be made tractable by high performance computing?
- What types of numerical uncertainty quantification approaches are able to scale on current or future parallel computers, without sticking to pure Monte Carlo methods?
- Might adaptivity, model reduction or similar techniques improve existing uncertainty quantification approaches, without breaking their parallel performance?
- Can we efficiently use Exascale computing for large-scale uncertainty quantification problems without being affected by performance, scalability and resilience problems?
- Does current research in uncertainty quantification fit the needs of industrial users? Would industrial users be willing and able to use HPC systems to solve uncertainty quantification problems?

Seminar outcome

Several presentations covered Bayesion inference / inversion (Ghattas, Marzouk, Najm, Peters), where seismology is an extremely computationally expensive problem that can only be solved by the largest parallel computers (Ghattas). While the parallelization is crucial, the numerical methods have to be adapted as well, such that fast convergence is achieved (Ghattas, Marzouk, Peters). The very computationally intensive optimization under uncertainties (Benner) becomes tractable by the use of tensor approximation methods (Benner, Osedelets). Tensor approximation methods as well as hierarchical matrices (Börm, Zaspel) are optimal complexity numerical methods for a series of applications in UQ. However their large-scale parallelization is still subject to research.

A series of talks considered mesh-free approximation methods (Rieger, Teckentrup, Zaspel) with examples in Gaussian process regression (Teckentrup) and kernel-based methods. It was possible to see that these methods have provable error bounds (Rieger, Teckentrup) and can be scaled on parallel computers (Rieger, Zaspel). Moreover these methods even fit well for inference (Teckentrup). Sparse grid techniques were considered as example for classical approximation methods for higher-dimensional problems (Stoyanov, Peters, Harbrecht, Pflüger). Here, recent developments in adaptivity and optimal convergence were discussed. Sparse grid techniques are usually considered in a non-intrusive setting such that parallel scalability is often guaranteed. Compressed sensing promises to reduce the amount of simulations in a non-intrusive framework (Dexter). Quasi-Monte Carlo methods are under investigation for optimal convergence (Nuyens). The latter methods are of high interest for excellent parallel scalability on parallel computers due to the full decoupling of all deterministic PDE solves while keeping convergence orders beyond classical Monte Carlo methods.

Adaptivity leads to strongly improved approximations using the same amount of deterministic PDE solutions (Pflüger, Stoyanov, Webster, …). However, a clear statement on how to parallelize adaptive schemes in an efficient way is still subject to research. The general class of multi-level schemes was also under investigation (Dodwell, Zhang), including but not being limited to multi-level Monte-Carlo and multi-level reduced basis approaches. These methods show excellent convergence properties. However their efficient and scalable parallelization is part of intensive studies, as well.

Performance considerations in the field of HPC (including future parallel computers) have been discussed (Heuveline, Legrand). Performance predictability is necessary to understand scaling behavior of parallel codes on future machines (Legrand). Parallel scalability of (elliptic) stochastic PDEs by domain decomposition has been discussed by LeMaître. His approach allows to increase parallel scalability and might show hints towards resilience.

Industrial applications were considered for the company Bosch (Schick), where intrusive and non-intrusive approaches are under investigation. High performance computing is still subject to discussion in this industrial context. One of the key applications, which is expected to become an industrial-like application, is UQ in medical engineering (Heuveline). Once introduced into the daily work cycle at hospitals, it will soon become a driving technology for our health.

Perspectives

Based on the survey and personal feedback from the invitees, the general consensus is that there is a high interest in deepening the discussions at the border of UQ and HPC. While some answers to the above questions could be given, there is still a lot more to learn, to discuss and to develop. A general wish is therefore to have similar meetings in the future.

Acknowledgements. The organizers would like to express their gratitude to all participants of the Seminar. Special thanks go to the Schloss Dagstuhl team for its extremely friendly support during the preparation phase and for the warm welcome at Schloss Dagstuhl.

6.55 SAT and Interactions

Organizers: Olaf Beyersdorff, Nadia Creignou, Uwe Egly, and Heribert Vollmer

Seminar No. 16381

Date: September 18–23, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.74

Participants: Olaf Beyersdorff, Joshua Blinkhorn, Ilario Bonacina, Sam Buss, Florent Capelli, Leroy Chew, Nadia Creignou, Arnaud Durand, Uwe Egly, Shiguang Feng, John Franco, Nicola Galesi, Anselm Haak, Miki Hermann, Marijn J. H. Heule, Edward A. Hirsch, Kazuo Iwama, Jan Johannsen, Peter Jonsson, Oliver Kullmann, Victor Lagerqvist, Florian Lonsing, Meena Mahajan, Barnaby Martin, Arne Meier, Stefan Mengel, Jakob Nordström, Steffen Reith, Rahul Santhanam, Dominik Scheder, Irena Schindler, Johannes Schmidt, Uwe Schöning, Anil Shukla, Sarah Sigley, Stefan Szeider, Jacobo Torán, Heribert Vollmer, Christoph M. Wintersteiger

Brief Introduction to the Topic

Propositional satisfiability (or Boolean satisfiability) is the problem of determining whether the variables of a Boolean formula can be assigned truth values in such a way as to make the formula true. This satisfiability problem, SAT for short, stands at the crossroads of logic, graph theory, computer science, computer engineering and computational physics. Indeed, many problems originating from one of these fields typically have multiple translations to satisfiability. Unsurprisingly, SAT is of central importance in various areas of computer science including algorithmics, verification, planning, hardware design and artificial intelligence. It can express a wide range of combinatorial problems as well as many real-world ones.

SAT is very significant from a theoretical point of view. Since the Cook-Levin theorem, which identified SAT as the first $\mathbf{NP}$-complete problem, it has become a reference for an enormous variety of complexity statements. The most prominent one is the question "is $\mathbf{P}$ equal to $\mathbf{NP}$?" Proving that SAT is not in $\mathbf{P}$ would answer this question negatively. Restrictions and generalizations of the propositional satisfiability problem play a similar rôle in the examination of other complexity classes and relations among them. In particular, quantified versions of SAT (QSAT, in which Boolean variables are universally or existentially quantified) as well as variants of SAT in which some notion of minimality is involved, provide prototypical complete problems for every level of the polynomial hierarchy.

During the past three decades, an impressive array of diverse techniques from mathematical fields, such as propositional and first-order logic, model theory, Boolean function theory, complexity, combinatorics and probability, has contributed to a better understanding of the SAT problem. Although significant progress has been made on several fronts, most of the central questions remain unsolved so far.

One of the main aims of the Dagstuhl seminar was to bring together researchers from different areas of activity in SAT so that they can communicate state-of-the-art advances and embark on a systematic interaction that will enhance the synergy between the different areas.

Concluding Remarks and Future Plans

The organizers regard the seminar as a great success. Bringing together researchers from different areas of theoretical computer science fostered valuable interactions and led to fruitful discussions. Feedback from the participants was very positive as well. Many attendants expressed their wish for a continuation.

Finally, the organizers wish to express their gratitude toward the Scientific Directorate of the Center for its support of this seminar, and hope to be able to continue this series of seminars on *SAT and Interactions* in the future.

6.56 Foundations of Unsupervised Learning

Organizers: Maria-Florina Balcan, Shai Ben-David, Ruth Urner, and Ulrike von Luxburg

Seminar No. 16382

Date: September 18–23, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.9.94

Participants: Sanjeev Arora, Pranjal Awasthi, Shai Ben-David, Olivier Bousquet, Kamalika Chaudhuri, Sanjoy Dasgupta, Debarghya Ghoshdastidar, Barbara Hammer, Matthias Hein, Christian Hennig, Adam Tauman Kalai, Ravindran Kannan, Samory Kpotufe, Marina Meila, Claire Monteleoni, Lev Reyzin, Heiko Röglin, Sivan Sabato, Melanie Schmidt, Karin Schnass, Hans Ulrich Simon, Christian Sohler, Ingo Steinwart, Ilya Tolstikhin, Ruth Urner, Ulrike von Luxburg, Robert C. Williamson

The success of Machine Learning methods for prediction crucially depends on data preprocessing such as building a suitable feature representation. With the recent explosion of data availability, there is a growing tendency to "let the data speak itself". Thus, unsupervised learning is often employed as a a first step in data analysis to build a good feature representation, but also, more generally, to detect patterns and regularities independently of any specific prediction task. There is a wide rage of tasks frequently performed for these purposes such as representation learning, feature extraction, outlier detection, dimensionality reduction, manifold learning, clustering and latent variable models.

The outcome of such an unsupervised learning step has far reaching effects. The quality of a feature representation will affect the quality of a predictor learned based on this representation, a learned model of the data generating process may lead to conclusions about causal relations, a data mining method applied to a database of people may identify certain groups of individuals as "suspects" (for example of being prone to developing a specific disease or of being likely to commit certain crimes).

However, in contrast to the well-developed theory of supervised learning, currently systematic analysis of unsupervised learning tasks is scarce and our understanding of the subject is rather meager. It is therefore more than timely to put effort into developing solid foundations for unsupervised learning methods. It is important to understand and be able to analyze the validity of conclusions being drawn from them. The goal of this Dagstuhl Seminar was to foster the development of a solid and useful theoretical foundation for unsupervised machine learning tasks.

The seminar hosted academic researchers from the fields of theoretical computer science and statistics as well as some researchers from industry. Bringing together experts from a variety of backgrounds, highlighted the many facets of unsupervised learning. The seminar included a number of technical presentations and discussions about the state of the art of research on statistical and computational analysis of unsupervised learning tasks.

We have held lively discussions concerning the development of objective criteria for the evaluation of unsupervised learning tasks, such as clustering. These converged to a consensus that such universal criteria cannot exist and that there is need to incorporate specific domain expertise to develop different objectives for different intended uses of the clusterings. Consequently, there was a debate concerning ways in which theoretical research could build useful tools for practitioners to assist them in choosing suitable methods for their tasks. One promising direction for progress towards better alignment of algorithmic objectives with application needs is the development of paradigms for interactive algorithms for such unsupervised learning tasks, that is, learning algorithms that incorporate adaptive "queries" to a domain expert. The seminar included presentations and discussions of various frameworks for the development of such active algorithms as well as tools for analysis of their benefits.

We believe, the seminar was a significant step towards further collaborations between different research groups with related but different views on the topic. A very active interchange of ideas took place and participants expressed their satisfactions of having gained new insights into directions of research relevant to their own. As a group, we developed a higher level perspective of the important challenges that research of unsupervised learning is currently facing.

6.57 Programming Language Techniques for Incremental and Reactive Computing

Organizers: Camil Demetrescu, Sebastian Erdweg, Matthew A. Hammer, and Shriram Krishnamurthi

Seminar No. 16402

Date: October 3–7, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.1

Participants: Pramod Bhatotia, Sebastian Burckhardt, Ezgi Cicek, Antony Courtney, Camil Demetrescu, Sebastian Erdweg, Deepak Garg, Philipp Haller, Matthew A. Hammer, Daco Harkes, Kyle Headley, Yit Phang Khoo, Shriram Krishnamurthi, Neel Krishnaswami, Nicholas Labich, Ruy Ley-Wild, Frank McSherry, Mira Mezini, Yaron Minsky, Ryan R. Newton, Marc Pouzet, Guido Salvaneschi, Rohin Shah, R. Benjamin Shapiro, Tamás Szabó, Kanat Tangwongsan

We sought to hold a Dagstuhl Seminar that would bring together programming language (PL) researchers focusing on incremental and reactive computing behavior. The meta-level purpose of this seminar was to take an initial step toward developing a community of experts from the disparate threads of successful research. In that this seminar provoked discussion about common and differing motivations, techniques, and future challenges, this event was successful in starting to cultivate this culture.

Short-term concrete outcomes

Thus far, there are been two concrete outcomes of this seminar:

1. *Wikipedia article outlines and edits*
2. *First Workshop on Incremental Computation (IC) at PLDI 2017*

The full report gives an overview of the event structure of the seminar, and details some of the event's outcomes, including outline brainstorming and Wikipedia editing, and the creation of a new Workshop on Incremental Computing (IC). In addition, the full report gives further background on research in reactive and incremental computing, and further details on the new IC Workshop.

Acknowledgments. We organizers are all thankful to the participants, who all brought a unique insight to the seminar, which in my humble opinion, succeeded in its aims.

Fig. 6.17
Impressions of a fire drill at Schloss Dagstuhl in summer 2016. Photos courtesy of Michael Wagner.

6.58 Algebraic and Combinatorial Methods in Computational Complexity

Organizers: Valentine Kabanets, Thomas Thierauf, Jacobo Tóran, and Christopher Umans

Seminar No. 16411

Date: October 9–14, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.13

Participants: Farid Ablayev, Vikraman Arvind, Markus Bläser, Andrej Bogdanov, Arkadev Chattopadhyay, Samir Datta, Stephen A. Fenner, Michael A. Forbes, Anna Gál, Frederic Green, Rohit Gurjar, Kristoffer Arnsfelt Hansen, William Hoza, Valentine Kabanets, Marek Karpinski, Neeraj Kayal, Pascal Koiran, Swastik Kopparty, Arpita Korwar, Michal Koucký, Andreas Krebs, Sophie Laplante, Nutan Limaye, Meena Mahajan, Pierre McKenzie, Or Meir, David A. Mix Barrington, Ryan O'Donnell, Rafael Oliveira, Chandan Saha, Rahul Santhanam, Shubhangi Saraf, Nitin Saxena, Uwe Schöning, Ronen Shaltiel, Amnon Ta-Shma, Thomas Thierauf, Jacobo Torán, Christopher Umans, Nikolay K. Vereshchagin, Amir Yehudayoff, Jeroen Zuiddam

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 25 talks, most of them lasting about 40 minutes, leaving ample room for discussions. In the following we describe the major topics of discussion in more detail.

Circuit Complexity

This is an area of fundamental importance to Complexity. Circuit Complexity was one of the main topics in the seminar. Still it remains a big challenge to prove strong upper and lower bounds. Also Polynomial Identity Testing (PIT) plays a central role.

The seminar started with a talk by *Steve Fenner*. In a breakthrough result, he showed how to solve the perfect matching problem in bipartite graphs (almost) efficiently in parallel, by circuits of quasi-polynomial size and $O(\log^2 n)$ depth (in quasi-NC). This solves a problem open since more than 30 years. *Rohit Gurjar* showed how to extend the result even further to *linear matroid intersection*, where bipartite perfect matching is a special case of.

Both of the above results can be read as a singularity test of certain symbolic matrices. We had several talks dealing with determining singularity or computing the rank of a symbolic matrix. *Rafael Oliveira* presented an efficient algorithm for the symbolic singularity problem in the *non-commutative* setting. In the *commutative* setting, the complexity is a major open problem. Many other important problems reduce to it. *Markus Bläser* presented an *approximation algorithm* (PTAS) for the rank of a symbolic matrix. Surprisingly, this is achieved with a greedy-algorithm. *Kristoffer Hansen* showed a different kind of approximation for low rank binary matrices.

We have seen some great work on *Polynomial Identity Testing* (PIT) and *c*ircuit lower bounds recently, in particular on depth-3 and depth 4 circuits, and on arithmetic branching programs, which has brought us very close to statements that are known to imply VP $\neq$ VNP, the analogue of the P vs. NP question in the arithmetic world. With respect to PIT, an ambitious goal is to come up with a hitting set construction for a specific model. 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 PIT problem in the *black box* model.

PIT is known to be efficiently solvable by *randomized* algorithms, for example when we consider arithmetic circuits. Things get a bit different when we consider *noncummutative* circuits. Now the standard test cannot be directly applied because the polynomials can have exponential degree, and hence doubly exponentially many monomials. *V. Arvind* presented a randomized polynomial identity test for noncommutative arithmetic circuits for the case when the polynomial has only exponentially many monomials.

One of the most successful methods for proving lower bounds for arithmetic circuits is to consider the dimension of the span of the *partial derivatives* of a polynomial. *Pascal Koiran* considered the complexity of the problem to compute this dimension. He showed that it is #P-hard. It remained open whether the problem is #P-complete.

Another important notion when proving lower bounds is the *algebraic independence* of arithmetic circuits. In 2015, Kumar and Saraf presented lower bounds and hitting sets for a class of depth-4 circuits that have low algebraic rank. Unfortunately, their technique requires base fields of characteristic zero, or at least exponentially large characteristic. *Nitin Saxena* closed this gap and showed how to make the approach work over *every* field.

Michael Forbes showed that lower bounds for certain algebraic circuits imply lower bounds in proof complexity.

Or Meir talked on one of the major open problems in

complexity theory: proving super-polynomial lower bounds on the size of formulas. Karchmer, Raz, and Wigderson suggested an approach to this problem. The *KRW-conjecture* states that the formula complexity of two functions f and g roughly adds up when we consider the composed function $g \circ f$. They showed that the conjecture implies super-polynomial formula lower bounds In his talk, Or Meir did a step to prove the conjecture: he proved a special case, namely when f is the parity-function. His proof uses techniques from communication complexity.

Valiant introduced the arithmetic analogue of classes P and NP. Very roughly, the class VP contains all multivariate polynomials that can be computed (non-uniformly) by polynomial-size arithmetic circuits, and the class VNP contains all multivariate polynomials that have coefficients computable by VP-circuits. The question whether VP is different from VNP plays the role of the P-NP question in algebraic complexity theory. Valiant showed that the permanent is complete for VNP. But for VP, only artificially constructed functions were known to be complete. In her talk, *Meena Mahajan* described several polynomial families complete for VP and for VNP, based on the notion of graph homomorphism polynomials.

Complexity

Since the famous AKS-primality test, prime numbers can be recognized efficiently. The *construction* of prime numbers is still a challenging task. The best known deterministic algorithm have only exponential running time. *Rahul Santhanam* presented a randomized subexponential time algorithm that outputs primes, and only primes, with high probability, and moreover, the output is mostly the same prime. This is called a *zero-error pseudo-deterministic* algorithm.

Since the famous Isolation Lemma of Mulmuley, Vazirani, Vazirani, researchers recognized the power of isolation. For example, the bipartite perfect matching and the matroid intersection algorithms mentioned above, both rely on isolating a minimum weight solution, *Nutan Limaye* studied the problem of isolating an s-t-path in a directed graph. She proved that a randomized logspace algorithm that isolates such a path can be used to show NL $\subseteq$ L/poly.

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 simulate randomized algorithms with only polynomial overhead. The polynomial overhead is fine for algorithms running in polynomial time. However, in case of subexponential randomized algorithms, this overhead makes the resulting deterministic algorithm more or less useless. *Ronen Shaltiel* showed how to overcome this problem by achieving a more modest overhead. He needs, however, stronger lower bounds to begin with. Further talks on pseudo-random generators and randomness extractors were given by *Amnon Ta-Shma* and *William Hoza*.

Chris Umans gave an evening talk presenting a recent breakthrough in additive combinatorics, the resolution of the so-called *cap-set conjecture* by Ellenberg and Gijswijt. This result has implications for the Cohn-Umans group-theoretic approach for matrix multiplication, and elsewhere in Complexity.

Coding Theory

Error-correcting 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. *Shubhangi Saraf* gave a talk on locally-correctable and locally-testable codes. *Swastik Kopparty* generalized the well known decoding algorithm for Reed-Solomon codes to higher dimensions. He presented an efficient algorithm to decode Reed-Muller codes when the evaluation points are an arbitrary product set S^m, for some m, when S is larger than the degree of the polynomials.

Quantum Complexity

Complexity issues arising in the context of quantum computation are an important area in complexity theory since several decades. In the seminar, we had two talks related to quantum complexity. *Farid Ablayev* talked bout the notion of quantum hash function and how to construct such functions. He also explained some of its applications for constructing quantum message authentication codes. *Ryan O'Donnel* explained about the *quantum tomography problem* and how this special case of *quantum spectrum estimation* can be solved combinatorially by understanding certain statistics of random words.

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 seminar. We look forward to our next meeting!

6.59 Automated Algorithm Selection and Configuration

Organizers: Holger H. Hoos, Frank Neumann, and Heike Trautmann

Seminar No. 16412

Date: October 9–14, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.33

Participants: Anne Auger, Thomas Bäck, Aymeric Blot, Benjamin Doerr, Carola Doerr, Michael Emmerich, Carlos M. Fonseca, Tobias Friedrich, Marcus Gallagher, Wanru Gao, Carlos Ignacio Hernández Castellanos, Holger H. Hoos, Frank Hutter, Laetitia Jourdan, Pascal Kerschke, Lars Kotthoff, Kevin Leyton-Brown, Marius Lindauer, Manuel López-Ibánez, Andres Munoz Acosta, Nysret Musliu, Samadhi Nethmini Nallaperuma, Frank Neumann, Mike Preuß, Günter Rudolph, Horst Samulowitz, Marc Schoenauer, Meinolf Sellmann, Thomas Stützle, Heike Trautmann, Joaquin Vanschoren, Markus Wagner, Hao Wang, Simon Wessing

The importance of high-performance algorithms, in particular for solving $\mathcal{NP}$-hard optimisation and decision problems, cannot be underestimated. Achievements in this area have substantial impact in sectors such as manufacturing, logistics, healthcare, finance, agriculture and energy systems – all of strategic importance to modern societies.

The development of effective automated algorithm selection and configuration techniques has been one of the major success stories in the area of empirical algorithmics in recent years. Building on a wide range of algorithmic approaches for problems such as propositional satisfiability (SAT) and mixed integer programming (MIP), these methods permit the selection of appropriate algorithms based on efficiently computable characteristic of a problem instance to be solved (algorithm selection) and the automatic determination of performance optimising parameter settings (algorithm configuration). In both cases, statistical models that enable performance predictions for previously unseen problem instances or parameter settings play a key enabling role; additionally, these models have other important uses, e.g., in load scheduling and distribution on large computer clusters.

The reach of those methods is illustrated by the fact that they have defined the state of the art in solving SAT, arguably the most prominent $\mathcal{NP}$-complete decision problem, for a decade (as witnessed by the results from the international SAT solver competitions (http://www.satcompetition.org), and more recently have been demonstrated to have the potential to achieve significant improvements over the long-standing state of the art in solving the TSP, one of the most widely studied $\mathcal{NP}$-hard optimisation problems [1]. Further very encouraging results have been achieved in recent years for continuous optimisation, AI planning and mixed integer programming problems.

The goal of the seminar was to foster research on algorithm selection and configuration, as well as on the underlying performance prediction methods, by bringing together researchers from the areas of artificial intelligence, theoretical computer science and machine learning in order to extend current studies to a much broader class of problems and build up the theoretical foundations of this important research area. On the foundational side, the seminar aimed at bridging the gap between experiments and theory in feature-based algorithm (runtime) analysis. In particular, we began investigating how mathematical and theoretical analyses can contribute to the experimentally driven research area of algorithm selection and configuration. We expect that studies following this initial exploration will bring together two of the currently most successful approaches for analysing heuristic search algorithms and ultimately achieve substantial impact in academic and industrial applications of algorithm configuration and selection techniques. Furthermore, we placed an emphasis on investigating automated algorithm selection and configuration approaches for multiobjective optimisation problems – an important, but largely unexplored area of investigation.

Background and Challenges: Algorithm Selection and Configuration for Combinatorial Problems

The design of algorithms for combinatorial optimisation and decision problems plays a key role in theoretical computer science as well as in applied algorithmics. These problems are frequently tackled using heuristic methods that perform extremely well on different classes of benchmark instances but usually do not have rigorous performance guarantees. Algorithm selection and configuration techniques have been applied to some of the most prominent $\mathcal{NP}$-hard combinatorial optimisation and decision problems, such as propositional satisfiability (SAT) and the travelling salesman problem (TSP).

Algorithm selection for SAT has first been explored in the seminal work on SATzilla [2–4], which was initially based on linear and ridge regression methods for performance prediction,

but later moved to more sophisticated models based on cost-sensitive random forest classification [5]. Other successful methods use clustering techniques to identify the algorithm to be run on a given instance [6, 7]. As clearly evident from the results of SAT competitions, which are regularly held to assess and document the state of the art in SAT solving, automated algorithm selection procedures effectively leverage the complementary strengths of different high-performance solvers and thus achieve substantial improvements over the best individual solvers [5].

As heuristic search algorithms often have numerous parameters that influence their performance, one of the classical questions is how to set parameters to optimise performance on a given class of instances. This per-set algorithm configuration problems can be solved using stochastic local search and model-based optimisation techniques [8–10], as well as racing techniques [11,12], and these configuration methods have been demonstrated to yield substantial performance improvements to state-of-the-art algorithms for SAT, TSP, MIP, AI planning and several other problems [13–16]. Algorithm configuration techniques are now routinely used for optimising the empirical performance of solvers for a wide range of problems in artificial intelligence, operations research and many application areas (see, e.g., [17, 18]).

Initial work on combining algorithm selection and configuration techniques has shown significant promise [19, 20]; such combinations allow configuring algorithms on a per-instance basis [6, 7] and configuring algorithm selection methods (which themselves make use of many heuristic design choices) [21]. However, we see much room for further work along these lines. Other challenges concern the automated selection and configuration of mechanisms that adapt parameter settings while an algorithm is running and the configuration of algorithms for optimised scaling behaviour. Finally, a better theoretical foundation of algorithm selection and configuration approaches is desired and necessary. Initial steps into this direction were an important goal of this Dagstuhl seminar. In the following, we motivate and outline some of the challenges addressed in the course of the seminar.

Background and Challenges: Algorithm Selection for Continuous Black-Box Optimisation

Black-box function optimisation is a basic, yet intensely studied model for general optimisation tasks, where all optimisation parameters are real-valued. Work in this area has important practical applications in parameter and design optimisation and has also inspired some of the most successful general-purpose algorithm configuration techniques currently available [9].

Despite many years of research in metaheuristics, especially evolutionary algorithms, aimed at optimising black-box functions effectively, it is currently hardly possible to automatically determine a good optimisation algorithm for a given black-box function, even if some of its features are known. In single-objective (SO) black-box optimisation, it is therefore of considerable interest to derive rules for determining how problem properties influence algorithm performance as well as for grouping test problems into classes for which similar performance of the optimisation algorithms can be observed. Recent benchmarking experiments [22, 23] provide at best high-level guidelines for choosing a suitable algorithm type based on basic features that are known a priori, such as the number of dimensions of the given problem. However, the preference rules for algorithm selection thus obtained are very imprecise, and even for slight algorithm or problem variations, the resulting performance-induced ordering of different algorithms can change dramatically.

Exploratory Landscape Analysis (ELA, [24]) aims at improving this situation by deriving cheaply computable problem features based on which models relating features to algorithm performance can be constructed using benchmark experiments. The final goal is an accurate prediction of the best suited algorithm for an arbitrary optimisation problem based on the computed features. The concept is not entirely knew; however, earlier approaches, such as fitness distance correlation (FDC) [25], have not been completely convincing.

A first idea to employ high-level (human expert designed) features, such as separability and modality, to characterize optimisation problems in an ELA context [26] was therefore refined by also integrating low-level features – e.g., based on convexity or the behaviour of local search procedures [27]. These effectively computable low-level features can be chosen from a wide range of easy to measure statistical properties. Suitably determined combinations of such features are expected to provide sufficient information to enable successful algorithm selection. Following recent results [27], this process is not necessarily costly in terms of function evaluations required for feature computation.

Additional, conceptually similar features were introduced in [28–31]. In [32], a representative portfolio of four optimisation algorithms was constructed from the complete list of BBOB 2009/2010 candidates. Based on the low-level features a sufficiently accurate prediction of the best suited algorithm within the portfolio for each function was achieved. Recently, the feature set was extended based on the cell mapping concept in [33] by which a finite subdivision of the domain in terms of hypercubes is constructed. Most recently, the ELA approach, extended by several specific features, was successfully used to experimentally detect funnel structured landscapes in unknown black-box optimisation problems [34]. As it can be assumed that this information can be efficiently exploited to speed up the optimisation process, we expect ELA to contribute importantly to automated algorithm selection in single-objective black-box optimisation. (See [35] for a survey of related work.)

One major challenge in this area is the construction of a suitable algorithm portfolio together with an algorithm selection mechanism for unknown instances that generalises well to practical applications. For this purpose, suitable benchmark sets have to be derived and the costs of feature computations have to be kept as small as possible. Furthermore, theoretical foundation of the approaches is desired and necessary. The seminar aimed to make first steps in this direction.

Special focus: Algorithm selection for multiobjective optimisation

Some of the most challenging real-world problems involve the systematic and simultaneous optimisation of multiple conflicting objective functions – for example, maximising product quality and manufacturing efficiency, while minimising production time and material waste. To solve such problems, a large number of multiobjective optimisation (MOO) algorithms has been reported. Like single-objective (SO) algorithms, new MOO algorithms are claimed to outperform others by comparing the results over a limited set of test problems. Knowles et al. [36] started working on systematically deriving performance measures for EMOA and evaluating EMOA performance. Mersmann et al. [37] recently derived a systematic benchmarking framework according to similar work of [38] on benchmarking classification algorithms.

However, it is unlikely that any algorithm would outperform all others on a broader set of problems, and it is possible that the algorithm fails miserably on some of them. These results go usually unreported, leaving the algorithm's limitations unknown. This knowledge is crucial to avoid deployment disasters, gain theoretical insights to improve algorithm design, and ensure that algorithm performance is robustly described. Therefore, we

see much value in the development of an algorithm selection and configuration framework for multiobjective optimisation. Successfully selecting the proper optimization algorithm for a multi-objective problem depends on detecting different problem characteristics, one of which is the multimodality of the induced landscape. In recent work [39], formal definitions were introduced for multimodality in multi-objective optimization problems in order to generalize the ELA framework to multi-objective optimization.

Significant progress has been made on single-objective (SO) problems of combinatorial and continuous nature as discussed above. However, these ideas are yet to be applied to the important class of MOO problems. We see five major avenues of exploration: (1) analysis on what makes MOO problems difficult; (2) design of features to numerically characterize MOO problems; (3) identification and visualization of strengths and weaknesses of state-of-the-art MOO algorithms; (4) methodology to assist the algorithm selection and configuration on (possibly expensive) real-world problems; (5) methodology to assist the design of tailored algorithms for real-world problems. An important aim of the seminar was to facilitate discussion of these directions.

Seminar structure and outcomes

The seminar was structured to balance short invited presentations with group breakout sessions and a generous amount of time set aside for informal discussions and spontaneously organised working groups at a ratio of about 2:1:1. Based on feedback obtained during and after the event, this structure worked well in fostering a vibrant atmosphere of intense and fruitful exchange and discussion. Presenters very successfully introduced important ideas, outlined recent results and open challenges, and facilitated lively discussion that provided much additional value. The afternoon group breakout sessions were particularly effective in addressing the challenges previously outlined as well as additional topics of interest that emerged during the seminar – thanks to the preparation and moderation by the session organisers as well as the lively participation of the attendees.

While it would be unreasonable to expect to exhaustively or conclusively address the substantial research challenges that inspired us to organise this Dagstuhl seminar, we believe that very significant progress has been achieved. As importantly, we feel that through this week-long event, an invaluable sharing of perspective and ideas has taken place, whose beneficial effects on the algorithm selection and configuration community and its work we hope to be felt for years to come. The following presentation abstracts and session summaries provided by the participants reflect the richness and depth of the scientific exchange facilitated by the seminar.

As organisers, we very much enjoyed working with presenters and session organisers, who greatly contributed to the success of the seminar, as did everyone who participated. Our thanks also go to the local team at Schloss Dagstuhl, who provided outstanding organisational support and a uniquely inspiring environment.

References

1 L. Kotthoff, P. Kerschke, H. Hoos, and H. Trautmann. Improving the state of the art in inexact TSP solving using per-instance algorithm selection. In C. Dhaenens, L. Jourdan, and M.-E. Marmion, editors, *Learning and Intelligent Optimization, 9^{th} International Conference*, pp. 202–217, 2015. Springer International Publishing.

2 E. Nudelman, K. Leyton-Brown, H. H Hoos, A. Devkar, and Y. Shoham. Understanding random SAT: Beyond the clauses-to-variables ratio. In *Principles and Practice of Constraint Programming–CP 2004*, pp. 438–452. Springer Berlin Heidelberg, 2004.

3 E. Nudelman, K. Leyton-Brown, A. Devkar, Y. Shoham, and H. Hoos. Satzilla: An algorithm portfolio for SAT. *Solver description, SAT competition*, 2004, 2004.

4 L. Xu, F. Hutter, H. Hoos, and K. Leyton-Brown. SATzilla: Portfolio-based algorithm selection for SAT. *Journal of Artificial Intelligence Research*, 32:565–606, 2008.

5 L. Xu, F. Hutter, H.H. Hoos, and K. Leyton-Brown. Evaluating component solver contributions to portfolio-based algorithm selectors. In *Theory and Applications of Satisfiability Testing–SAT 2012*, pp. 228–241. Springer Berlin Heidelberg, 2012.

6 S. Kadioglu, Y. Malitsky, M. Sellmann, and K. Tierney. Isac-instance-specific algorithm configuration. *ECAI*, 215:751–756, 2010.

7 Y. Malitsky, A. Sabharwal, H. Samulowitz, and M. Sellmann. Algorithm portfolios based on cost-sensitive hierarchical clustering. In *Proc. of the 23rd Int'l Joint Conf. on Artificial Intelligence*, pp. 608–614. AAAI Press, 2013.

8 F. Hutter, H.H. Hoos, K. Leyton-Brown, and K.P. Murphy. An experimental investigation of model-based parameter optimisation: Spo and beyond. In *GECCO'09: Proc. of the 11th Annual Conf. on Genetic and evolutionary computation*, pp. 271–278, New York, NY, USA, 2009. ACM.

9 F. Hutter, H.H. Hoos, and K. Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *Learning and Intelligent Optimization*, pp. 507–523. Springer Berlin Heidelberg, 2011.

10 C. Ansótegui, M. Sellmann, and K. Tierney. A gender-based genetic algorithm for the automatic configuration of algorithms. *Principles and Practice of Constraint Programming-CP 2009*, pp. 142–157, 2009.

11 M. Birattari, T. Stützle, L. Paquete, and K. Varrentrapp. A racing algorithm for configuring metaheuristics. In *GECCO '02: Proc. of the Genetic and Evolutionary Computation Conf.*, pp. 11–18, 2002.

12 M. Birattari, Z. Yuan, P. Balaprakash, and T. Stützle. F-race and iterated F-race: An overview. In T. Bartz-Beielstein, M. Chiarandini, L. Paquete, and M. Preuss, editors, *Empirical Methods for the Analysis of Optimization Algorithms*. Springer, 2010.

13 F. Hutter, M.T. Lindauer, A. Balint, S. Bayless, H.H. Hoos, and K. Leyton-Brown. The configurable SAT solver challenge (CSSC). *Artificial Intelligence*, Accepted for publication, 2015.

14 J. Styles and H. Hoos. Ordered racing protocols for automatically configuring algorithms for scaling performance. In *Proc. of the 15th Conf. on Genetic and Evolutionary Computation (GECCO-13)*, pp. 551–558. ACM, 2013.

15 F. Hutter, H.H. Hoos, and K. Leyton-Brown. Automated configuration of mixed integer programming solvers.

In *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, pp. 186–202. Springer Berlin Heidelberg, 2010.

16 M. Vallati, C. Fawcett, A.E. Gerevini, H.H. Hoos, and A. Saetti. Automatic generation of efficient domain-optimized planners from generic parametrized planners. In *Sixth Annual Symp. on Combinatorial Search (SoCS-13)*, pp. 184–192, 2013.

17 P. Lengauer and H. Mössenböck. The taming of the shrew: Increasing performance by automatic parameter tuning for java garbage collectors. In *Proc. of the 5th ACM/SPEC Int'l Conf. on Performance Engineering*, ICPE'14, pp. 111–122, New York, NY, USA, 2014. ACM.

18 J.P. Dickerson, A.D. Procaccia, and T. Sandholm. Dynamic matching via weighted myopia with application to kidney exchange. In *Proc. of the 28th National Conf. on Artificial Intelligence (AAAI-14)*, pp. 1340–1346, 2012.

19 L. Xu, H.H. Hoos, and K. Leyton-Brown. Hydra: Automatically configuring algorithms for portfolio-based selection. *Proc. of the 24th Conf. on Artificial Intelligence (AAAI-10)*, 10:210–216, 2010.

20 L. Xu, F. Hutter, H.H. Hoos, and K. Leyton-Brown. Hydra-MIP: Automated algorithm configuration and selection for mixed integer programming. *RCRA workshop on experimental evaluation of algorithms for solving problems with combinatorial explosion at the Int'l Joint Conf. on Artificial Intelligence (IJCAI)*, pp. 16–30, 2011.

21 M. Lindauer, H. Hoos, F. Hutter and T. Schaub: AutoFolio: An Automatically Configured Algorithm Selector. J. Artif. Intell. Res. (JAIR) 53:745-778, 2015.

22 N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization benchmarking 2009: Experimental setup. Technical Report RR-6828, INRIA, 2009.

23 N. Hansen, A. Auger, S. Finck, and R. Ros. Real-parameter black-box optimization benchmarking 2010: Experimental setup. Technical Report RR-7215, INRIA, 2010.

24 O. Mersmann, B. Bischl, H. Trautmann, M. Preuss, C. Weihs, and G. Rudolph. Exploratory landscape analysis. In *Proc. of the 13th Annual Conf. on Genetic and Evolutionary Computation*, GECCO'11, pp. 829–836, New York, NY, USA, 2011. ACM.

25 T. Jones and S. Forrest. Fitness distance correlation as a measure of problem difficulty for genetic algorithms. In *Proceedings of the 6th Ijnt'l Conf. on Genetic Algorithms*, pp. 184–192. Morgan Kaufmann, 1995.

26 M. Preuss and Th. Bartz-Beielstein. Experimental analysis of optimization algorithms: Tuning and beyond. In Y. Borenstein and A. Moraglio, editors, *Theory and Principled Methods for Designing Metaheuristics*. Springer, 2011.

27 O. Mersmann, M. Preuss, H. Trautmann, B. Bischl, and C. Weihs. Analyzing the BBOB results by means of benchmarking concepts. *Evolutionary Computation Journal*, 23(1):161–185, 2015.

28 M. A. Muñoz, M. Kirley, and S. K. Halgamuge. A meta-learning prediction model of algorithm performance for continuous optimization problems. In C. A. Coello Coello et al., editors, *Parallel Problem Solving from Nature – PPSN XII*, volume 7491 of *LNCS*, pp. 226–235. Springer, 2012.

29 T. Abell, Y. Malitsky, and K. Tierney. Features for exploiting black-box optimization problem structure. In Giuseppe Nicosia and Panos Pardalos, editors, *Learning and Intelligent Optimization*, LNCS, pp. 30–36. Springer, 2013.

30 R. Morgan and M. Gallagher. Using landscape topology to compare continuous metaheuristics: A framework and case study on EDAs and ridge structure. *Evolutionary Computation*, 20(2):277–299, 2012.

31 M.A. Munoz, M. Kirley, and S.K. Halgamuge. Exploratory landscape analysis of continuous space optimization problems using information content. *Evolutionary Computation, IEEE Transactions on*, 19(1):74–87, Feb 2015.

32 B. Bischl, O. Mersmann, H. Trautmann, and M. Preuss. Algorithm selection based on exploratory landscape analysis and cost-sensitive learning. In *Proc. of the 14th Annual Conf. on Genetic and Evolutionary Computation*, GECCO '12, pp. 313–320. ACM, 2012.

33 P. Kerschke, M. Preuss, C. Hernández, O. Schütze, J. Sun, C. Grimme, G. Rudolph, B. Bischl, and H. Trautmann. Cell mapping techniques for exploratory landscape analysis. In A. Tantar et al., editors, *EVOLVE — A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation V*, volume 288 of *Advances in Intelligent Systems and Computing*, pp. 115–131. Springer, 2014.

34 P. Kerschke, M. Preuss, S. Wessing, and H. Trautmann. Detecting funnel structures by means of exploratory landscape analysis. In *Proc. of the*, pp. 265–272, New York, NY, USA, 2015. ACM.

35 M. A. Muñoz, Y. Sun, M. Kirley, and S. K. Halgamuge. Algorithm selection for black-box continuous optimization problems: A survey on methods and challenges. *Information Sciences*, 317:224 – 245, 2015.

36 J. Knowles, L. Thiele, and E. Zitzler. A Tutorial on the Performance Assessment of Stochastic Multiobjective Optimizers. TIK Report 214, Computer Engineering and Networks Laboratory, ETH Zurich, 2006.

37 O. Mersmann, H. Trautmann, B. Naujoks, and C. Weihs. Benchmarking evolutionary multiobjective optimization algorithms. In *IEEE Congress on Evolutionary Computation*, pp. 1–8. IEEE, 2010.

38 K. Hornik and D. Meyer. Deriving consensus rankings from benchmarking experiments. In R. Decker and H.-J. Lenz, editors, *Advances in Data Analysis (Proc. of the 30th Ann. Conf. of the Gesellschaft für Klassifikation*, pp. 163–170. Springer, Berlin, 2007.

39 P. Kerschke, H. Wang, M. Preuss, C. Grimme, A. Deutz, H. Trautmann and M. Emmerich. Towards Analyzing Multimodality of Multiobjective Landscapes. In *Proceedings of the 14^{th} Int'l Conf. on Parallel Problem Solving from Nature (PPSN XIV)*, pp. 962–972. Lecture Notes in Computer Science, Springer, 2016.

6.60 Universality of Proofs

Organizers: Gilles Dowek, Catherine Dubois, Brigitte Pientka, and Florian Rabe

Seminar No. 16421

Date: October 16–21, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.75

Participants: Andreas Martin Abel, Jesús María Aransay Azofra, Christoph Benzmüller, Jasmin Christian Blanchette, Frédéric Blanqui, Peter Brottveit Bock, Venanzio Capretta, Benjamin Delaware, Gilles Dowek, Catherine Dubois, William M. Farmer, Amy Felty, Thibault Gauthier, Frédéric Gilbert, Georges Gonthier, Stéphane Graham-Lengrand, Hugo Herbelin, Olivier Hermant, Guy Katz, Chantal Keller, Michael Kohlhase, Ramana Kumar, Dale Miller, Alberto Momigliano, César A. Muñoz, Adam Naumowicz, Brigitte Pientka, Elaine Pimentel, Florian Rabe, Claudio Sacerdoti Coen, Ivan Scagnetto, Gert Smolka, René Thiemann, Josef Urban, Laurent Voisin, Freek Wiedijk, Bruno Woltzenlogel Paleo

Proof systems are software systems that allow us to build formal proofs, either interactively or automatically, and to check the correctness of such proofs. Building such a formal proof is always a difficult task – for instance the Feit-Thompson odd order theorem, the CompCert verified C compiler, the seL4 verified operating system micro-kernel, and the proof of the Kepler conjecture required several years with a medium to large team of developers to be completed. Moreover, the fact that each of these proofs is formalized in a specific logic and the language of a specific proof tool is a severe limitation to its dissemination within the community of mathematicians and computer scientists. Compared to many other branches of computer science, for instance software engineering, we are still very far from having off-the-shelf and ready-to-use components, "proving in the large" techniques, and interoperability of theory and systems. However, several teams around the world are working on this issue and partial solutions have been proposed including point-to-point translations, proof standards, and logical frameworks. Yet, a lot still remains to be done as there is currently no overarching general foundation and methodology.

This seminar has been organized to bring together researchers from different communities, such as automated proving, interactive proving and SAT/SMT solving as well as from logic, proof engineering, program verification and formal mathematics. An essential goal has been to form a community around these issues in order to learn about and reconcile these different approaches. This will allow us to develop a common objective and framework for proof developments that support the communication, reuse, and interoperability of proofs.

The program of the seminar included introductions to different methods and techniques, the definition of precise objectives, and the description of recent achievements and current trends. It consisted of 30 contributed talks from experts on the above topics and six breakout sessions on major problems: theory graph – based reasoning, benchmarks, conflicting logics and system designs, proof certificates, design of a universal library of elementary mathematics, and a standard for system integration and proof interchange. The contributed talks took place in the morning, and two parallel breakout sessions each took place on Monday, Tuesday and Thursday afternoon, followed by plenary discussions organized by each session's moderator.

The organizers would like to thank the Dagstuhl team and all the participants for making this first seminar a success and, hopefully, an event to be repeated.

Fig. 6.18
"Morning run at Dagstuhl" Twitter post by 16412 Dagstuhl Seminar participant Joaquin Vanschoren. https://twitter.com/joavanschoren/status/786816337360351232. Photo courtesy of Joaquin Vanschoren.

6.61 Computation over Compressed Structured Data

Organizers: Philip Bille, Markus Lohrey, Sebastian Maneth, and Gonzalo Navarro

Seminar No. 16431

Date: October 23–28, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.99

Participants: Diego Arroyuelo, Hideo Bannai, Djamal Belazzougui, Philip Bille, Stefan Böttcher, Rodrigo Cánovas, Patrick Hagge Cording, Héctor Ferrada, Johannes Fischer, Travis Gagie, Adrià Gascón, Pawel Gawrychowski, Simon Gog, Inge Li Gørtz, Cecilia Hernández Rivas, Danny Hucke, Tomohiro I, Shunsuke Inenaga, Artur Jez, Juha Kärkkäinen, Susana Ladra González, Markus Lohrey, Sebastian Maneth, Ian Munro, Gonzalo Navarro, Yakov Nekrich, Patrick K. Nicholson, Alberto Ordóñez, Fabian Peternek, Nicola Prezza, Rajeev Raman, Wojciech Rytter, Hiroshi Sakamoto, Srinivasa Rao Satti, Markus Schmid, Manfred Schmidt-Schauss, Rahul Shah, Ayumi Shinohara, Tatiana Starikovskaya, Alexander Tiskin, Rossano Venturini

The Dagstuhl Seminar "Computation over Compressed Structured Data" took place from October 23rd to 28th, 2016. The aim was to bring together researchers from various research directions in data compression, indexing for compressed data, and algorithms for compressed data. Compression, and the ability to index and compute directly over compressed data, is a topic that is gaining importance as digitally stored data volumes are increasing at unprecedented speeds. In particular, the seminar focused on techniques for compressed *structured data*, i.e., string, trees, and graphs, where compression schemes can exploit complex structural properties to achieve strong compression ratios.

The seminar was meant to inspire the exchange of theoretical results and practical requirements related to compression of structured data, indexing, and algorithms for compressed structured data. The following specific points were addressed.

Encoding Data Structures. The goal is to encode data structures with the minimal number of bits needed to support only the desired operations, which is also called the effective entropy. The best known example of such an encoding is the $2n$-bit structure that answers range minimum queries on a permutation of $[1, n]$, whose ordinary entropy is $n \log(n)$ bits. Determining the effective entropy and designing encodings that reach the effective entropy leads to challenging research problems in enumerative combinatorics, information theory, and data structures.

Computation-Friendly Compression. Existing state-of-the-art compression schemes encode data by extensive and convoluted references between pieces of information. This leads to strong compression guarantees, but often makes it difficult to efficiently perform compressed computation. Recent developments have moved towards designing more computation-friendly compression schemes that achieve both strong compression and allow for efficient computation. Precise bounds on the worst-case compression of these schemes are mostly missing so far.

Repetitive Text Collections. Many of the largest sequence collections that are arising are formed by many documents that are very similar to each other. Typical examples arise from version control systems, collaborative editing systems (wiki), or sequencing of genomes from the same species. Statistical-compression does not exploit this redundancy. Recently, compressed indexes based on grammar-based compressors have been developed for repetitive text collections. They achieve a considerable compression, but on the downside operations are much slower.

Recompression. Recompression is a new technique that was successfully applied for the approximation of smallest string grammars and to solve several algorithmic problems on grammar-compressed strings. Recently, recompression has been extended from strings to trees. The long list of problems that were solved in a relatively short period using recompression indicates that there exist more applications of recompression.

Graph Compression. A lot of recent work deals with succinct data structures for graphs and with graph compression, in particular for web and network graphs. At the same time, simple queries such as in- and out-neighbors can be executed efficiently on these structures. There is a wide range of important open problems and future work. For instance, there is a strong need to support more complex graph queries, like for instance regular path queries, on compressed graphs.

The seminar fully satisfied our expectations. The 41 participants from 16 countries (Algiers, Canada, Chile, Denmark, Finland, France, Germany, Great Britain, Ireland, Italy, Israel, Japan, Korea, Poland, Spain, and US) had been invited by the organizers to give survey talks about their recent research related to the topic of the seminar. The talks covered topics related to

compression (e.g., grammar-based compression of string, trees, and graphs, Lempel-Ziv compression), indexing of compressed data (e.g., set-intersection, longest common extensions, labeling schemes), algorithms on compressed data (e.g., streaming, regular expression matching, parameterized matching) and covered a wide range of applications including databases, WWW, and bioinformatics. Most talks were followed by lively discussions. Smaller groups formed naturally which continued these discussions later.

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

6.62 Adaptive Isolation for Predictability and Security

Organizers: Tulika Mitra, Jürgen Teich, and Lothar Thiele

Seminar No. 16441

Date: October 30 to November 4, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.120

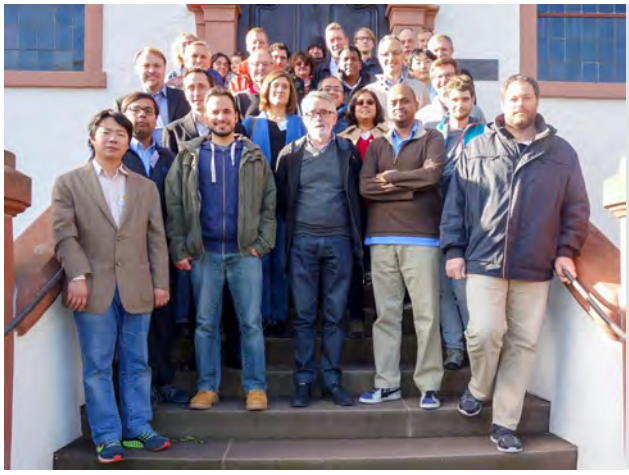

Participants: Davide Bertozzi, Björn B. Brandenburg, David Broman, Samarjit Chakraborty, Sudipta Chattopadhyay, Jian-Jia Chen, Albert Cohen, Ruan de Clercq, Heiko Falk, Felix Freiling, Johannes Götzfried, Gernot Heiser, Andreas Herkersdorf, Karine Heydemann, Patrick Koeberl, Pieter Maene, Claire Maiza, Peter Marwedel, Tulika Mitra, Sibin Mohan, Frank Mueller, Sri Parameswaran, Jan Reineke, Christine Rochange, Zoran Salcic, Patrick Schaumont, Martin Schoeberl, Wolfgang Schröder-Preikschat, Takeshi Sugawara, Jürgen Teich, Lothar Thiele, Theo Ungerer, Reinhard von Hanxleden, Stefan Wildermann, Reinhard Wilhelm

Semiconductor industry has shifted from processor clock speed optimization (having reached its physical limits) to parallel and heterogeneous many-core architectures. Indeed, the continuous technological scaling enables today the integration of hundred and more cores and, thus, enormous parallel processing capabilities. Whereas higher (average) performance has been and still is the major driver for any MPSoC platform design, there is a huge hesitation and fear to install such platforms in embedded systems that require predictable (boundable) guarantees of non-functional properties of execution rather than average properties for a mix of applications. Moreover, it may be observed that in an embedded system, each application running on a platform typically a) requires different qualities to be satisfied. For example, one application might demand for authentification, thus requiring the guarantee of unmodified data and program but have no requirements on speed of execution. Another application might rather require the execution to meet a set of real-time properties such as a deadline or a target data rate. To give an example, consider a driver assistance video processing application in a car that must detect obstacles in front of the car fast enough so to activate the brake system in a timely manner. It must therefore be possible to enforce a set of non-functional qualities of execution on a multi-core platform on a per-application/job basis. b) The above requirements on execution qualities may even change over time or during the program execution of a single application or being dependent on user or environmental settings. For example, one user might not care about sending or distributing personal information over the communication interfaces of a mobile phone whereas another one cares a lot, even in the presence of side channels.

Unfortunately, the way MPSoCs are built and programmed today, the embedded system engineers often experience even worse execution qualities than in the single core case, the reason being the sharing of resources such as cores, buses and/or memory in an unpredictable way. Another obstacle for a successful deployment of multi-core technology in embedded systems is the rather unmanageable complexity. This holds particularly true for the analysis complexity of a system for predictable execution qualities at either compile-time or run-time or using hybrid analysis techniques. The complexity is caused here by an abundant number of resources on the MPSoC and the increasing possibilities of interference created by their concurrent execution and multiple layers of software controlling program executions on a platform. Such layers are often designed for contradictory goals. For example, the power management firmware of an MPSoC may be designed to reduce the energy/power consumption or avoid temperature hot spots. The OS scheduler, on the other hand, may be designed to maximize the average CPU utilization for average performance. Providing tight bounds on execution qualities of individual applications sharing an execution platform is therefore not possible on many MPSoC platforms available today.

One remedy out of this dilemma that has been proposed a long time before the introduction of any MPSoC technology is isolation. With isolation, a set of techniques is subsumed to separate the execution of multiple programs either spatially (by allocating disjoint resources) or temporally (by separating the time intervals shared resources are used). Additionally, in order to provide isolation on demand, there is the need for adaptivity in all hardware as well as software layers from application program to executing hardware platform. Indeed, adaptivity is considered a key topic in order to reduce or bound execution quality variations actively on a system and in an on-demand manner for the reason to neither overly restrict nor to underutilize available resources.

Adaptive Isolation, the topic of the proposed Dagstuhl seminar, may be seen as a novel and important research topic for providing predictability of not only timing but also security and may be even other properties of execution on a multi-core

platform on a per application/job basis while easing and trading off compile-time and run-time complexity.

First, a common understanding of which techniques may be used for isolation including hardware units design, resource reservation protocols, virtualization techniques, and including novel hybrid and dynamic resource assignment techniques were discussed. Second, a very interdisciplinary team of experts including processor designers, OS and compiler specialists, as well as experts for predictability and security analysis were brought together for evaluating these opportunities and presenting novel solutions. The competencies, experiences, and existing solutions of the multiple communities stimulated discussions and co-operations that hopefully will manifest in innovative research directions for enabling predictability on demand on standard embedded MPSoCs.

6.63 Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR)

Organizers: Roger K. Moore, Serge Thill, and Ricard Marxer

Seminar No. 16442

Date: October 30 to November 4, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.10.154

Participants: Andrey Anikin, Timo Baumann, Tony Belpaeme, Elodie Briefer, Nick Campbell, Fred Cummins, Angela Dassow, Robert Eklund, Julie E. Elie, Sabrina Engesser, Sarah Hawkins, Ricard Marxer, Roger K. Moore, Julie Oswald, Bhiksha Raj, Rita Singh, Dan Stowell, Zheng-Hua Tan, Serge Thill, Petra Wagner, Benjamin Weiss

Almost all animals exploit vocal signals for a range of ecologically-motivated purposes. For example, predators may use vocal cues to detect their prey (and vice versa), and a variety of animals (such as birds, frogs, dogs, wolves, foxes, jackals, coyotes, etc.) use vocalisation to mark or defend their territory. Social animals (including human beings) also use vocalisation to express emotions, to establish social relations and to share information, and humans beings have extended this behaviour to a very high level of sophistication through the evolution of speech and language – a phenomenon that appears to be unique in the animal kingdom, but which shares many characteristics with the communication systems of other animals.

Also, recent years have seen important developments in a range of technologies relating to vocalisation. For example, systems have been created to analyse and playback animals calls, to investigate how vocal signalling might evolve in communicative agents, and to interact with users of spoken language technology (voice-based human-computer interaction using speech technologies such as automatic speech recognition and text-to-speech synthesis). Indeed, the latter has witnessed huge commercial success in the past 10-20 years, particularly since the release of *Naturally Speaking* (Dragon's continuous speech dictation software for a PC) in 1997 and Siri (Apple's voice-operated personal assistant and knowledge navigator for the iPhone) in 2011. Research interest in this area is now beginning to focus on voice-enabling autonomous social agents (such as robots).

Therefore, whether it is a bird raising an alarm, a whale calling to potential partners, a dog responding to human commands, a parent reading a story with a child, or a businessperson accessing stock prices using an automated voice service on their mobile phone, vocalisation provides a valuable communications channel through which behaviour may be coordinated and controlled, and information may be distributed and acquired.

Indeed, the ubiquity of vocal interaction has given rise to a wealth of research across an extremely diverse array of fields from the behavioural and language sciences to engineering, technology and robotics. This means that there is huge potential for crossfertilisation between the different disciplines involved in the study and exploitation of vocal interactivity. For example, it might be possible to use contemporary advances in machine learning to analyse animal activity in different habitats, or to use robots to investigate contemporary theories of language grounding. Likewise, an understanding of animal vocal behaviour might inform how vocal expressivity might be integrated into the next generation of autonomous social agents. Some of these issues have already been addressed by relevant sub-sections of the research community. However, many opportunities remain unexplored, not least due to the lack of a suitable forum to bring the relevant people together.

Our Dagstuhl seminar on the topic of "Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR)" provided the unique and timely opportunity to bring together scientists and engineers from a number of different fields to appraise our current level of knowledge. Our broad aim was to focus discussion on the general principles of vocal interactivity as well as evaluating the state-of-the-art in our understanding of vocal interaction within-and-between humans, animals and robots. Some of these sub-topics, such as human spoken language or vocal interactivity between animals, have a long history of scientific research. Others, such as vocal interaction between robots or between robots and animals, are less well studied – mainly due to the relatively recent appearance of the relevant technology. What is interesting is that, independent of whether the sub-topics are well established fields or relatively new research domains, there is an abundance of open research questions which may benefit from a comparative interdisciplinary analysis of the type addressed in this seminar.

6.64 Structure and Hardness in P

Organizers: Moshe Lewenstein, Seth Pettie, and Virginia Vassilevska Williams

Seminar No. 16451

Date: November 6–11, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.1

Participants: Amir Abboud, Alexandr Andoni, Arturs Backurs, Andreas Björklund, Karl Bringmann, Maike Buchin, Yi-Jun Chang, Raphaël Clifford, Radu Curticapean, Marek Cygan, Søren Dahlgaard, Holger Dell, Fedor V. Fomin, Omer Gold, Isaac Goldstein, Thore Husfeldt, Mathias Bæk Tejs Knudsen, Young Kun Ko, Tsvi Kopelowitz, Robert Krauthgamer, Sebastian Krinninger, Marvin Künnemann, Moshe Lewenstein, Andrea Lincoln, Veronika Loitzenbauer, Aleksander Madry, Kurt Mehlhorn, Ivan Mikhailin, Danupon Nanongkai, Ramamohan Paturi, Seth Pettie, Ely Porat, Vijaya Ramachandran, Liam Roditty, Aviad Rubinstein, Thatchaphol Saranurak, Stefan Schneider, Arkadiusz Socala, Virginia Vassilevska Williams, Joshua R. Wang, Oren Weimann, Philip Wellnitz, Huacheng Yu

The complexity class **P** (polynomial time) contains a vast variety of problems of practical interest and yet relatively little is known about the structure of P, or of the complexity of many individual problems in P. It is known that there exist contrived problems requiring $\Omega(n1.5)$ time or $\Omega(n2)$ time, and yet to date no unconditional nonlinear lower bounds have been proved for any problem of practical interest. However, the last few years have seen a new resurgence in conditional lower bounds, whose validity rests on the conjectured hardness of some archetypal computational problem. This work has imbued the class P with new structure and has valuable explanatory power.

To cite a small fraction of recent discoveries, it is now known that classic dynamic programming problems such as Edit Distance, LCS, and Fréchet distance require quadratic time (based on the conjectured hardness of k-CNF-SAT), that the best known triangle enumeration algorithms are optimal (based on the hardness of 3-SUM), that Valiant's context-free grammar parser is optimal (based on the hardness of k-CLIQUE), and that the best known approximate Nash equilibrium algorithm is optimal (based on the hardness of 3-SAT).

This Dagstuhl Seminar will bring together top researchers in diverse areas of theoretical computer science and include a mixture of both experts and non-experts in conditional lower bounds. Some specific goals of this seminar are listed below.

- Numerous important problems (such as Linear Programming) seem insoluble in linear time, and yet no conditional lower bounds are known to explain this fact. A goal is to discover conditional lower bounds for key problems for which little is currently known.
- Recent work has been based on both traditional hardness assumptions (such as the ETH, SETH, 3SUM, and APSP conjectures) and a variety of newly considered hardness assumptions (such as the OMv conjecture, the k-CLIQUE conjecture, and the Hitting Set conjecture). Almost nothing is known about the relative plausibility of these conjectures, or if multiple conjectures are, in fact, equivalent. A goal is to discover formal relationships between the traditional and newer hardness assumptions.
- A key goal of the seminar is to disseminate the techniques used to prove conditional lower bounds, particularly to researchers from areas of theoretical computer science that have yet to benefit from this theory. To this end the seminar will include a number of tutorials from top experts in the field.

6.65 Beyond-Planar Graphs: Algorithmics and Combinatorics

Organizers: Seok-Hee Hong, Michael Kaufmann, Stephen G. Kobourov, and János Pach

Seminar No. 16452

Date: November 6–11, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.35

Participants: Muhammad Jawaherul Alam, Patrizio Angelini, Evmorfia Argyriou, Michael Bekos, Franz J. Brandenburg, Sabine Cornelsen, Giordano Da Lozzo, Giuseppe Di Battista, Walter Didimo, Stefan Felsner, Kathrin Hanauer, Seok-Hee Hong, Michael Kaufmann, Stephen G. Kobourov, Giuseppe Liotta, Fabrizio Montecchiani, Quan Nguyen, Martin Nöllenburg, Yoshio Okamoto, Janos Pach, Sergey Pupyrev, Chrysanthi Raftopoulou, Md. Saidur Rahman, Vincenzo Roselli, Ignaz Rutter, Géza Tóth, Torsten Ueckerdt, Pavel Valtr, Alexander Wolff

Relational data sets, containing a set of objects and relations between them, are commonly modeled by graphs/networks, with the objects as the vertices and the relations as the edges. A great deal is known about the structure and properties of special types of graphs, in particular *planar graphs*. The class of planar graphs is fundamental for both Graph Theory and Graph Algorithms, and extensively studied. Many structural properties of planar graphs are known and these properties can be used in the development of efficient algorithms for planar graphs, even where the more general problem is NP-hard.

Most real world graphs, however, are *non-planar*. In particular, many scale-free networks, which can be used to model web-graphs, social networks and biological networks, consists of sparse non-planar graphs. To analyze and visualize such real-world networks, we need to solve fundamental mathematical and algorithmic research questions on *sparse non-planar* graphs, which we call *beyond-planar graphs*. The notion of beyond-planar graphs has been established as non-planar graphs with topological constraints such as specific types of crossings or with some forbidden crossing patterns, although it has not been formally defined. Examples of beyond-planar graphs include:

- *k-planar* graphs: graphs which can be embedded with at most k crossings per edge.
- *k-quasi-planar* graphs: graphs which can be embedded without k mutually crossing edges.
- *bar k-visibility* graphs: graphs whose vertices are represented as horizontal segments (bars) and edges as vertical lines connecting bars, intersecting at most k other bars.
- *fan-crossing-free* graphs: graphs which can be embedded without fan-crossings.
- *fan-planar* graphs: graphs which can be embedded with crossings sharing the common vertices.
- *RAC (Right Angle Crossing)* graphs: a graph which has a straight-line drawing with right angle crossings.

The aim of the seminar was to bring together world-renowned researchers in graph algorithms, computational geometry and graph theory, and collaboratively develop a research agenda for the study of beyond-planar graphs. The plan was to work on specific open problems about the structure, topology, and geometry of beyond-planar graphs. One of the outcomes of the workshop might be an annotated bibliography of this new field of study.

On Sunday afternoon, 29 participants met at Dagstuhl for an informal get-together. Fortunately, there were no cancelations and everybody who registered was able to attend. On Monday morning, the workshop officially kicked off. After a round of introductions, where we discovered that eight participants were first-time Dagstuhl attendees, we enjoyed three overview talks about beyond-planar graphs from three different points of view. First, Géza Tóth from the Rényi Institute in Budapest talked about the combinatorics of beyond-planar graphs in connection to graph theory. Next, Giuseppe Liotta from the University of Perugia gave an overview about the connections between graph drawing and beyond-planar graphs and presented a taxonomy of related topics and questions. Finally, Alexander Wolff from the University of Würzburg discussed beyond-planar graphs in the context of geometry and geometric graph representations.

On Monday afternoon, we had lively open problem sessions, where we collected 20 problems covering the most relevant topics. The participants split into four groups based on common interest in subsets of the open problems. The last three days of the seminar were dedicated to working group efforts. Most of the groups kept their focus on the original problems as stated in the open problem session, while one group modified and expanded the problems. We had two progress reports sessions, including one on Friday morning, where group leaders were officially designated and plans for follow-up work were made. Work from one of the groups has been submitted to an international conference, and we expect further research publications to result directly from the seminar.

Arguably the best, and most-appreciated, feature of the seminar was the opportunity to engage in discussion and interactions

with experts in various fields with shared passion about graphs, geometry and combinatorics. We received very positive feedback from the participants (e.g., scientific quality: 10.5/11, inspired new ideas: 23/25, inspired joint projects: 21/25) and it is our impression that the participants enjoyed the unique scientific atmosphere at the seminar and benefited from the scientific program. In summary, we regard the seminar as a success, and we are grateful for having had the opportunity to organize it and take this opportunity to thank the scientific, administrative, and technical staff at Schloss Dagstuhl.

6.66 Assessing ICT Security Risks in Socio-Technical Systems

Organizers: Tyler W. Moore, Christian W. Probst, Kai Rannenberg, and Michel van Eeten

Seminar No. 16461

Date: November 13–18, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.63

Participants: Ross Anderson, Johannes M. Bauer, Zinaida Benenson, Rainer Böhme, L. Jean Camp, Tristan Caulfield, Nicolas Christin, Richard Clayton, Serge Egelman, Carlos H. Ganan, Dieter Gollmann, Hannes Hartenstein, Florian Kammüller, Vincent Koenig, Barbara Kordy, Stewart Kowalski, Kwok-Yan Lam, Stefan Laube, Gabriele Lenzini, Thomas Maillart, Fabio Massacci, Kanta Matsuura, Tyler W. Moore, Sebastian Pape, Simon Parkin, Wolter Pieters, Christian W. Probst, Kai Rannenberg, Martina Angela Sasse, Christian Sillaber, Sven Übelacker, Michel van Eeten, Maarten van Wieren, Melanie Volkamer, Edgar A. Whitley, Jeff Yan

The Dagstuhl Seminar 16461 "Assessing ICT Security Risks in Socio-Technical Systems" is part of a series of seminars that explore aspects of risk and security in socio-technical systems. After initial work on insider threats, the focus has turned towards understanding of relevant metrics and their application in novel security risk assessment methods.

Classical Risk Assessment

Assessing risk in classic mechanical or electrical production systems is difficult but possible, as experience shows. Historical knowledge provides us with approximation of likelihood of failures of machines and components. We can combine these likelihoods in modular ways, mapping the impact of the loss of a system component to the processes that component contributes to.

This approach works for traditional engineering systems, since the dependencies between components are expected to be known, and their behaviour is assumed to be deterministic. To reach a comparable level of predictability for risk assessment in areas that are less governed by machines, for example economics, standard procedure is to fix those factors that are not deterministic. The behaviour of buyers and sellers on a market, for example, is assumed to be gain-oriented and rational; dealing with irrational actors is less explored and hard to model.

Risk in ICT Systems and the Immaterial Nature of Information. The techniques that work well for assessing risk in classic mechanical or electrical production systems do not scale to ICT systems. A primary reason is that there is no clear connection between the usage of a system and the risk of failure. Many risks that we need to consider in ICT are hard to trace due to the immaterial nature of information [2]. Examples of such risks are unauthorized or illegitimate information flows. Information flows are more difficult to trace than material flows, as the flow is usually caused by a copying operation; the information is not missing at the source, which in the material world is an indicator of an illegitimate flow, *e.g.*, when goods are stolen. Moreover, the damage is barely related to the measurable amount of information [5, 9]. Several megabytes of, *e.g.*, white noise can be relatively harmless, while a single health data record or financial record of limited size can be a major problem, *e.g.*, in terms of loss of trust, reputation, or damage compensation payments.

Another reason why classic approaches have struggled to assess risk is that the threat from strategic adversaries is harder to model than random failures. Whereas events triggered by nature occur randomly, attackers can readily identify and target the weakest links present in systems, and adapt to evade defenses.

A final reason why assessing security risks is hard is that there is often an incentive to hide failures from public view, due to fears of reputational damage. This makes collecting data to empirically estimate loss probabilities very difficult.

Risk in Socio-technical Systems. Beyond ICT systems, socio-technical systems also contain human actors as integral parts of the system. In such socio-technical systems there may occur unforeseen interactions between the system, the environment, and the human actors, especially insiders [8].

Assessing the risk of the ICT system for human actors is difficult [4]; the assessment must take into account the effect of the ICT system on the environment, and it must quantify the likelihood for this risk to materialize. Assessing the risk of the human actor for the ICT system is difficult, too. As mentioned above, economics models human actors by assuming them to be gain-oriented and rational; dealing with irrational actors is less explored and hard to model. However, one of the biggest risks from human actors for an ICT system is irrational behaviour, or an unknown gain function.

Economics of Risk Assessment. The economic aspect both of the risk identified and the process of assessing risk often prohibits either risk mitigation or the assessment itself. Protection against irrational threats requires appropriate

preventive measure, be it too restrictive policies or too intense surveillance. Neither the cost nor the effects of these measures are easily predictable [1].

Even worse, the cost for risk assessment itself can also be prohibitive. For example, trying to identify the actual risk for irrational behaviour or its impact on the system can be impossible or at least imply a too high price [4].

Security Metrics. As we concluded after the previous Dagstuhl Seminar 14491, well-defined data sources and clear instructions for use of the metrics are key assets "to understand security in today's complex socio-technical systems, and to provide decision support to those who can influence security".

Security metrics obviously cannot be applied on their own, but must be embedded in a well-defined framework of sources for metrics and computations on them [7]. Important topics include understanding the aspects surrounding metrics, such as sources, computations on metrics, relations to economics, and the analyses based on metrics.

Assessing ICT Security Risks in Socio-Technical Systems. Making risk in socio-technical systems assessable requires an understanding of how to address issues in these systems in a systematic way [3, 6]. In this seminar, we built upon the work in the predecessor seminars on insider threats and security metrics, and explores the embedding of human behaviour and security metrics into methods to support risk assessment.

Main findings

We established five working groups in the seminar, that discussed several times during the week and reported back in plenum. The results are presented in the full report, and briefly summarized here.

Which data do we need to collect? In a working group on "Collecting Data for the Security Investment Model", we considered the relationship between efforts to secure an organisation, the actual security level achieved through these efforts, and their effect on moderating attacks and the induced losses. We identified relevant, measurable indicators for the components in the model that relate metrics about components to the expected risk. Model outcomes could be used to guide security investments.

Which security risks should we consider? To identify relevant risk assessment methods, we discussed the kind of risk relevant to measure in two working groups. On the one hand we explored "New Frontiers of Socio-Technical Security", where we considered disrupting new technologies and how they influence and change our perception of risk, or its limitations. The main example were orphaned devices in IoT systems, which often cannot be switched off, but pose a threat to the overall system if they remain unmaintained. A similar problem space was explored in the working group on "Software Liability Regimes", which considered liability or lack thereof of producers to identify and fix problems.

Which attacker and user traits do we need to consider? To understand relevant aspects of human actors involved in socio-technical systems, we established two more working groups. The group on "Unpacking the motivations of attackers" discussed how to understand attacker motivations in highly integrated socio-technical systems, where purpose and means play a fundamental role in the way and at which level(s) the cyberspace can be disrupted.

Conclusions

Assessing risk in socio-technical systems is and remains difficult, but can be supported by techniques and understanding of limitations and properties inherent to the system and the risk assessment methods applied. This seminar has explored how to identify these limitations and properties by exploring the different layers of socio-technical systems, their interactions, and their defining attributes.

A total of 36 researchers participated in the seminar across from different communities, which together span the range relevant to developing novel security risk assessment methods and to ensure the continuation from the previous seminars' results: cyber security, information security, data-driven security, security architecture, security economics, human factors, (security) risk management, crime science, formal methods, and social science.

References

1 Anderson, R., Barton, C., Böhme, R., Clayton, R., van Eeten, M. J. G., Levi, M., Moore, T., and Savage, S. (2012). Measuring the Cost of Cybercrime. Paper presented at the Workshop of Economics and Information Security (WEIS 2012), Berlin.

2 Blakley, B., McDermott, E., Morgan, J. P., and Geer, D. (2001). Information security is information risk management. Proc. 2001 Workshop on New Security Paradigms (NSPW'01), pp. 97–104, ACM.

3 Gollmann, D. (2012). From Insider Threats to Business Processes that are Secure-by-Design. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, 3(1/2), 4–12.

4 Hunker, J. and Probst, C. W. (2009). The Risk of Risk Analysis, And its relation to the Economics of Insider Threats. Proc. 2009 Workshop on Economics of Information Security (WEIS'09).

5 Johnson, M. E, Goetz, E., and Pfleeger, S. L. (2009). Security through Information Risk Management, IEEE Security and Privacy, 7(3):45–52.

6 Littlewood, B., Brocklehurst, S., Fenton, N., Mellor, P., Page, S., Wright, D., Dobson, J., McDermid, J., and Gollmann, D. (1993). Towards Operational Measures of Computer Security. J. Computer Security, 2(2–3):211–229.

7 Pieters, W., Van der Ven, S. H. G., and Probst, C. W. (2012). A move in the security measurement stalemate: Elo-style ratings to quantify vulnerability. Proc. 2012 New Security Paradigms Workshop (NSPW'12), pp. 1–14, ACM. doi: 10.1145/2413296.2413298.

8 Probst, C. W., Hunker, J., Bishop, M. and Gollmann, D. (2009). Countering Insider Threats. ENISA Quarterly Review, 5(2).

9 Tschersich, M., Kahl, C., Heim, S., Crane, S., Böttcher, K., Krontiris, I., Rannenberg, K. (2011). Towards Privacy-enhanced Mobile Communities – Architecture, Concepts and User Trials. J. Systems and Software, Volume 84, Issue 11, pp. 1947–1960, 2011.

6.67 Inpainting-Based Image Compression

Organizers: Christine Guillemot, Gerlind Plonka-Hoch, Thomas Pock, and Joachim Weickert

Seminar No. 16462

Date: November 13–18, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.90

Participants: Sarah Andris, Johannes Ballé, Zakaria Belhachmi, Aurélien Bourquard, Eva-Maria Brinkmann, Dorin Bucur, Martin Buhmann, Gene Cheung, Nira Dyn, Gabriele Facciolo, Jalal Fadili, Irena Galic, Yann Gousseau, Christine Guillemot, Laurent Hoeltgen, Armin Iske, Claire Mantel, Simon Masnou, Peter Ochs, Pascal Peter, Gerlind Plonka-Hoch, Thomas Pock, Daniela Rosca, Naoki Saito, Tomas Sauer, Carola-Bibiane Schönlieb, Joan Serra Sagristà, Joachim Weickert, Hao-Min Zhou

Since the amount of visual data is rapidly increasing, there is a high demand for powerful methods for compressing digital images. A well-known example is the lossy JPEG standard that is based on the discrete cosine transform. Unfortunately its quality deteriorates substantially for high compression rates, such that better alternatives are needed.

The goal of this seminar was to pursue a completely different strategy than traditional, transform-based codecs (coders and decoders): We studied approaches that rely on so-called inpainting methods. They store only a small, carefully selected subset of the image data. In the decoding phase, the missing data is reconstructed by interpolation with partial differential equations (PDEs) or by copying information from patches in other image regions. Such codecs allow a very intuitive interpretation, and first experiments show their advantages for high compression rates where they can beat even advanced transform-based methods.

However, inpainting-based codecs are still in an early stage and require to solve a number of challenging fundamental problems, in particular:

1. Which data gives the best reconstruction?
2. What are the optimal inpainting operators?
3. How should the selected data be encoded and decoded?
4. What are the most efficient algorithms for real-time applications?

These problems are highly interrelated. Moreover, they require interdisciplinary expertise from various fields such as image inpainting, data compression and coding, approximation theory, and optimisation. To design these codecs in an optimal way, one must also understand their connections to related areas such as sparsity and compressed sensing, harmonic analysis, scattered data approximation with radial basis functions, and subdivision strategies.

Our seminar constituted the first symposium on this topic. It brought together 29 researchers from 11 countries, covering a broad range of expertise in the different fields mentioned above. Many of them have met for the first time, which resulted in a very fruitful interaction.

In order to have a good basis for joint discussions, first all participants introduced themselves and briefly described their background and interests. Then the seminar proceeded with six tutorial talks (45 minutes plus 15 minutes discussion), given by the four organisers as well as by Simon Masnou and Nira Dyn. In this way all participants could acquire a general overview on the achievements and challenges of inpainting-based image compression and its various aspects such as coding, inpainting, convex optimisation, subdivision, and computational harmonic analysis.

Afterwards we decided to cluster the talks thematically into six sessions, each consisting of 3–4 talks (ca. 30 minutes plus 15 minutes discussion) and lasting half a day:

1. Harmonic Analysis
   (talks by Gerlind Plonka-Hoch, Naoki Saito, and Hao-Min Zhou)
2. Approximation Theory
   (talks by Martin Buhmann, Armin Iske, Nira Dyn, and Tomas Sauer)
3. Inpainting
   (talks by Aurelien Bourquard, Carola-Bibiane Schönlieb, and Yann Gousseau)
4. Compression
   (talks by Gene Cheung, Joan Serra Sagrista, and Claire Mantel)
5. Optimisation of Data and Operators
   (talks by Zakaria Belhachmi, Laurent Hoeltgen, Peter Ochs, and Pascal Peter)
6. Algorithms, Biological Vision, and Benchmarking
   (talks by Jalal Fadili, Johannes Ballé, and Sarah Andris)

These sessions triggered interesting discussions during the talks, in the breaks, and in the evening, and they allowed the different communities to learn many new things from each other.

Our program featured also an evening panel discussion on open research questions on the interface between image inpainting and image compression. It was a lively interaction between the five panel members and the audience, involving also controversial statements and views about the future of inpainting-based codecs.

The participants had a very positive impression of this seminar as an inspiring forum to bring together different fields. As a consequence, this symposium also created several new collaborations, e.g. regarding interpolation with radial basis functions, subdivision-based coding, and diffusion-based coding. There was a general consensus that it would be desirable to have another seminar on this topic in 2–3 years. Moreover, it is planned to compile a related monograph which will be the first in its field.

6.68 Concurrency with Weak Memory Models: Semantics, Languages, Compilation, Verification, Static Analysis, and Synthesis

Organizers: Jade Alglave, Patrick Cousot, and Caterina Urban

Seminar No. 16471

Date: November 20–25, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.108

Participants: Jade Alglave, Giovanni Tito Bernardi, Annette Bieniusa-Middelkoop, Richard Bornat, Stephen Brookes, Simon Castellan, Andrea Cerone, Pierre Clairambault, Patrick Cousot, Andrei Marian Dan, Will Deacon, David Delmas, Delphine Demange, Stephan Diestelhorst, Charles Anthony Richard Hoare, Vincent Jacques, Bernhard Kragl, Ori Lahav, Daniel Lustig, Yatin Manerkar, Luc Maranget, Paul McKenney, Paul-Andre Mellies, Roland Meyer, Maged M. Michael, Antoine Miné, Vincent Nimal, Andrea Parri, Gustavo Petri, Susmit Sarkar, Helmut Seidl, Suzanne Shoaraee, Daryl Stewart, Caroline J. Trippel, Caterina Urban, Viktor Vafeiadis, Derek Williams, Glynn Winskel, Sizhuo Zhang

In the last decade, research on weak memory has focussed on modeling accurately and precisely existing systems such as hardware chips. These laudable efforts have led to definitions of models such as IBM Power, Intel x86, Nvidia GPUs and others.

Now that we have faithful models, and know how to write others if need be, we can focus on how to use these models for verification, for example to assess the correctness of concurrent programs.

The goal of our seminar was to discuss how to get there. To do so, we gathered people from various horizons: hardware vendors, theoreticians, verification practitioners and hackers. We asked them what issues they are facing, and what tools they would need to help them tackle said issues.

The first day was dedicated to theory. We had overviews of classic semanticists tools such as event structures, message sequence charts, and pomsets. The remaining days were mostly dedicated to models and verification practices, whether from a user point of view, or a designer point of view. We chose to close the days early, so that our guests would have ample time to come back to an interesting point they had heard during one of the talks, or engage in deep discussions. The feedback we got was quite positive, in that the seminar help spark discussions with, for example, a PhD student in concurrency theory, and a verification practitioner from ARM.

6

6.69 QoE Vadis?

Organizers: Markus Fiedler, Sebastian Möller, Peter Reichl, and Min Xie
Seminar No. 16472

Date: November 20–25, 2016 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.6.11.129

Participants: Jan-Niklas Antons, Luigi Atzori, Katrien De Moor, Touradj Ebrahimi, Sebastian Egger-Lampl, Markus Fiedler, Jörgen Gustafsson, Tobias Hoßfeld, Lucjan Janowski, Kalevi Kilkki, Udo Krieger, Effie Lai-Chong Law, Sebastian Möller, Marianna Obrist, Peter Reichl, Virpi Hannele Roto, Henning Schulzrinne, Lea Skorin-Kapov, Jan Van Looy, Martín Varela, Katarzyna Wac, Felix Wu, Min Xie, Hans-Jürgen Zepernick

During the recent decade, the transition from the technology-oriented notion of QoS (Quality of Service) to the user-centric concept of QoE (Quality of Experience) has become an important paradigm change in communication networking research. Simultaneously, the field of QoE as such has significantly developed and matured. This is amongst others reflected in the series of three Dagstuhl Seminars 09192 "From Quality of Service to Quality of Experience" (2009), 12181 "Quality of Experience: From User Perception to Instrumental Metrics" (2012), and 15022 "Quality of Experience: From Assessment to Application" (2015).

The QoE-related Dagstuhl Seminars had a significant impact on the understanding, definition and application of the QoE notion and concepts in the QoE community, for instance with respect to redefining fundamental concepts of quality. That work was performed in close collaboration with the COST Action IC1003 Qualinet [1] that has been concentrating on QoE in multimedia systems and services, and is still actively convening experts from all over the world to regular meetings and exchanges. In particular, this collaboration has led to the widely regarded Qualinet White Paper on "Definitions of QoE and related concepts" [1] and to the launch of a new journal entitled "Quality and User Experience" [2], fostering the scientific exchange within and between QoE and User Experience (UX) communities.

Realising the urgent need of jointly and critically reflecting the future perspectives and directions of QoE research, the QoE-related Dagstuhl Seminars were complemented by the present Dagstuhl Perspectives Workshop 16472 "QoE Vadis?", whose output is compiled in a Dagstuhl Manifesto. Besides of having brought together the two communities much closer, and besides triggering new events such as special sessions at conferences, the main outcome of the workshop has been concretized in terms of 11 recommendations to be communicated to stakeholders in the QoE and UX domains.

The workshop was organised around the writing process of the Manifesto draft: Starting from personal statements instead of talks, two sets of group works were arranged, whose output was criticially reviewed by "Advocatii Diaboli" and then refined and extended. A final review round by one representative of each the QoE and the UX group was performed before the Manifesto draft was completed by the end of the week.

References

1 European Network on Quality of Experience in Multimedia Systems and Services (COST IC 1003 Qualinet), http://www.qualinet.eu (last seen 2017-02-24).

2 Quality and Experience (QUEX), a journal published by Springer, http://link.springer.com/journal/41233 (last seen 2017-02-24).

6.70 New Directions for Learning with Kernels and Gaussian Processes

Organizers: Arthur Gretton, Philipp Hennig, Carl Edward Rasmussen, and Bernhard Schölkopf

Seminar No. 16481

Date: November 27 to December 2, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.142

Participants: Jan-Niklas Antons, Luigi Atzori, Katrien De Moor, Touradj Ebrahimi, Sebastian Egger-Lampl, Markus Fiedler, Jörgen Gustafsson, Tobias Hoßfeld, Lucjan Janowski, Kalevi Kilkki, Udo Krieger, Effie Lai-Chong Law, Sebastian Möller, Marianna Obrist, Peter Reichl, Virpi Hannele Roto, Henning Schulzrinne, Lea Skorin-Kapov, Jan Van Looy, Martín Varela, Katarzyna Wac, Felix Wu, Min Xie, Hans-Jürgen Zepernick

Machine learning is a young field that currently enjoys rapid, almost dizzying advancement both on the theoretical and the practical side. On account of either, the until quite recently obscure discipline is increasingly turning into a central area of computer science. Dagstuhl seminar 16481 on *"New Directions for Learning with Kernels and Gaussian Processes"* attempted to allow a key community within machine learning to gather its bearings at this crucial moment in time.

Positive definite kernels are a concept that dominated machine learning research in the first decade of the millennium. They provide infinite-dimensional hypothesis classes that deliver expressive power in an elegant analytical framework. In their probabilistic interpretation as Gaussian process models, they are also a fundamental concept of Bayesian inference:

A *positive definite kernel* $k : \mathbb{X} \times \mathbb{X} \to \Re$ on some input domain $\mathbb{X}$ is a function with the property that, for all finite sets $\{x_1, \dots, x_N\} \subset \mathbb{X}$, the matrix $K \in \Re^{N \times N}$, with elements $k_{ij} = k(x_i, x_j)$, is positive semidefinite. According to a theorem by Mercer, given certain regularity assumptions, such kernels can be expressed as a potentially *infinite* expansion

$$k(x, x') = \sum_{i=1}^{\infty} \lambda_i \phi_i(x) \phi_i^*(x'), \qquad \text{with} \qquad \sum_{i=1}^{\infty} \lambda_i < \infty,$$

where $*$ is the conjugate transpose, $\lambda_i \in \Re_+$ is a non-negative *eigenvalue* and ϕ_i is an *eigenfunction* with respect to some measure $\nu(x)$: a function satisfying

$$\int k(x, x') \phi_i(x) d\nu(x) = \lambda_i \phi_i(x').$$

Random functions $f(x)$ drawn by independently sampling Gaussian weights for each eigenfunction,

$$f(x) = \sum_{j=1}^{\infty} f_j \phi_j(x) \qquad \text{where} \qquad f_j \sim \mathcal{N}(0, \lambda_i),$$

are draws from the centered *Gaussian process* (GP) $p(f) = \mathcal{GP}(f; 0, k)$ with *covariance function* k. The logarithm of this Gaussian process measure is, up to constants and some technicalities, the square of the norm $\|f\|_k^2$ associated with the *reproducing kernel Hilbert space* (RKHS) of functions reproduced by k.

Supervised machine learning methods that *infer* an unknown function f from a data set of input-output pairs $(X, Y) := \{(x_i, y_i)\}_{i=1,\dots,N}$ can be constructed by minimizing an empirical risk $\ell(f(X); Y)$ regularized by $\| \cdot \|_k^2$. Or, algorithmically equivalent but with different philosophical interpretation, by computing the *posterior* Gaussian process measure arising from conditioning $\mathcal{GP}(f; 0, k)$ on the observed data points under a likelihood proportional to the exponential of the empirical risk.

The prominence of kernel/GP models was founded on this conceptually and algorithmically compact yet statistically powerful description of inference and learning of nonlinear functions. In the past years, however, hierarchical ('deep') parametric models have bounced back and delivered a series of impressive empirical successes. In areas like speech recognition and image classification, deep networks now far surpass the predictive performance previously achieved with nonparametric models. One central goal of the seminar was to discuss how the superior adaptability of deep models can be transferred to the kernel framework while retaining at least some analytical clarity. Among the central lessons from the 'deep resurgence' identified by the seminar participants is that the kernel community has been too reliant on theoretical notions of universality. Instead, representations must be learned on a more general level than previously accepted. This process is often associated with an 'engineering' approach to machine learning, in contrast to the supposedly more 'scientific' air surrounding kernel methods. But its importance must not be dismissed. At the same time, participants also pointed out that deep learning is often misrepresented, in particular in popular expositions, as an almost magic kind of process; when in reality the concept

is closely related to kernel methods, and can be understood to some degree through this connection: Deep models provide a hierarchical parametrization of the feature functions $\phi_i(x)$ in terms of a finite-dimensional family. The continued relevance of the established theory for kernel/GP models hinges on how much of the power of deep models can be understood from within the RKHS view, and how much new concepts are required to understand the expressivity of a deep learning machine.

There is also unconditionally good news: In a separate but related development, kernels have had their own renaissance lately, in the young areas of probabilistic programming ('computing of probability measures') and probabilistic numerics ('probabilistic descriptions of computing'). In both areas, kernels and Gaussian processes have been used as a descriptive language. And, similar to the situation in general machine learning, only a handful of comparably simple kernels have so far been used. The central question here, too, is thus how kernels can be designed for challenging, in particular high-dimensional regression problems. In contrast to the wider situation in ML, though, kernel design here should take place at compile-time, and be a structured algebraic process mapping source code describing a graphical model into a kernel. This gives rise to new fundamental questions for the theoretical computer science of machine learning.

A third thread running through the seminar concerned the internal conceptual schism between the probabilistic (Gaussian process) view and the statistical learning theoretical (RKHS) view on the model class. Although the algorithms and algebraic ides used on both sides overlap *almost* to the point of equivalence, their philosophical interpretations, and thus also the required theoretical properties, differ strongly. Participants for the seminar were deliberately invited from both "denominations" in roughly equal number. Several informal discussions in the evenings, and in particular a lively break-out discussion on Thursday helped clear up the mathematical connections (while also airing key conceptual points of contention from either side). Thursday's group is planning to write a publication based on the results of the discussion; this would be a highly valuable concrete contribution arising from the seminar, that may help drawing this community closer together.

Despite the challenges to some of the long-standing paradigms of this community, the seminar was infused with an air of excitement. The participants seemed to share the sensation that machine learning is still only just beginning to show its full potential. The mathematical concepts and insights that have emerged from the study of kernel/GP models may have to evolve and be adapted to recent developments, but their fundamental nature means they are quite likely to stay relevant for the understanding of current and future model classes. Far from going out of fashion, mathematical analysis of the statistical and numerical properties of machine learning model classes seems slated for a revival in coming years. And much of it will be leveraging the notions discussed at the seminar.

6.71 Algorithms and Effectivity in Tropical Mathematics and Beyond

Organizers: Stéphane Gaubert, Dima Grigoriev, Michael Joswig, and Thorsten Theobald

Seminar No. 16482

Date: November 27 to December 2, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.168

Participants: Marianne Akian, Xavier Allamigeon, Fuensanta Aroca, Frédéric Bihan, Manuel Bodirsky, Timo de Wolff, Cristhian Garay, Stéphane Gaubert, Dima Grigoriev, Alexander Guterman, Christian Haase, Simon Hampe, Thomas Dueholm Hansen, Anders Jensen, Thorsten Jörgens, Michael Joswig, Bo Lin, Georg Loho, Marie MacCaig, Diane Maclagan, Hannah Markwig, Thomas Markwig, Vladimir Podolskii, Felipe Rincon, J. Maurice Rojas, Benjamin Schröter, Sergei Sergeev, Mateusz Skomra, Luis Tabera, Thorsten Theobald, Cynthia Vinzant

Tropical mathematics is a uniting name for different research directions which involve the semi-ring of real numbers endowed with the operations min, + (called the tropical semi-ring). It has emerged in several areas of computer science and of pure and applied mathematics. For the first time, this seminar brought together the computer science and the mathematics viewpoints. A focus was on effective methods, algorithms and complexity bounds in tropical mathematics, and on their relations with open questions in various areas of computer science, including optimization, game theory and circuit complexity.

One of the oldest open algorithmic challenges in tropical mathematics is the complexity of solving systems of tropical linear equalities and inequalities. It is known to be equivalent to solving mean payoff games. The solvability of these problems is among the few known problems which are contained in the intersection NP ∩ co-NP, but not currently known to be in P. This leads to new approaches in linear programming or convex semialgebraic programming over nonarchimedean fields.

According to the organizers' points of view the seminar was quite successful. In addition to 28 talks there were many informal discussions and exchange of ideas in small groups. We expect several new common papers of the participants conceived during the seminar. An important feature was to bring together experts with different backgrounds who often knew other participants just by their publications. According to the opinions expressed, the participants learned a lot of new things. The seminar was especially useful for the young people.

Every talk, in addition to new results, also contained open problems. This created a lot of interaction in subsequent discussions. The audience was very active, many questions were posed to the speakers during the talks and the breaks.

The talks can be conditionally partitioned into the following groups, although there were many interrelations between different groups:

- Algorithmical problems of foundations of tropical mathematics (H. Markwig, D. Maclagan, F. Rincon, V. Podolskii);
- Complexity of games and of tropical linear and convex algebra (M. Bodirsky, S. Gaubert, T. Hansen, M. Joswig, G. Loho, M. MacCaig, B. Schröter, S. Sergeev, M. Skomra);
- Algorithms and complexity bounds on tropical varieties (F. Bihan, D. Grigoriev, S. Hampe, A. Jensen, T. Jörgens, L. Tabera, T. Theobald, T. de Wolff). We mention that S. Hampe has made a demonstration of the software Polymake for computations in tropical algebra;
- Algorithms in tropical differential algebra (F. Aroca, C. Garay);
- Interactions of tropical mathematics with algorithmic issues in classical mathematics (M. Akian, X. Allamigeon, Bo Lin, M. Rojas, C. Vinzant).

During the seminar a manuscript appeared (on the Internet) with the very strong claim of a quasi-polynomial complexity algorithm for parity games. It was a lucky coincidence that so many experts with various backgrounds were present. So a special evening session on Thursday was created to analyze this result and its ramifications. This was one more highlight.

6.72 Symbolic-Numeric Methods for Reliable and Trustworthy Problem Solving in Cyber-Physical Domains

Organizers: Sergiy Bogomolov, Martin Fränzle, Kyoko Makino, and Nacim Ramdani

Seminar No. 16491

Date: December 4–9, 2016 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.11.168

Participants: Erika Ábrahám, Julien Alexandre dit Sandretto, Roman Barták, Sergiy Bogomolov, Martin Brain, Mauricio Castillo-Effen, Alexandre Chapoutot, Xin Chen, Chih-Hong Cheng, Eva Darulova, Parasara Sridhar Duggirala, Georgios Fainekos, Martin Fränzle, Mirco Giacobbe, Eric Goubault, Michael W. Hofbaur, Kyoko Makino, Marius Mikucionis, Ned Nedialkov, Markus Neher, Junkil Park, Pavithra Prabhakar, John D. Pryce, Sylvie Putot, Nacim Ramdani, Stefan Ratschan, Rajarshi Ray, Anne Remke, Karsten Scheibler, Zhikun She, Walid Taha, Yuichi Tazaki, Tino Teige, Louise Travé-Massuyès, Stavros Tripakis, Alexander Wittig

With the advent of cyber-physical systems increasingly penetrating our life, we are facing an ever-growing and permanent dependency on their reliable availability, continued function, and situationally adequate behavior even in highly sensitive application domains. As cyber-physical systems comprise complex, heteromorphic software systems, their reliability engineering calls for combinations of theories and methods traditionally considered separate. While we have recently seen some of the necessary combinations blossom, e.g. the theory of hybrid systems bridging continuous control with reactive systems, other areas remain less developed and explored. A prominent one is the role of numerics in cyber-physical systems: while it is obvious that cyber-physical systems increasingly rely on numerical software components, e.g., in signal processing or in state representation and extrapolation during situation assessment and planning, specific methods for addressing the issues associated, like consequences of numerical inaccuracy and methods for confining propagation of errors, are just in their infancy. This is in stark contrast to the use of numerics in more mature branches of computing, like signal processing or numerical analysis, where quantization effects as well as genesis and propagation of numerical error is well-understood and dedicated methods for controlling it in critical application, like various forms of interval-based numerical algorithms, are readily available. The aforementioned "traditional" methods are, however, not versatile enough to cope with the cyber-physical setting, where numerical results, like state extrapolations over significant temporal horizons, enter into complex and safety-critical decision making, rendering error propagation potentially highly discontinuous. It seems that future critical applications, like automated driving contributing to the EU's "Vision Zero" of eliminating fatalities in road-bound traffic, consequently call for novel means of analyzing and controlling the impact of numerics on system correctness, complemented by pertinent means of verification for establishing the safety case. The germs of such methods obviously have to be sought in the fields of design and verification of cyber-physical systems, i.e. in particular, (1) hybrid discrete-continuous systems, as well as (2) verified numerics, arithmetic constraint solving also involving symbolic reasoning, and (3) planning and rigorous optimization in arithmetic domains. The seminar gathered prominent researchers from all these fields in order to address the pressing problems induced by our societal dependence on cyber-physical systems.

As argued above, bringing together researchers dealing with hybrid discrete-continuous systems, with verified numerics in arithmetic constraint solving, and with planning and optimization in arithmetic domains can help improve the state of the art in rigorously interpreting and controlling cyber-physical phenomena. In the sequel, we review existing and potential contributions of the three fields to problem solving in cyber-physical domains and sketch potentials for cross-fertilization, which was the aim of the proposed seminar.

Fig. 6.19
Impressions of the creation of a second large lecture hall at Schloss Dagstuhl.

7 Öffentlichkeitsarbeit
Public Relations and Outreach

7.1 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 Dagstuhl[40] 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. 1 040 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 über 330 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ gepflegt (derzeit etwa 640 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 website[40].

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 1,040 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 more than 330 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with about 640 members), which supports the networking of participants in Dagstuhl Seminars. Additionally, interesting news about Schloss Dagstuhl are announced there.

7.2 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

[40] http://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 2016 gibt es auf der Webseite der Veranstaltung[41].

Wissenschaftsjournalismus

Unter dem Motto „Schreiben über Informatik“ fand vom 2. bis 4. Mai 2016 bereits zum achten Mal ein Trainings-Workshop für junge Wissenschafts-Journalisten und -Volontäre statt. Der Workshop richtete sich aber auch an etablierte Redakteure, die ihren Themenschwerpunkt erweitern möchten. Leitung des Workshops hatte auch in diesem Jahr der Wissenschaftsjournalist und Medientrainer Tim Schröder zusammen mit Gordon Bolduan, dem Experten für Wissenschaftskommunikation am *Kompetenzzentrum Informatik Saarland*.

Anhand aktueller Beispiele aus der Informatikforschung lernten die Workshop-Teilnehmer, wie abstrakte und technisch anspruchsvolle Informatik-Themen allgemein verständlich und spannend aufbereitet werden können. Alle Teilnehmer sowie die Dozenten waren höchst zufrieden mit den Inhalten und Ergebnissen des Workshops. Weitere Informationen sind auf der Webseite des Workshops[42] abrufbar.

participate. Details on the workshop in 2016 are available at the event webpage[41].

Computer science journalism

From May 2nd to May 4th, 2016, the 8th training workshop on computer science journalism was held at Schloss Dagstuhl. The workshop is designed as training opportunity for young journalists and trainees, as well as for established journalists who want to expand their focus. This year's workshop was again organized by science journalist and media trainer Tim Schröder and scientific communication expert Gordon Bolduan (Saarland Informatics Campus).

On the basis of contemporary topics from computer science research, the workshop's participants learned how to prepare and present abstract and technically sophisticated computer science topics in a comprehensible and exciting manner. Participants as well as trainers and referees were very satisfied with the workshop. See the event webpage[42] for further details.

[41] http://www.dagstuhl.de/16503
[42] http://www.dagstuhl.de/16184

Fig. 7.1
"Immer wieder schön hier in @dagstuhl . Diesmal #Informatik #Lehrerbildung mit #It2school" Twitter post by participant Ira Diethelm of the 2016 teacher training. https://twitter.com/elaine_miller/status/809734919740071937. Photo courtesy of Ira Diethelm.

„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.

Zum Auftakt am 28. Oktober 2016 präsentierte Prof. Holger Hermanns (Universität des Saarlandes) unter dem Titel „Wenn Schweine schwitzen“ interessante und unterhaltsame Anekdoten über die breite Vielfalt der Informatik und deren praktische Anwendungen in Alltag und Wirtschaft. Der anregende Vortrag stimulierte alle Anwesenden zu regen Diskussionen, die sich auch beim anschießenden gemütlichen Ausklang fortsetzten. Eine Fortführung der Reihe in 2017 ist bereits fest eingeplant.

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.

On the of October 28th, 2016, Prof. Holger Hermanns (University of the Saarland) started the event series with his talk titled “Wenn Schweine schwitzen” (“When Pigs Perspire”). By presenting interesting and entertaining anecdotes from a variety of computer science applications in everyday life, the talk stimulated a lively discussion among all participants. A continuation of the series in 2017 is already planned.

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 hauptsächlich in Saarbrücken.

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 primarily located in Saarbrücken.

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, einen 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.

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 über Nacht untergebracht werden 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 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.

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 Mikrophonen 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.

2016 wurde damit begonnen, einen zweiten großen Hörsaal zu schaffen. Im Rahmen des Umbaus 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.

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.

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.

2016 saw the beginning of construction works for a second large lecture hall. 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.

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.

Die beiden größten Hörsäle sind jeweils mit mehreren Tafeln ausgestattet, während in den anderen Tagungsräumen jeweils große Whiteboards an den Wänden montiert sind. In einem Seminarraum kann durch eine spezielle Wandfarbe sogar eine ganze Wand als Whiteboard (über $12\mathrm{m}^2$) benutzt werden.

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.

Dagstuhls 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.

Dagstuhls Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden jeden Tag von unseren 10 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 häufig 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 Frühstücksbüffet, dienstags bis donnerstags ein

Whereas the two main lecture halls are equipped with several blackboards, whiteboards are provided in the other rooms. One of the conference rooms features a complete "whiteboard wall" painted with a special paint which allows the use this whole wall (over $12\mathrm{m}^2$) as one large whiteboard.

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 10-person staff and 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 frequently invited 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. 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

Mittagsbuffet sowie ein Menü am Abend an. Montags und freitags wird aus logistischen Gründen auch am Mittag ein 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. In der Cafeteria und dem Weinkeller können Gäste Snacks erwerben. Abends gibt es in diesen beiden beliebten Räumen Brot und eine 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. Für Seminarteilnehmer übernimmt Schloss Dagstuhl die Kosten für Verpflegung und Unterkunft der Kinder und der Begleitperson.

Dagstuhls Angebot der Kinderbetreuung für Eltern wird immer mehr genutzt. Im Jahr 2016 wurden 18 Kinder durch eine Tagesmutter und 12 weitere durch Verwandte betreut. Insgesamt beherbergte Schloss Dagstuhl 30 Kinder von Teilnehmern an 23 Veranstaltungen während 20 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.

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 hall. 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 in the cafe and the wine cellar – two popular after-hours hangouts. Bread and cheese is served there 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 both for the children and the accompanying person.

Guests make increasing use of Dagstuhl's childcare offer for parents. In 2016, Dagstuhl hosted 30 children, 18 of whom were cared for by a nanny on site and 12 by relatives. Participants of 23 events in 20 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

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.

Dagstuhl Office at Saarbrücken

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

Die für die Bibliographiedatenbank dblp zuständigen wissenschaftlichen Mitarbeiter haben ihren Standort an der Universität Trier. Grund dafür ist die 2010 gestartete Zusammenarbeit zwischen Schloss Dagstuhl und der Universität Trier, die Ende 2016 um weitere zwei Jahre bis zum 31.Dezember 2018 verlängert wurde.

Dagstuhl at University of Trier

The scientific staff working on the *dblp computer science bibliography* is located at the Dagstuhl offices at the University of Trier. This is due to the cooperation between Schloss Dagstuhl and the University of Trier which was first established in November 2010 and was renewed until the December 31, 2018 in the end of 2016.

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 Bibliotheksverbünden 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 33 000 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 6 000 deutschlandweite und internationale Zeitungen und Magazine aus über 100 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 33,000 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 6,000 national and international newspapers and magazines from more than 100 countries.

Zusammenarbeit

Schloss Dagstuhls 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 ist 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. 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 Dagstuhls Fachbibliothek ist institutionelles Mitglied des Deutschen Bibliotheksverbandes (DBV).

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 is 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. 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).

Spenden an die Bibliothek

Die Bibliothek von Schloss Dagstuhl profitiert von zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2016 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 684 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. In 2016, 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 684 volumes during the year 2016 as donations from publishing houses and seminar participants.

Birkhäuser Verlag http://www.birkhaeuser-science.com
SIAM – Society for Industrial and Applied Mathematics http://www.siam.org
Springer-Verlag GmbH | Springer Science+Business Media http://www.springer.com

Fig. 9.1
Donations from publishers to the Dagstuhl library.

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. Im Vergleich zu anderen Forschungsorganisationen stellt das Ausrichten von Informatik-Seminaren die IT-Dienste vor keine besonderen Herausforderungen, abgesehen davon, dass große Nachfrage nach einer drahtlosen Internetverbindung besteht. Teilnehmer bringen in der Regel Laptop und Smartphone mit und erwarten immer und überall einfachen, schnellen und verlässlichen Zugang zum Internet.

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-Leitungen sichergestellt, die durch den DFN e.V. (Deutsches Forschungsnetz) betrieben werden.
- 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 (http://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 (http://dblp.dagstuhl.de und http://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, TeamDrive (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. Hosting seminars in computer science does not create unusual demands in IT services compared with other research organizations, except that wireless connectivity to the internet is valued at a premium. Attendees routinely bring along a laptop and a smartphone each and expect easy, fast, and reliable access to the internet anywhere and at any time.

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* service[43] (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 connections that are managed by DFN e.V. (National Science Network).
- 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 http://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 http://dblp.dagstuhl.de and at http://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, TeamDrive (a cloud-based storage system), and several others.

[43] *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 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 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). Am Freitag, 9. September 2016 fand in der Galerie der HBKsaar die Kick-Off-Veranstaltung „Troika für die Kunst“ für diese Kooperation statt.

Die drei Ausstellungen (siehe Fig. 10.1), die im Jahr 2016 stattfanden, sind nachfolgend beschrieben. Die jeweils aktuellen Ausstellungen sind nach Anmeldung auch für die interessierte Öffentlichkeit zugänglich.

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 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. 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. Being a major art sponsor, Saartoto is in possession of a substantial art collection, which Schloss Dagstuhl can now access to create a series of exhibitions. In the context of this collaboration, HBKsaar inventories and documents Saartoto's art collection, while simultaneously, there were, and will be, Dagstuhl exhibitions where the 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). On Friday, September 9, 2016, a kick-off event called “Troika für die Kunst” (a troika for art) was held at HBKsaar.

The three exhibitions (cf. Fig. 10.1) hosted by Schloss Dagstuhl in 2016 are described below. Current exhibitions are open to the interested public upon request.

»August Clüsserath (1899 – 1966)« Works by August Clüsserath curated by Beate Kolodziej, M.A. | January 11 to April 8, 2016
»Farbe und Form, Abstraktion und Expression« Works from a private art collection curated by Beate Kolodziej, M.A | May 30 to July 22, 2016
»Zarte Linien, Starke Flächen« Works from the art collection of Saartoto curated by Nadine Brettar | September 19 to December 16, 2016

Fig. 10.1
Art exhibitions in 2016.

»August Clüsserath (1899 – 1966)«

August Clüsserath wurde 1899 in Völklingen-Fenne geboren und studierte von 1926 bis 1932 an der Staatlichen Schule für Kunst und Kunstgewerbe Saarbrücken. Unter anderem war er Mitbegründer der „Neuen Gruppe Saar".

August Clüsserath zählt zu den wichtigsten Künstlern des Saarlandes der zweiten Hälfte des 20. Jahrhunderts. Zunächst von gegenständlichen Strömungen wie der Neuen Sachlichkeit inspiriert, gelangte er nach der Auseinandersetzung mit Positionen des Bauhauses und des Kubismus ab Ender der 1950er Jahre zu einem vollkommen abstrakten gestischen Stils. Formen lösen sich auf, Linien und Strukturen bewegen sich rhythmisch frei in seinen Bildern. Diese Arbeiten der letzten Sachaffensperiode Clüsseraths waren Schwerpunkt der von der Kunsthistorikerin Beate Kolodziej kuratierten Ausstellung.

»Farbe und Form, Abstraktion und Expression«

In der von der Kunsthistorikerin Beate Kolodziej kuratierten Ausstellung aus dem privaten Fundus eines baden-württembergischen Sammlers wurden vor allem Arbeiten von Künstlern aus dem süddeutschen-Raum sowie auch aus Italien und China gezeigt.

Die knapp 40 in der Ausstellung gezeigten Gemälde und Grafiken vermitteln den Querschnitt einer leidenschaftlichen Sammlertätigkeit, die ihren Anfang Mitte der 1950er Jahre nahm, und die den Zeitgeist und die künstlerischen Strömungen jener Zeit einfängt. Die Künstler, fast alle in den ersten Jahrzehnten des 20. Jahrhunderts geboren. entfalten die ganze Vielfalt an Möglichkeiten des Mediums Malerei. Man findet gegenständliche Arbeiten, die Tendenzen des Fauvismus und Expressionismus aufgreifen, ebenso wie Werke mit abstrakten Kompositionen von Gemälden und Collagen mit surrealen Motiven. Auch trifft konzeptuelle Malerei auf fernöstliche Ungegenständlichkeit.

»Zarte Linien, Starke Flächen«

Dies war die erste Ausstellung im Rahmen der Zusammenarbeit von Saartoto, HBKsaar und Schloss Dagstuhl. Kuratiert von Nadine Brettar wurde vorübergehend ein kleiner Teil der sich seit fünf Jahrzehnten entwickelnden Grafikkabinetts des regionalen Kunstförderers Saartoto gezeigt.

Die gesammelten Radierungen, Holz-, Linolschnitte und Siebdrucke treten in mehreren Etappen in einen erfrischenden Dialog mit Druckgrafiken zeitgenössischer Künstler aus Luxemburg, Frankreich und Belgien. Die aktuelle Edition „Drucksachen III" die Studierende unter Anleitung von Frau Prof. Langendorf, Rektorin der HBKsaar, in den Werkstätten von Ulrich Kerker und Dirk Rausch, geschaffen haben, runden das Ensemble der Exponate ab. Die Druckgrafik präsentiert sich, in ihrer Fülle an Verfahrensmöglichkeiten sowie experimentellen und kreativen Freiräumen, als hochaktuelles künstlerisches Medium, dessen Reiz auch im Zeitalter hochtechnisierter Repro-

»August Clüsserath (1899 – 1966)«

August Clüsserath was born in Völklingen-Fenne in 1899, and studied at Staatliche Schule für Kunst und Kunstgewerbe (federal school of arts and crafts) in Saarbrücken from 1926 to 1932. He was co-founder of "Neuen Gruppe Saar", a group of artists based on Bauhaus ideas.

August Clüsserath is one of Saarland's most significant artists of the second half of the 20th century. At first inspired by concrete movements such as New Objectivity, he dealt with Bauhaus ideas and cubism, and finally found an entirely abstract, gestural style in the late 1950s. Shapes dissolve, lines and structures move rhythmically unbound in his paintings. Curated by art historian Beate Kolodziej, the exhibition's focus was on these works from the last period of Clüsserath's work.

»Farbe und Form, Abstraktion und Expression«

The exhibition, curated by art historian Beate Kolodziej, mainly contained works from southern Germany, but also from Italy and China. They belong to a private collector based in Baden-Wuerttemberg.

The almost 40 paintings and graphics represent a cross-section of passionate collecting activities that started out in the middle of the 1950s. Furthermore, they capture the zeitgeist and artistic movements of those years. The artists, almost all born in the first decades of the 20th century, cover the wide variety of possibilities that come with the medium of painting. There are concrete works, addressing tendencies of fauvism and expressionism, as well as works with abstract compositions of paintings, and collages with surreal themes. Aside from that, the exhibition combines conceptual painting with Far Eastern nonrepresentational art.

»Zarte Linien, Starke Flächen«

This was the first exhibition created in the context of the cooperation between Saartoto, HBKsaar, and Schloss Dagstuhl. Curated by Nadine Brettar, it showed a small part of Saartoto's graphics collection that has been evolving for five decades for a limited period of time.

In several steps, the etchings, woodcuts, linocuts, and screen prints enter into a dynamic dialogue with printed graphics by contemporary artists from Luxembourg, France, and Belgium. The exhibits' ensemble is rounded out by the latest edition of "Drucksachen III" (printed papers), created by students under the guidance of Prof. Langendorf, rector of HBKsaar, in Ulrich Kerker and Dirk Rausch's studios. With their variety of procedures as well as experimental and creative freedom, printed graphics prove to be a highly topical artistic medium appealing to generations of artists and art enthusiasts, even in an era of high-tech reproductions.

duktionsmöglichkeiten Generationen von Künstlern und Kunstliebhabern besticht.

10.2 Kunstankauf durch Spenden

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 2016 erhielt Schloss Dagstuhl insgesamt 380 Euro von verschiedenen Spendern. Wir möchten diese Stelle nutzen, allen Spendern, die 2016 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Nähere Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite[44] von Dagstuhl.

10.2 *Art Sponsorship and Donations*

Dagstuhl's website contains a page featuring an internet gallery enabling participants, individuals, and groups to make contributions to Dagstuhl for art donations. The works of art are featured 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 on 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 2016, Schloss Dagstuhl received a total of 380€ from various donors. We would like to thank all donors who contributed to Dagstuhl's art collection in 2016.

For further information and current news about Dagstuhl's art program, please visit Dagstuhls's art webpage[44].

10.3 Dagstuhls permanente Kunstausstellung

Die von Gästen immer wieder positiv hervorgehobene Kunstsammlung geht auf den Gründungsdirektor Professor Wilhelm zurück. Seine Idee war es, den 1995 neueröffneten Speisesaal und den etwa ein Jahr älteren Neubau, durch Kunstwerke zu beleben. Dazu startete er die oben beschrieben Kunstaustellungen. Unter Mitwirkung der Künstler wird aus jeder Ausstellung ein Werk ausgewählt, für das dann Spender gesucht werden. In den letzten 20 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.

10.3 *Dagstuhl's Permanent Art Exhibition*

The art collection, continually praised by guests, was intiated by Founding Director Professor Wilhelm. It was his idea to use works of art in order to enliven the 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 120 works of art could be acquired over the last 20 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.

[44] http://www.dagstuhl.de/art/

11 Struktur der Gesellschaft
Structure of the Company

Gründung und Gesellschafter 11.1 *Formation and Shareholders*

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit derzeit 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 evaluiert. Die Ergebnisse dieser Evaluierung werden im ersten Halbjahr 2017 erwartet.

Schloss Dagstuhl is operated as a non-profit organization by currently eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e. V.[45] (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 Forschungsförderung[46] (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”[47] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[48] 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 again been evaluated. The results of this evaluation are expected in the first half of 2017.

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.

[45] engl.: German Informatics Society
[46] engl.: Federal/State Government Commission for Educational Planning and Research Promotion
[47] engl.: International Conference and Research Center for Computer Science
[48] 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 Aufsichtsratmitglieder beträgt mindestens vier volle, abgeschlossene Geschäftsjahre und endet mit der Entlastung des Aufsichtsrats 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. Er ist Mitglied des Wissenschaftlichen Direktoriums und leitet dieses. Seit Mai 2014 ist Prof. Raimund Seidel, Ph. D., der wissenschaftliche Direktor von Schloss Dagstuhl.

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 other shareholders in proportion to its larger number of shares. Decisions are made in shareholders' meetings which take place at least once the 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 each of the three founding universities, two representatives of the four universities that subsequently joined, and one representative 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 of the actions of the Supervisory Board 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. He is a member and the chairperson of the Scientific Directorate. Since May 2014, Prof. Raimund Seidel, Ph.D., is the Scientific Director of Schloss Dagstuhl.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee

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.

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 auch mehrfach möglich.

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 twice-yearly 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. Multiple reelections are possible.

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 Wissenschaftliche Beirat auf zwei Wegen bei der Qualitätssiche-

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 which the Scientific Advisory Board is involved in quality assurance.

rung 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 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.

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 of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board and the Senatsausschuss.

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. Doreen Becker Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government | *tenure ended in August 2016*
Prof. Alejandro P. Buchmann, Ph. D. Technische Universität Darmstadt, Germany | Representative of Technische Universität Darmstadt
Dr. Christian Heimann Ministerium für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state | *tenure ended in November 2016*
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.
Prof. Dr. Volker Linneweber Universität des Saarlandes, Saarbrücken, Germany | Representative of Universität des Saarlandes
Prof. Dr. Erhard Plödereder Universität Stuttgart, Germany | Representative of Universität Stuttgart
Prof. Dr. Arnd Poetzsch-Heffter Technische Universität Kaiserslautern, Germany | Representative of Technische Universität Kaiserslautern
Alexander Rabe Gesellschaft für Informatik e. V., Berlin, Germany | Representative of Gesellschaft für Informatik e. V. | *tenure ended in September 2016*
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | Representative of the Saarland
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. | *tenure started in October 2016*
Marcus Wittrin Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government | *tenure started in October 2016*
N. N. Ministerium für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state | *tenure will start in May 2017*

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 Institue, Madrid, Spain
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Stefan Diehl Universität Trier, Germany
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany | *tenure started in November 2016*
Prof. Dr. Hans Hagen Technische Universität Kaiserslautern, Germany
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany
Prof. Dr.-Ing. Oliver Kohlbacher Eberhard Karls Universität Tübingen, Germany
Dr. Stephan Merz Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France
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
Prof. Dr. Albrecht Schmidt Universität Stuttgart, Germany | *tenure started in November 2016*
Prof. Dr. Nicole Schweikardt Humboldt-Universität zu Berlin, Germany | *tenure ended in October 2016*
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Ir. Arjen P. de Vries Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
Prof. Dr.-Ing. Klaus Wehrle Rheinisch-Westfälische Technische Hochschule Aachen, Germany
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Universität des Saarlandes, Saarbrücken, Germany | *tenure ended in October 2016*
Prof. Dr. Verena Wolf Universität des Saarlandes, Saarbrücken, Germany | *tenure started in November 2016*

Fig. 11.4
Scientific Directorate.

Wissenschaftlicher Beirat | Scientific Advisory Board
Prof. Dr. Christel Baier Technische Universität Dresden, 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 HITS GmbH, Heidelberg, Germany
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | *Guest*

Fig. 11.5
Scientific Advisory Board.

Kuratorium | Industrial Curatory Board
Dr. Udo Bub EIT ICT Labs, Berlin, Germany
Dr.-Ing. Elmar Dorner SAP SE, Karlsruhe, Germany
Dr. Jo Ebergen Oracle Labs, Redwood Shores, United States
Dr.-Ing. Uwe Franke Daimler AG, Böblingen, Germany
Dr. Goetz Graefe Google, Madison, Wisconsin, United States
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“

Association “Friends of Dagstuhl”

Holger Hermanns (Universität des Saarlandes, Germany)
Erich Reindel (Universität des Saarlandes, Germany)

Seit dem 6. Mai 2014 haben die Freunde von Dagstuhl endlich eine Heimat. An diesem Tag fanden sich 16 Freunde von Schloss Dagstuhl zusammen, um den Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V. zu gründen. Der sehr technische und holprig klingende Name spiegelt aber exakt den Vereinszweck wider: die Förderung von Wissenschaft und Forschung im Leibniz-Zentrum für Informatik in Schloss Dagstuhl. Für die Webpräsenz wurde mit „Friends of Dagstuhl“ ein wesentlich geschmeidigerer Name gewählt (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.

Im Jahr 2016 wurden entscheidende Schritte zur Anlage des Stiftungsvermögens vorgenommen. Nachdem zuvor unter Einbeziehung des Stiftungsrates (siehe Fig. 12.2) verschiedene Möglichkeiten geprüft wurden, das Kapital trotz der andauernden Niedrigzinsphase sicher und dennoch nicht ganz ohne Rendite anzulegen, wurden Verträge mit einer professionellen und auf Stiftungskapital spezialisierten Vermögensverwaltung geschlossen.

Inzwischen gehören dem Verein 38 persönliche sowie 4 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.

As of May 6, 2014, Dagstuhl supporters finally have a home. On that day, 16 friends of Dagstuhl gathered in order to found the registered association in support of Schloss Dagstuhl – Leibniz Center for Informatics (Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.). 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).

In 2016, crucial steps were taken regarding the foundation assets. Under involvement of the foundation council (see Fig. 12.2), the association’s chair evaluated several opportunities 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.

Currently, the association has 38 individual and 4 institutional members. Especially with regard to the small number of institutional members, Friends of Dagstuhl look forward to welcoming 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
Kurt Mehlhorn Max Planck Institute for Informatics (MPII), Saarbrücken, Germany
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. 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 eingereichter 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 Seminarteilnehmer, 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 participant 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 fraction 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 these 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.

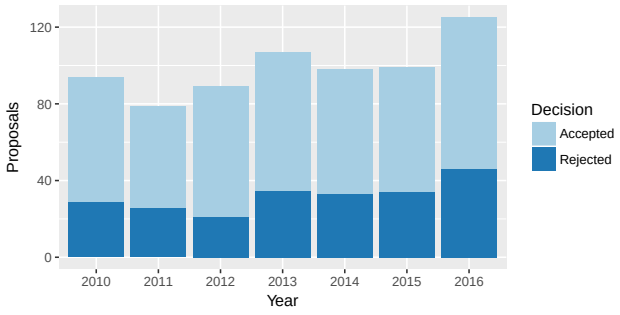

(a) Chart for 2010–2016

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
2010	94	65	69.1	29	30.9
2011	79	53	67.1	26	32.9
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

(b) Detailed numbers for 2010–2016

Fig. 13.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

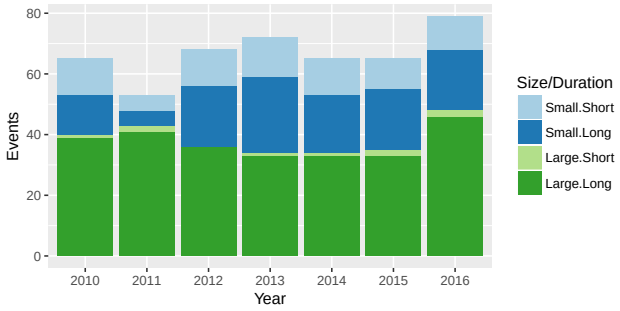

(a) Chart for 2010–2016

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2010	12	13	1	39	65
2011	5	5	2	41	53
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

(b) Detailed numbers for 2010–2016

Fig. 13.2
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2010–2016. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

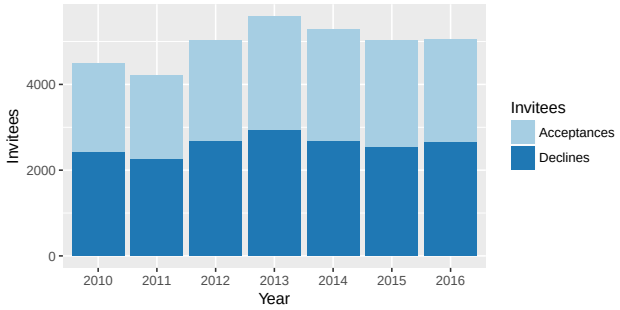

(a) Chart for 2010–2016

Year	Invitees	Acceptances		Declines	
	#	#	%	#	%
2010	4499	2053	45.6	2446	54.4
2011	4223	1958	46.4	2265	53.6
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

(b) Detailed numbers for 2010–2016

Fig. 13.3
Total number of invitees, acceptances, and declines for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

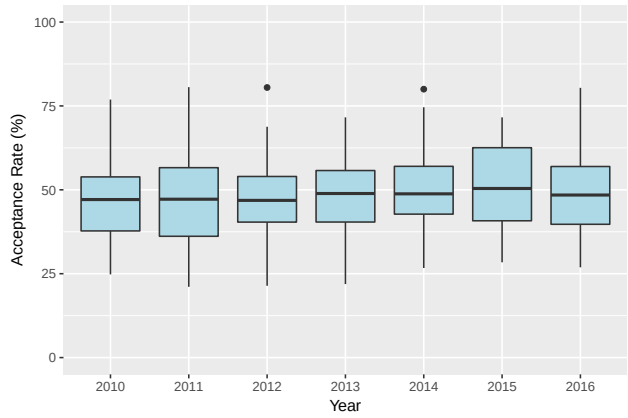

(a) Chart for 2010–2016

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2010	24.8	76.9	46.5	12.2
2011	21.1	80.6	47.7	14.0
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

(b) Detailed numbers for 2010–2016

Fig. 13.4
Distribution of the acceptance rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2010–2016. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

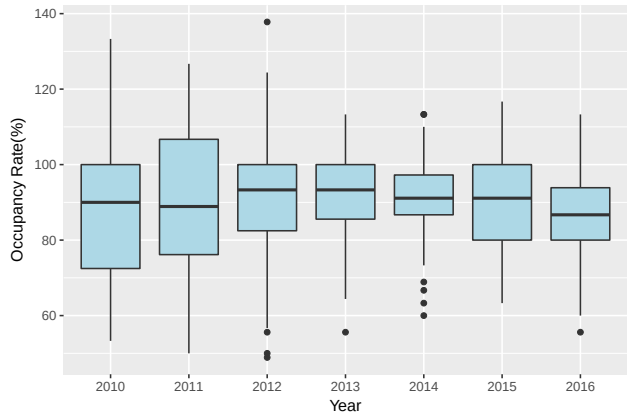

(a) Chart for 2010–2016

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2010	53.3	133.3	85.9	18.0
2011	50.0	126.7	90.4	19.3
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	89.8	12.3
2016	55.6	113.3	86.7	11.8

(b) Detailed numbers for 2010–2016

Fig. 13.5
Distribution of the occupancy rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2010–2016. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

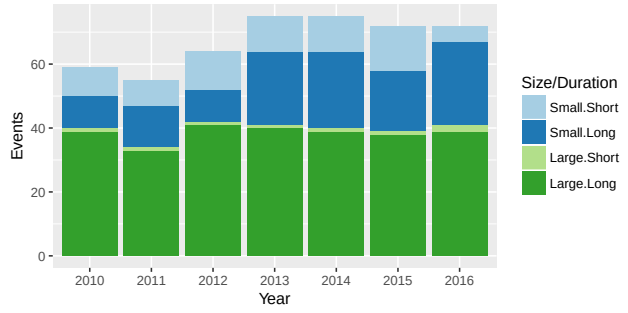

(a) Chart for 2010–2016

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2010	9	10	1	39	59
2011	8	13	1	33	55
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

(b) Detailed numbers for 2010–2016

Fig. 13.6
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2010–2016. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

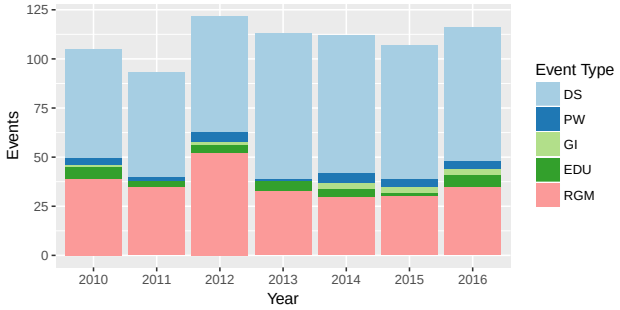

(a) Chart for 2010–2016

Year	DS	PW	GI	EDU	RGM	Total
2010	55	4	1	6	39	105
2011	53	2	0	3	35	93
2012	59	5	2	4	52	122
2013	74	1	0	5	33	113
2014	70	5	3	4	30	112
2015	68	4	3	2	30	107
2016	68	4	3	6	35	116

(b) Detailed numbers for 2010–2016

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.

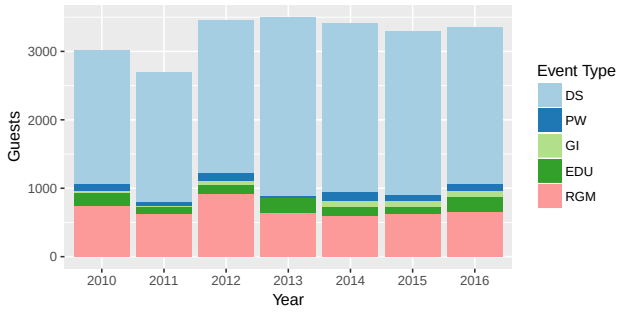

(a) Chart for 2010–2016

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	#
2010	1950	64.7	103	3.4	25	0.8	192	6.4	743	24.7	3013
2011	1894	70.2	64	2.4	0	0.0	103	3.8	637	23.6	2698
2012	2226	64.4	120	3.5	48	1.4	144	4.2	916	26.5	3454
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503
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

(b) Detailed numbers for 2010–2016

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
Germany	576	735	1311
United States	530	35	565
United Kingdom	208	22	230
France	173	7	180
Netherlands	111	6	117
Switzerland	71	17	88
Austria	55	12	67
Israel	62	2	64
Denmark	50	9	59
Canada	50	8	58
Belgium	47	6	53
Italy	50	3	53
Japan	49	2	51
Australia	40	5	45
Sweden	40	4	44
Luxembourg	13	30	43
Finland	24	12	36
India	29	1	30
Norway	20	1	21
Poland	17	3	20
Spain	16	4	20
Brazil	16	2	18
Pakistan	1	16	17
Czech Republic	13	2	15
China	10	4	14
Greece	11	0	11
Chile	10	0	10
Ireland	10	0	10
Russian Federation	10	0	10
Slovenia	7	1	8
Hong Kong	7	0	7
Hungary	6	1	7
Republic of Korea	7	0	7
Singapore	7	0	7
Portugal	6	0	6
Croatia	4	1	5
New Zealand	5	0	5
Serbia	0	4	4
Taiwan	4	0	4
Iran	1	2	3
Qatar	3	0	3
Romania	3	0	3
Slovak Republic	3	0	3
Thailand	3	0	3
Turkey	3	0	3
Estonia	2	0	2
Iceland	2	0	2
Mexico	2	0	2
Ukraine	1	1	2
Algeria	1	0	1
Argentina	1	0	1
Bangladesh	1	0	1
Morocco	0	1	1
Nigeria	0	1	1
South Africa	1	0	1
United Arab Emirates	1	0	1
Total	2393	960	3353

(a) Details for 2016 by country

Region	A		B		Total	
	#	%	#	%	#	%
Germany	576	24.1	735	76.6	1311	39.1
Europe (w/o Germany)	976	40.8	146	15.2	1122	33.5
North America	580	24.2	43	4.5	623	18.6
Asia	185	7.7	27	2.8	212	6.3
Australia	45	1.9	5	0.5	50	1.5
South America	29	1.2	2	0.2	31	0.9
Africa	2	0.1	2	0.2	4	0.1
Total	2393	100	960	100	3353	100

(b) Details for 2016 by region

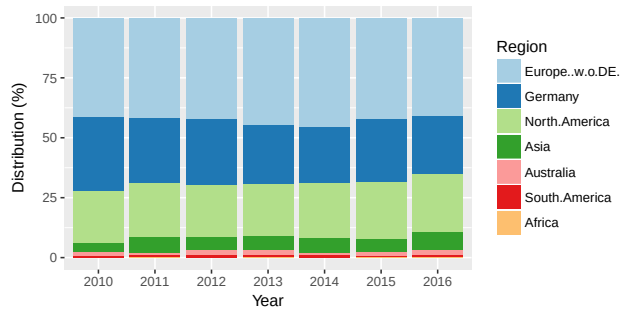

(c) Graphical distribution of seminar type A in 2010–2016 by year and region

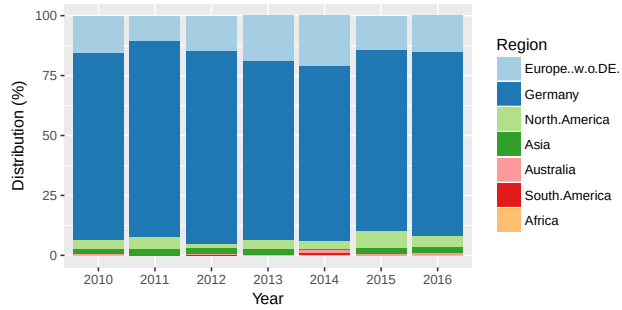

(d) Graphical distribution of seminar type B in 2010–2016 by year and region

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).

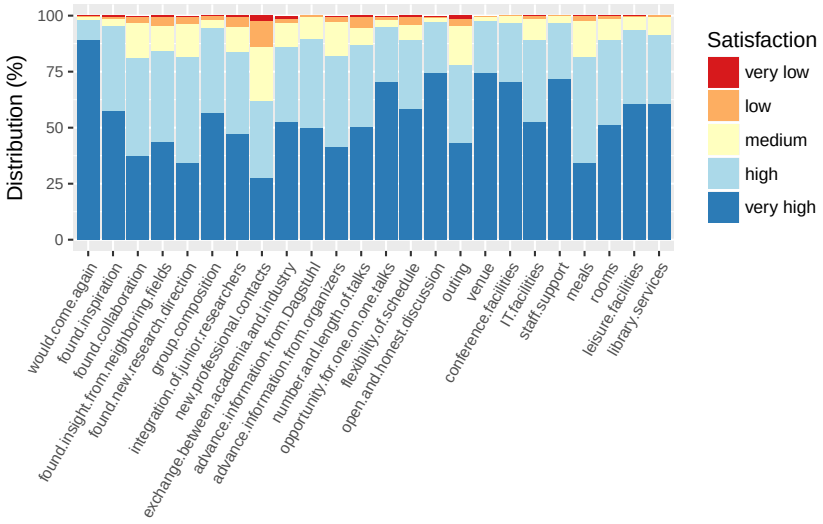

(a) Graphical distribution for 2016

	2010	2011	2012	2013	2014	2015	2016	2016 – 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	1	4	22	130	1295	1452
found inspiration	4.3	4.4	4.4	4.5	4.4	4.5	4.5	7	12	46	555	835	1455
found collaboration	3.9	4.0	4.0	4.0	4.1	4.1	4.1	6	43	225	624	543	1441
found insight from neighboring fields	4.0	4.1	4.1	4.1	4.2	4.3	4.2	8	55	166	585	632	1446
found new research direction	3.9	4.0	4.0	4.0	4.0	4.1	4.1	5	49	207	684	495	1440
group composition	4.4	4.4	4.4	4.4	4.4	4.5	4.5	4	20	56	551	828	1459
integration of junior researchers	4.2	4.2	4.2	4.2	4.2	4.2	4.3	9	58	163	530	677	1437
new professional contacts	3.5	3.7	3.6	3.7	3.7	3.6	3.7	34	161	346	487	394	1422
exchange between academia and industry	4.3	4.1	4.3	4.2	4.2	4.3	4.3	11	19	96	301	473	900
advance information from Dagstuhl	4.3	4.3	4.4	4.4	4.4	4.4	4.4	0	9	136	579	727	1451
advance information from organizers	4.0	4.2	4.1	4.1	4.1	4.1	4.2	6	33	216	580	591	1426
number and length of talks	4.1	4.1	4.1	4.2	4.1	4.2	4.3	7	69	114	530	733	1453
opportunity for one on one talks	4.5	4.5	4.4	4.5	4.5	4.5	4.6	6	25	42	358	1027	1458
flexibility of schedule	4.3	4.2	4.2	4.2	4.3	4.3	4.4	7	49	100	442	841	1439
open and honest discussion	4.7	4.7	4.6	4.7	4.7	4.7	4.7	5	8	25	332	1081	1451
outing	4.0	4.2	4.1	4.1	4.1	4.1	4.2	16	37	199	395	497	1144
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	5	31	336	1086	1458
conference facilities	4.7	4.8	4.7	4.6	4.7	4.6	4.7	0	3	41	385	1021	1450
IT facilities	4.5	4.6	4.4	4.4	4.4	4.3	4.4	2	14	126	476	683	1301
staff support	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	2	39	352	1004	1397
meals	4.2	4.3	4.2	4.1	4.1	4.1	4.1	4	27	239	682	497	1449
rooms	4.5	4.5	4.4	4.4	4.4	4.4	4.4	3	15	137	549	741	1445
leisure facilities	4.5	4.6	4.6	4.6	4.6	4.6	4.5	1	7	75	429	798	1310
library services	4.6	4.5	4.5	4.5	4.5	4.5	4.5	0	4	55	213	417	689

(b) Averages for 2010–2016 and detailed numbers for 2016: 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.

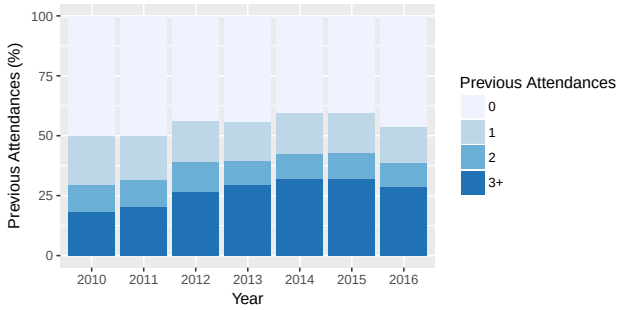

(a) Graphical distribution for 2010–2016

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
2010	442	50	185	21	98	11	162	18	887
2011	413	50	154	19	94	11	168	20	829
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

(b) Detailed numbers for 2010–2016

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.

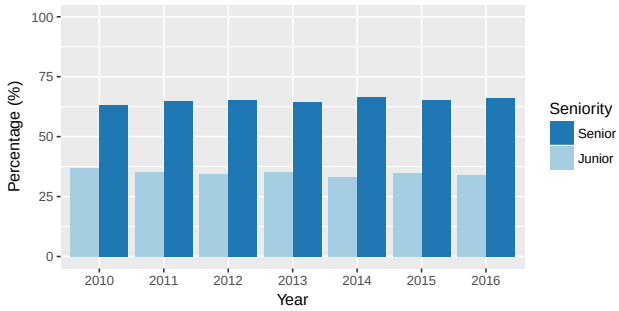

(a) Chart for 2010–2016

Year	Junior		Senior		Total
	#	%	#	%	#
2010	291	36.8	500	63.2	791
2011	266	35.2	489	64.8	755
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

(b) Detailed numbers for 2010–2016

Fig. 13.12

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

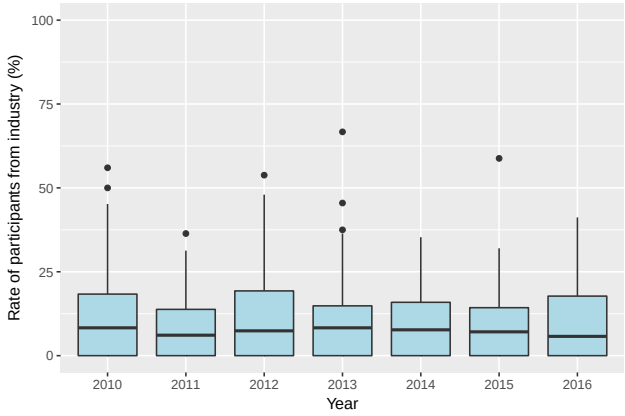

(a) Chart for 2010–2016

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2010	0.0	56.0	12.2	13.1
2011	0.0	36.4	8.2	9.1
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

(b) Detailed numbers for 2010–2016

Fig. 13.13

Distribution of the rate of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2010–2016, 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”.

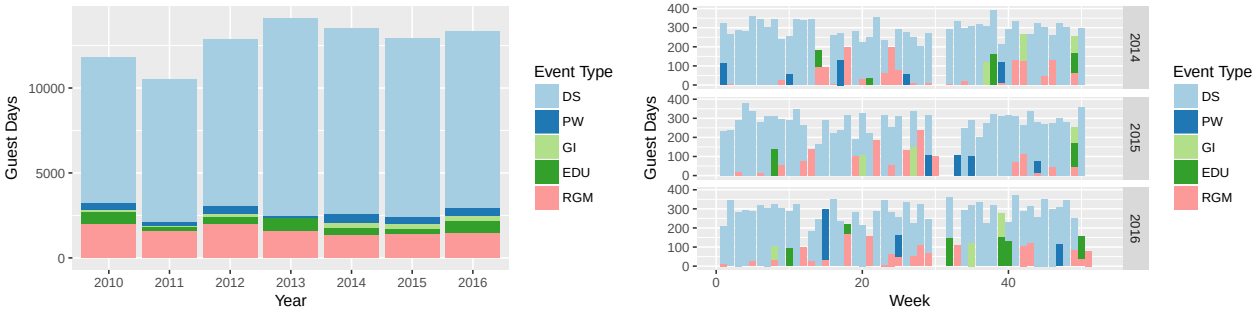

(a) Chart for 2010–2016

(b) Graphical distribution for 2014–2016 by week

Year	DS	PW	GI	EDU	RGM	Total
2010	8572	381	125	722	2002	11802
2011	8415	228	0	266	1604	10513
2012	9798	458	190	393	2031	12870
2013	11612	130	0	753	1614	14109
2014	10939	475	348	390	1370	13522
2015	10491	380	344	261	1424	12900
2016	10362	495	315	703	1462	13337

(c) Detailed numbers for 2010–2016

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.

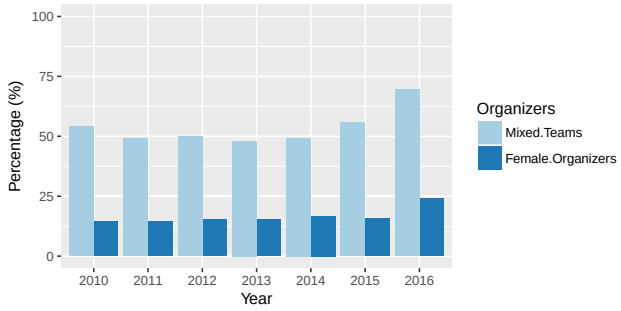

(a) Chart for 2010–2016

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2010	59	32	54.2	233	34	14.6
2011	55	27	49.1	213	31	14.6
2012	64	32	50.0	256	39	15.2
2013	75	36	48.0	282	43	15.2
2014	75	37	49.3	303	51	16.8
2015	72	40	55.6	284	45	15.8
2016	72	50	69.4	278	67	24.1

(b) Detailed numbers for 2010–2016

Fig. 13.15

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

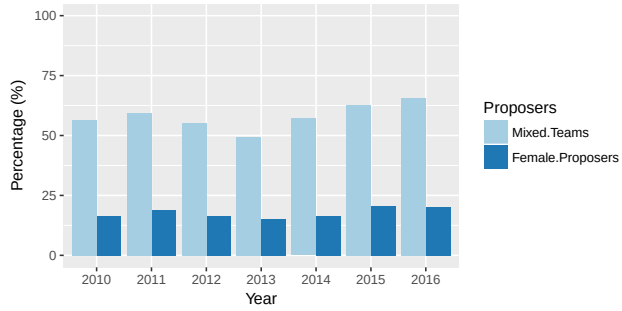

(a) Chart for 2010–2016

	Proposer Teams			Proposers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2010	94	53	56.4	366	60	16.4
2011	79	47	59.5	311	59	19.0
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

(b) Detailed numbers for 2010–2016

Fig. 13.16
Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.

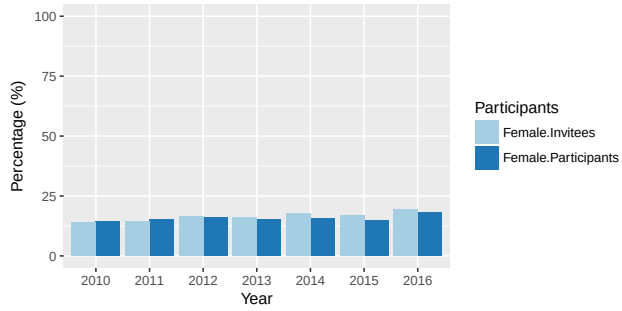

(a) Chart for 2010–2016

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2010	4499	632	14.0	2053	294	14.3
2011	4223	604	14.3	1958	295	15.1
2012	5033	821	16.3	2346	377	16.1
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	977	19.3	2393	436	18.2

(b) Detailed numbers for 2010–2016

Fig. 13.17
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

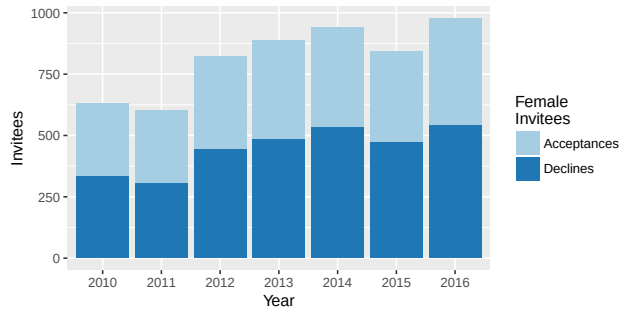

(a) Chart for 2010–2016

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2010	632	294	46.5	338	53.5
2011	604	295	48.8	309	51.2
2012	821	377	45.9	444	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	977	436	44.6	541	55.4

(b) Detailed numbers for 2010–2016

Fig. 13.18
Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.

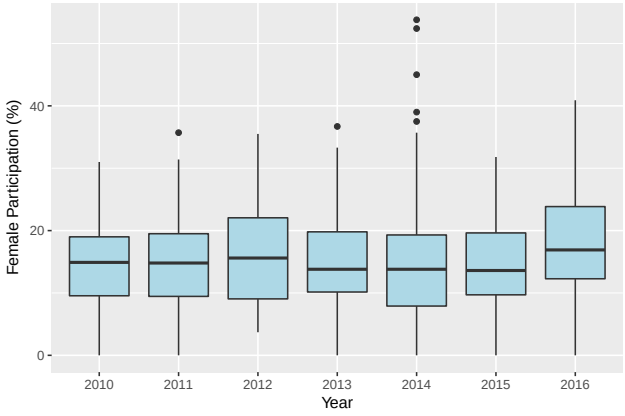

(a) Chart for 2010–2016

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2010	0.0	31.0	14.5	7.7
2011	0.0	35.7	14.7	7.4
2012	3.7	35.5	16.1	7.7
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

(b) Detailed numbers for 2010–2016

Fig. 13.19

Distribution of female participants rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2010–2016. 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.

This section provides statistical data about the dblp computer science bibliography. Fig. 13.20 show the average usage statistics of the dblp servers of the past years. An overview of the development of the dblp database can be found in Fig. 13.21 and Fig. 13.22.

	Trier 1			Trier 2			Dagstuhl		
	2014	2015	2016	2014	2015	2016	2014	2015	2016
user sessions (visits) per day	21,057	28,327	26,911	4,703	662	1,428	326	510	1,254
page views per day	174,247	452,089	501,208	47,531	8,839	26,355	14,964	14,868	35,406
page views per user session	8.2	15.9	18.6	10.1	13.3	18.4	45.8	29.1	28.2
distinct users (IPs) per month	327,299	416,413	393,273	76,566	11,474	25,249	4,399	7,241	20,416
data served per month	825.2 GB	861.8 GB	1,187.6 GB	345.8 GB	22.0 GB	72.7 GB	27.2 GB	75.8 GB	120.7 GB

Fig. 13.20

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

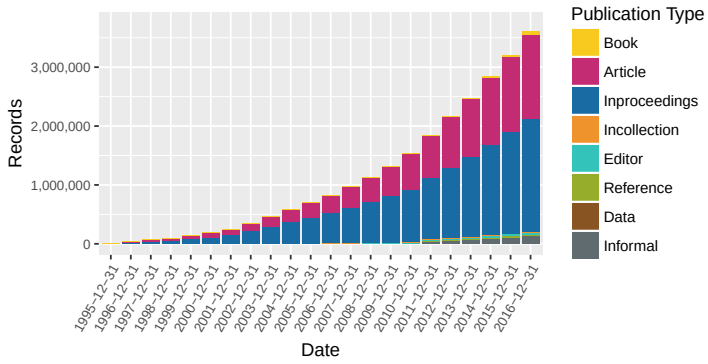

(a) Chart for 1996–2016

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2010	1,436	0.1	606,746	39.6	893,940	58.4	15,117	1.0	14,709	1.0	0	0.0	0	0.0	47	0.0	1,531,995
2011	9,356	0.5	705,716	38.3	1,048,640	57.0	7,453	0.4	17,151	0.9	12,207	0.7	0	0.0	39,925	2.2	1,840,448
2012	16,037	0.7	850,578	39.3	1,199,845	55.5	9,631	0.4	20,154	0.9	13,125	0.6	0	0.0	53,809	2.5	2,163,179
2013	16,816	0.7	997,549	40.2	1,350,620	54.4	12,797	0.5	22,770	0.9	13,125	0.5	0	0.0	69,896	2.8	2,483,573
2014	17,530	0.6	1,128,921	39.8	1,544,958	54.5	14,470	0.5	26,136	0.9	14,690	0.5	0	0.0	88,201	3.1	2,834,906
2015	18,314	0.6	1,280,881	40.0	1,724,145	53.9	16,288	0.5	30,042	0.9	19,103	0.6	12	0.0	110,955	3.5	3,199,740
2016	51,049	1.4	1,428,987	39.7	1,912,771	53.1	19,772	0.5	33,779	0.9	20,174	0.6	26	0.0	134,324	3.7	3,600,882

(b) Detailed numbers for 2010–2016

Fig. 13.21
Development of the total size of the dblp database.

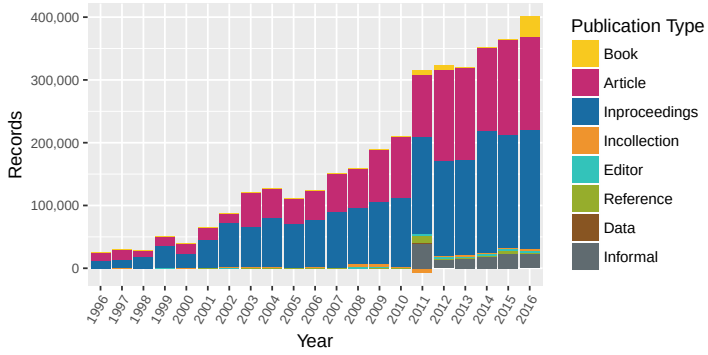

(a) Chart for 1996–2016

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2010	86	0.0	97,941	46.6	109,951	52.3	591	0.3	1,543	0.7	0	0.0	0	0.0	0	0.0	210,112
2011	7,920	2.6	98,970	32.1	154,700	50.2	−7,664	-2.5	2,442	0.8	12,207	4.0	0	0.0	39,878	12.9	308,453
2012	6,681	2.1	144,862	44.9	151,205	46.9	2,178	0.7	3,003	0.9	918	0.3	0	0.0	13,884	4.3	322,731
2013	779	0.2	146,971	45.9	150,775	47.1	3,166	1.0	2,616	0.8	0	0.0	0	0.0	16,087	5.0	320,394
2014	714	0.2	131,372	37.4	194,338	55.3	1,673	0.5	3,366	1.0	1,565	0.4	0	0.0	18,305	5.2	351,333
2015	784	0.2	151,960	41.7	179,187	49.1	1,818	0.5	3,906	1.1	4,413	1.2	12	0.0	22,754	6.2	364,834
2016	32,735	8.2	148,106	36.9	188,626	47.0	3,484	0.9	3,737	0.9	1,071	0.3	14	0.0	23,369	5.8	401,142

(b) Detailed numbers for 2010–2016

Fig. 13.22
Distribution of newly included publications in dblp. The negative number of new *Incollection* records in 2011 results from relabeling several thousand existing records with the newly introduced *Reference* type. Similarily, in the same year, several thousand *Article* and *Inproceedings* records have been labeled as *Informal*.

Statistiken zu 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. 13.23 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen, Fig. 13.24 die der Publikationen in der Reihe Dagstuhl Manifestos und schließlich Fig. 13.25 die der veröffentlichten Bände in der Reihe Dagstuhl Follow-Ups.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 13.26 fasst die Daten in der Reihe OASIcs und Fig. 13.27 die der Reihe LIPIcs zusammen.

Fig. 13.28 fasst die Kennzahlen der Zeitschrift LITES zusammen.

Die verschiedenen Publikationsserien wurden in verschiedenen Jahren zwischen 2009 und 2015 gegründet. Wir stellen in den Statistiken dennoch stets den gesamten Zeitraum dar.

13.3 Statistics of Dagstuhl Publishing

In this section the statistical data of Dagstuhl Publishing are presented.

The first three figures present the development of the seminar-focussed series: Fig. 13.23 summarizes the data of the periodical Dagstuhl Reports, Fig. 13.24 the data of the Dagstuhl Manifestos series, and, finally, Fig. 13.25 those of the volumes published in the Dagstuhl Follow-Ups series.

The statistical data to the service-focussed series are presented afterwards. Fig. 13.26 presents numbers related to OASIcs and Fig. 13.27 numbers related to LIPIcs.

We summarize the publications of the journal LITES in Fig. 13.28.

Please note that the publication series were established in different years, but all in the period between 2009 and 2015. However, we always consider this complete period.

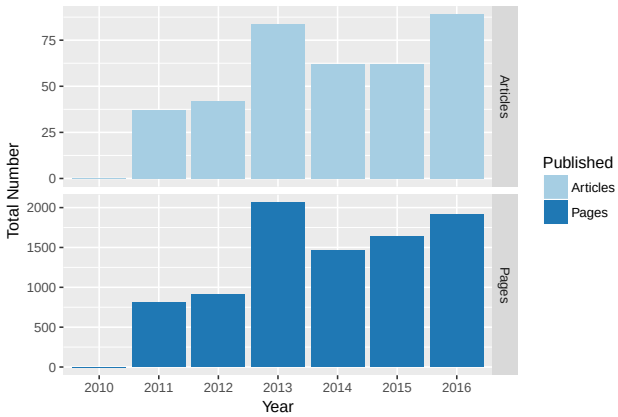

(a) Graphical distribution for 2011–2016

Year	Articles	Pages
2010	0	0
2011	37	806
2012	42	913
2013	84	2059
2014	62	1464
2015	62	1636
2016	89	1910

(b) Detailed numbers for 2011–2016

Fig. 13.23
Statistics about Dagstuhl Reports published between 2011 to 2016.

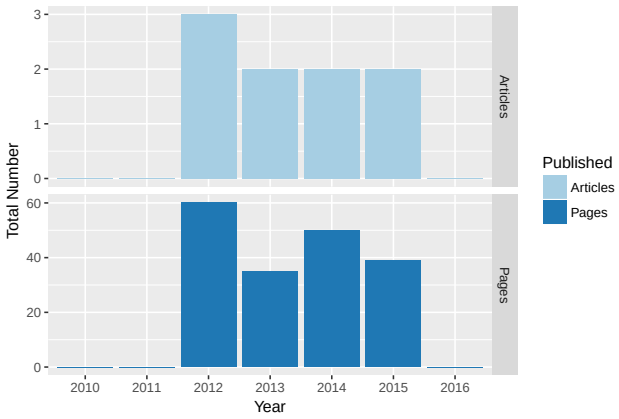

(a) Graphical distribution for 2012–2016

Year	Articles	Pages
2010	0	0
2011	0	0
2012	3	60
2013	2	35
2014	2	50
2015	2	39
2016	0	0

(b) Detailed numbers for 2012–2016

Fig. 13.24
Statistics about Dagstuhl Manifestos published between 2012 to 2016.

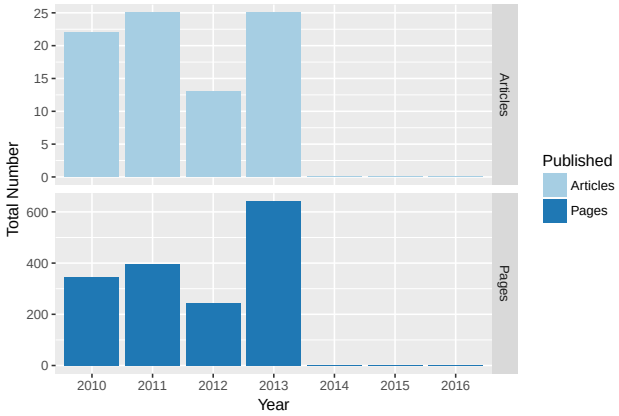

(a) Graphical distribution for 2010–2016

Year	Volumes	Articles	Pages
2010	1	22	345
2011	1	25	395
2012	1	13	246
2013	3	25	641
2014	0	0	0
2015	0	0	0
2016	0	0	0

(b) Detailed numbers for 2010–2016

Fig. 13.25
Statistics about Dagstuhl Follow-Ups volumes published between 2010 to 2016.

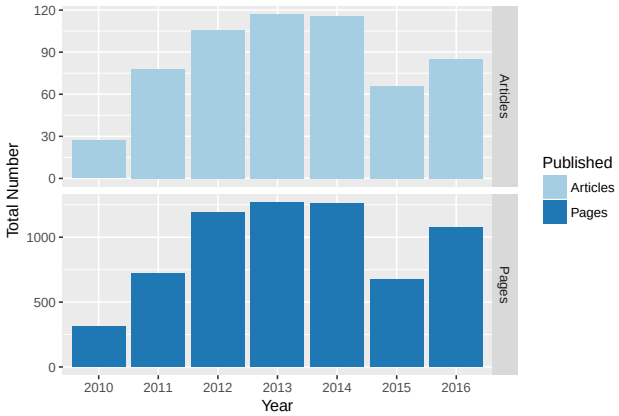

(a) Graphical distribution for 2010–2016

Year	Volumes	Articles	Pages
2010	2	27	315
2011	5	78	717
2012	8	106	1192
2013	7	117	1265
2014	8	116	1264
2015	6	66	674
2016	6	85	1078

(b) Detailed numbers for 2010–2016

Fig. 13.26
Statistics about OASIcs volumes published between 2010 to 2016.

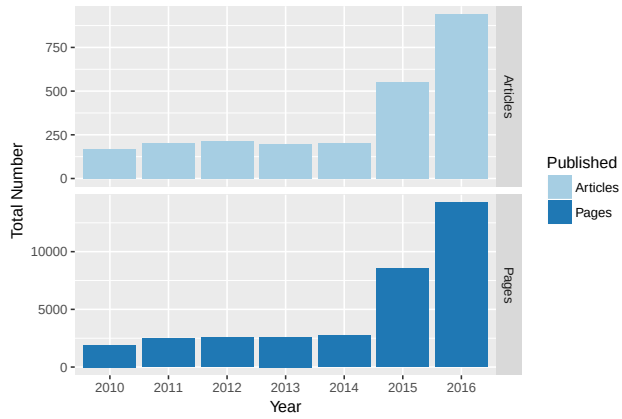

(a) Graphical distribution for 2010–2016

Year	Volumes	Articles	Pages
2010	4	167	1907
2011	5	205	2439
2012	5	215	2591
2013	6	195	2607
2014	5	204	2752
2015	16	553	8565
2016	19	939	14222

(b) Detailed numbers for 2010–2016

Fig. 13.27
Statistics about LIPIcs volumes published between 2010 to 2016.

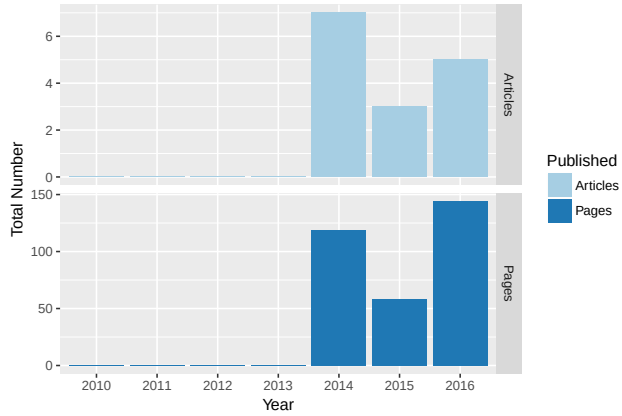

(a) Graphical distribution for 2014–2016

Year	Articles	Pages
2010	0	0
2011	0	0
2012	0	0
2013	0	0
2014	7	119
2015	3	58
2016	5	144

(b) Detailed numbers for 2014–2016

Fig. 13.28
Statistics about LITES articles published between 2014 to 2016.

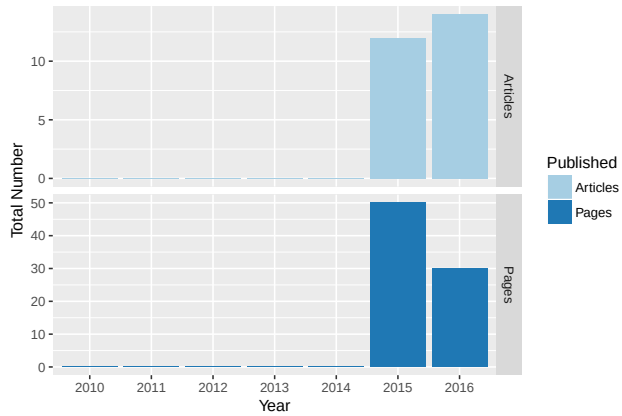

(a) Graphical distribution for 2011–2016

Year	Articles	Pages
2010	0	0
2011	0	0
2012	0	0
2013	0	0
2014	0	0
2015	12	50
2016	14	30

(b) Detailed numbers for 2011–2016

Fig. 13.29
Statistics about DARTS artifacts published between 2011 to 2016.

14 Veranstaltungen 2016
Schedule of Events 2016

Dagstuhl-Seminare 14.1 *Dagstuhl Seminars*

16011 – Evolution and Computing
Nick Barton (IST Austria – Klosterneuburg, AT), Bernard Chazelle (Princeton University, US), Per Kristian Lehre (University of Nottingham, GB), Nisheeth K. Vishnoi (EPFL – Lausanne, CH)
January 4–8, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16011

16012 – Global Measurements: Practice and Experience
Vaibhav Bajpai (Jacobs University Bremen, DE), Arthur W. Berger (Akamai Technologies – Cambridge, US), Philip Eardley (BT Research – Ipswich, GB), Jörg Ott (TU München, DE), Jürgen Schönwälder (Jacobs University Bremen, DE)
January 4–7, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16012

16021 – Symmetric Cryptography
Frederik Armknecht (Universität Mannheim, DE), Tetsu Iwata (Nagoya University, JP), Kaisa Nyberg (Aalto University, FI), Bart Preneel (KU Leuven, BE)
January 10–15, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16021

16022 – Geometric and Graph-based Approaches to Collective Motion
Giuseppe F. Italiano (University of Rome "Tor Vergata", IT), Bettina Speckmann (TU Eindhoven, NL), Guy Theraulaz (Université Paul Sabatier – Toulouse, FR), Marc van Kreveld (Utrecht University, NL)
January 10–15, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16022

16031 – Well Quasi-Orders in Computer Science
Jean Goubault-Larrecq (ENS – Cachan, FR), Monika Seisenberger (Swansea University, GB), Victor Selivanov (A. P. Ershov Institute – Novosibirsk, RU), Andreas Weiermann (Ghent University, BE)
January 17–22, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16031

16032 – Privacy and Security in Smart Energy Grids
George Danezis (University College London, GB), Stefan Katzenbeisser (TU Darmstadt, DE), Christiane Peters (IBM Belgium, BE), Bart Preneel (KU Leuven, BE)
January 17–20, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16032

16041 – Reproducibility of Data-Oriented Experiments in e-Science
Juliana Freire (New York University, US), Norbert Fuhr (Universität Duisburg-Essen, DE), Andreas Rauber (TU Wien, AT)
January 24–29, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16041

16042 – Eyewear Computing – Augmenting the Human with Head-mounted Wearable Assistants
Andreas Bulling (MPI für Informatik – Saarbrücken, DE), Ozan Cakmakci (Google Inc. – Mountain View, US), Kai Kunze (Keio University – Yokohama, JP), James M. Rehg (Georgia Institute of Technology – Atlanta, US)
January 24–29, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16042

16051 – Modern Cryptography and Security: An Inter-Community Dialogue
Kristin Lauter (Microsoft Research – Redmond, US), Ahmad-Reza Sadeghi (TU Darmstadt, DE), Radu Sion (National Security Institute – Stony Brook, US), Nigel P. Smart (University of Bristol, GB)
January 31 to February 5, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16051

16052 – Dark Silicon: From Embedded to HPC Systems
Hans Michael Gerndt (TU München, DE), Michael Glaß (Universität Erlangen-Nürnberg, DE), Sri Parameswaran (UNSW – Sydney, AU), Barry L. Rountree (LLNL – Livermore, US)
January 31 to February 3, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16052

16061 – Data-Driven Storytelling
Sheelagh Carpendale (University of Calgary, CA), Nicholas Diakopoulos (University of Maryland – College Park, US), Nathalie Henry Riche (Microsoft Research – Redmond, US), Christophe Hurter (ENAC – Toulouse, FR)
February 7–12, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16061

16062 – Modeling and Analysis of Semiconductor Supply Chains
Chen-Fu Chien (National Tsing Hua University, TW), Hans Ehm (Infineon Technologies – München, DE), John Fowler (Arizona State University – Tempe, US), Lars Mönch (FernUniversität in Hagen, DE)
February 7–12, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16062

16071 – Pattern Avoidance and Genome Sorting
Michael Albert (University of Otago, NZ), Miklós Bóna (University of Florida – Gainesville, US), István Miklós (Alfréd Rényi Institute of Mathematics – Budapest, HU), Einar Steingrimsson (University of Strathclyde – Glasgow, GB)
February 14–19, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16071

16072 – Assessing Learning In Introductory Computer Science
Michael E. Caspersen (Aarhus University, DK), Kathi Fisler (Worcester Polytechnic Institute, US), Jan Vahrenhold (Universität Münster, DE)
February 14–19, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16072

16081 – Scheduling
Nikhil Bansal (TU Eindhoven, NL), Nicole Megow (TU München, DE), Clifford Stein (Columbia University, US)
February 21–26, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16081

16091 – Computational Challenges in Cooperative Intelligent Urban Transport
Caitlin Doyle Cottrill (University of Aberdeen, GB), Jan Fabian Ehmke (FU Berlin, DE), Franziska Klügl (University of Örebro, SE), Sabine Timpf (Universität Augsburg, DE)
February 28 to March 4, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16091

16092 – Computational Music Structure Analysis
Juan Pablo Bello (New York University, US), Elaine Chew (Queen Mary University of London, GB), Meinard Müller (Universität Erlangen-Nürnberg, DE)
February 28 to March 4, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16092

16101 – Data Structures and Advanced Models of Computation on Big Data
Alejandro Lopez-Ortiz (University of Waterloo, CA), Ulrich Carsten Meyer (Goethe-Universität – Frankfurt a. M., DE), Markus E. Nebel (TU Kaiserslautern, DE), Robert Sedgewick (Princeton University, US)
March 6–11, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16101

16111 – Rethinking Experimental Methods in Computing
Daniel Delling (Apple Inc. – Cupertino, US), Camil Demetrescu (Sapienza University of Rome, IT), David S. Johnson (verstorben 03/2016, US), Jan Vitek (Northeastern University – Boston, US)
March 13–18, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16111

16112 – From Theory to Practice of Algebraic Effects and Handlers
Andrej Bauer (University of Ljubljana, SI), Martin Hofmann (LMU München, DE), Matija Pretnar (University of Ljubljana, SI), Jeremy Yallop (University of Cambridge, GB)
March 13–18, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16112

16131 – Language Based Verification Tools for Functional Programs
Marco Gaboardi (University at Buffalo, US), Suresh Jagannathan (Purdue University – West Lafayette, US), Ranjit Jhala (University of California – San Diego, US), Stephanie Weirich (University of Pennsylvania – Philadelphia, US)
March 28 to April 1, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16131

16141 – Analysis, Interpretation and Benefit of User-Generated Data: Computer Science Meets Communication Studies
Thorsten Quandt (Universität Münster, DE), German Shegalov (Twitter – San Francisco, US), Helle Sjøvaag (University of Bergen, NO), Gottfried Vossen (Universität Münster, DE)
April 3–8, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16141

16142 – Multidisciplinary Approaches to Multivalued Data: Modeling, Visualization, Analysis
Ingrid Hotz (Linköping University, SE), Evren Özarslan (Linköping University, SE), Thomas Schultz (Universität Bonn, DE)
April 3–8, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16142

16161 – Natural Language Argumentation: Mining, Processing, and Reasoning over Textual Arguments
Elena Cabrio (Laboratoire I3S – Sophia Antipolis, FR), Graeme Hirst (University of Toronto, CA), Serena Villata (Laboratoire I3S – Sophia Antipolis, FR), Adam Wyner (University of Aberdeen, GB)
April 17–22, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16161

16162 – Managing Technical Debt in Software Engineering
Paris Avgeriou (University of Groningen, NL), Philippe Kruchten (University of British Columbia – Vancouver, CA), Ipek Ozkaya (Carnegie Mellon University – Pittsburgh, US), Carolyn Seaman (University of Maryland, Baltimore County, US)
April 17–22, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16162

16171 – Algorithmic Methods for Optimization in Public Transport
Leo G. Kroon (Erasmus University – Rotterdam, NL), Anita Schöbel (Universität Göttingen, DE), Dorothea Wagner (KIT – Karlsruher Institut für Technologie, DE)
April 24–29, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16171

16172 – Machine Learning for Dynamic Software Analysis: Potentials and Limits
Amel Bennaceur (The Open University – Milton Keynes, GB), Dimitra Giannakopoulou (NASA – Moffett Field, US), Reiner Hähnle (TU Darmstadt, DE), Karl Meinke (KTH Royal Institute of Technology – Stockholm, SE)
April 24–27, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16172

16191 – Fresh Approaches to Business Process Modeling
Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US), Agnes Koschmider (KIT – Karlsruher Institut für Technologie, DE), Hajo A. Reijers (Free University Amsterdam, NL), William Wong (Middlesex University, GB)
May 8–13, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16191

16192 – Supporting Organizational Efficiency and Agility: Models, Languages and Software Systems
Tony Clark (Sheffield Hallam University, GB), Ulrich Frank (Universität Duisburg-Essen, DE), Vinay Kulkarni (Tata Consultancy Services – Pune, IN)
May 8–13, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16192

16201 – Synergies among Testing, Verification, and Repair for Concurrent Programs
Julian Dolby (IBM TJ Watson Research Center – Yorktown Heights, US), Orna Grumberg (Technion – Haifa, IL), Peter Müller (ETH Zürich, CH), Omer Tripp (IBM TJ Watson Research Center – Yorktown Heights, US)
May 16–20, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16201

16202 – Hardware Security
Osnat Keren (Bar-Ilan University, IL), Ilia Polian (Universität Passau, DE), Mark M. Tehranipoor (University of Florida – Gainesville, US), Pim Tuyls (Intrinsic-ID – Mol, BE)
May 16–20, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16202

16221 – Algorithms for Optimization Problems in Planar Graphs
Jeff Erickson (University of Illinois – Urbana-Champaign, US), Philip N. Klein (Brown University – Providence, US), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU), Claire Mathieu (ENS – Paris, FR)
May 29 to June 3, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16221

16222 – Engineering Moral Agents – from Human Morality to Artificial Morality
Michael Fisher (University of Liverpool, GB), Christian List (London School of Economics, GB), Marija Slavkovik (University of Bergen, NO), Alan FT Winfield (University of the West of England – Bristol, GB)
May 29 to June 3, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16222

16231 – Immersive Analytics
Tim Dwyer (Monash University – Caulfield, AU), Nathalie Henry Riche (Microsoft Research – Redmond, US), Karsten Klein (Universität Konstanz, DE), Wolfgang Stuerzlinger (Simon Fraser University – Vancouver, CA), Bruce Thomas (University of South Australia – Mawson Lakes, AU)
June 5–10, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16231

16232 – Fair Division
Yonatan Aumann (Bar-Ilan University – Ramat Gan, IL), Steven J. Brams (New York University, US), Jérôme Lang (University Paris-Dauphine, FR), Ariel D. Procaccia (Carnegie Mellon University – Pittsburgh, US)
June 5–10, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16232

16241 – Graph Polynomials: Towards a Comparative Theory
Jo Ellis-Monaghan (Saint Michael's College – Colchester, US), Andrew Goodall (Charles University – Prague, CZ), Johann A. Makowsky (Technion – Haifa, IL), Iain Moffatt (Royal Holloway University of London, GB)
June 12–17, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16241

16251 – Information-centric Networking and Security
Edith Ngai (Uppsala University, SE), Börje Ohlman (Ericsson Research – Stockholm, SE), Gene Tsudik (University of California – Irvine, US), Ersin Uzun (Xerox PARC – Palo Alto, US)
June 19–22, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16251

16261 – Integration of Expert Knowledge for Interpretable Models in Biomedical Data Analysis
Gyan Bhanot (Rutgers University – Piscataway, US), Michael Biehl (University of Groningen, NL), Thomas Villmann (Hochschule Mittweida, DE), Dietlind Zühlke (Seven Principles AG – Köln, DE)
June 26 to July 1, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16261

16262 – Automotive User Interfaces in the Age of Automation
Susanne Boll (Universität Oldenburg, DE), Andrew Kun (University of New Hampshire – Durham, US), Andreas Riener (TH Ingolstadt, DE)
June 26 to July 1, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16262

16271 – Algorithmic Foundations of Programmable Matter
Sándor Fekete (TU Braunschweig, DE), Andréa Richa (Arizona State University – Tempe, US), Kay Römer (TU Graz, AT), Christian Scheideler (Universität Paderborn, DE)
July 3–8, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16271

16281 – Network Latency Control in Data Centres
Mohammad Alizadeh Attar (MIT – Cambridge, US), Jon Crowcroft (University of Cambridge, GB), Lars Eggert (NetApp Deutschland GmbH – Kirchheim, DE), Klaus Wehrle (RWTH Aachen, DE)
July 10–13, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16281

16282 – Topological Methods in Distributed Computing
Dmitry Feichtner-Kozlov (Universität Bremen, DE), Damien Imbs (Universität Bremen, DE)
July 10–15, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16282

16291 – Data, Responsibly
Serge Abiteboul (ENS – Cachan, FR), Gerome Miklau (University of Massachusetts – Amherst, US), Julia Stoyanovich (Drexel Univ. – Philadelphia, US), Gerhard Weikum (MPI für Informatik – Saarbrücken, DE)
July 17–22, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16291

16321 – Coding Theory in the Time of Big Data
Martin Bossert (Universität Ulm, DE), Eimear Byrne (University College Dublin, IE), Emina Soljanin (Rutgers University – Piscataway, US)
August 7–12, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16321

16341 – Integrating Process-Oriented and Event-Based Systems
David Eyers (University of Otago, NZ), Avigdor Gal (Technion – Haifa, IL), Hans-Arno Jacobsen (TU München, DE), Matthias Weidlich (HU Berlin, DE)
August 21–26, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16341

16342 – Foundations of Secure Scaling
Lejla Batina (Radboud University Nijmegen, NL), Swarup Bhunia (University of Florida – Gainesville, US), Patrick Schaumont (Virginia Polytechnic Institute – Blacksburg, US), Jean-Pierre Seifert (TU Berlin, DE)
August 21–26, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16342

16351 – Next Generation Sequencing – Algorithms, and Software For Biomedical Applications
Gene Myers (MPI – Dresden, DE), Mihai Pop (University of Maryland – College Park, US), Knut Reinert (FU Berlin, DE), Tandy Warnow (University of Illinois – Urbana-Champaign, US)
August 28 to September 2, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16351

16361 – Network Attack Detection and Defense – Security Challenges and Opportunities of Software-Defined Networking
Marc C. Dacier (QCRI – Doha, QA), Sven Dietrich (City University of New York, US), Frank Kargl (Universität Ulm, DE), Hartmut König (BTU Cottbus, DE)
September 4–9, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16361

16362 – Robustness in Cyber-Physical Systems
Martin Fränzle (Universität Oldenburg, DE), James Kapinski (Toyota Technical Center – Gardena, US), Pavithra Prabhakar (Kansas State University – Manhattan, US)
September 4–9, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16362

16371 – Public-Key Cryptography
Marc Fischlin (TU Darmstadt, DE), Alexander May (Ruhr-Universität Bochum, DE), David Pointcheval (ENS – Paris, FR), Tal Rabin (IBM TJ Watson Research Center – Yorktown Heights, US)
September 11–16, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16371

16372 – Uncertainty Quantification and High Performance Computing
Vincent Heuveline (HITS & Universität Heidelberg), Michael Schick (Robert Bosch GmbH – Stuttgart, DE), Clayton Webster (Oak Ridge National Laboratory, US), Peter Zaspel (HITS & Universität Heidelberg)
September 11–16, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16372

16381 – SAT and Interactions
Olaf Beyersdorff (University of Leeds, GB), Nadia Creignou (Aix-Marseille University, FR), Uwe Egly (TU Wien, AT), Heribert Vollmer (Leibniz Universität Hannover, DE)
September 18–23, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16381

16382 – Foundations of Unsupervised Learning
Maria-Florina Balcan (Carnegie Mellon University – Pittsburgh, US), Shai Ben-David (University of Waterloo, CA), Ruth Urner (MPI für Intelligente Systeme – Tübingen, DE), Ulrike von Luxburg (Universität Tübingen, DE)
September 18–23, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16382

16402 – Programming Language Techniques for Incremental and Reactive Computing
Camil Demetrescu (Sapienza University of Rome, IT), Sebastian Erdweg (TU Delft, NL), Matthew A. Hammer (University of Colorado – Boulder, US), Shriram Krishnamurthi (Brown University – Providence, US)
October 3–7, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16402

16411 – Algebraic Methods in Computational Complexity
Valentine Kabanets (Simon Fraser University – Burnaby, CA), Thomas Thierauf (Hochschule Aalen, DE), Jacobo Torán (Universität Ulm, DE), Christopher Umans (CalTech – Pasadena, US)
October 9–14, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16411

16412 – Automated Algorithm Selection and Configuration
Holger H. Hoos (University of British Columbia – Vancouver, CA), Frank Neumann (University of Adelaide, AU), Heike Trautmann (Universität Münster, DE)
October 9–14, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16412

16421 – Universality of Proofs
Gilles Dowek (INRIA & ENS Cachan, FR), Catherine Dubois (ENSIIE – Evry, FR), Brigitte Pientka (McGill University – Montreal, CA), Florian Rabe (Jacobs University Bremen, DE)
October 16–21, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16421

16431 – Computation over Compressed Structured Data
Philip Bille (Technical University of Denmark – Lyngby, DK), Markus Lohrey (Universität Siegen, DE), Sebastian Maneth (University of Edinburgh, GB), Gonzalo Navarro (University of Chile – Santiago de Chile, CL)
October 23–28, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16431

16441 – Adaptive Isolation for Predictability and Security
Tulika Mitra (National University of Singapore, SG), Jürgen Teich (Universität Erlangen-Nürnberg, DE), Lothar Thiele (ETH Zürich, CH), Ingrid Verbauwhede (KU Leuven, BE)
October 30 to November 4, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16441

16442 – Vocal Interactivity in-and-between Humans, Animals and Robots (VIHAR)
Ricard Marxer (University of Sheffield, GB), Roger K. Moore (University of Sheffield, GB), Serge Thill (University of Skövde, SE), Clémentine Vignal (Université Jean Monnet – Saint-Étienne, FR)
October 30 to November 4, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16442

16451 – Structure and Hardness in P
Moshe Lewenstein (Bar-Ilan University – Ramat Gan, IL), Seth Pettie (University of Michigan – Ann Arbor, US), Virginia Vassilevska Williams (Stanford University, US)
November 6–11, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16451

16452 – Beyond-Planar Graphs: Algorithmics and Combinatorics
Seokhee Hong (The University of Sydney, AU), Michael Kaufmann (Universität Tübingen, DE), Stephen G. Kobourov (University of Arizona – Tucson, US), Janos Pach (EPFL – Lausanne, CH)
November 6–11, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16452

16461 – Assessing ICT Security Risks in Socio-Technical Systems
Tyler W. Moore (University of Tulsa, US), Christian W. Probst (Technical University of Denmark – Lyngby, DK), Kai Rannenberg (Goethe-Universität Frankfurt am Main, DE), Michel van Eeten (TU Delft, NL)
November 13–18, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16461

16462 – Inpainting-Based Image Compression
Christine Guillemot (INRIA – Rennes, FR), Gerlind Plonka-Hoch (Universität Göttingen, DE), Thomas Pock (TU Graz, AT), Joachim Weickert (Universität des Saarlandes, DE)
November 13–18, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16462

16471 – Concurrency with Weak Memory Models: Semantics, Languages, Compilation, Verification, Static Analysis, and Synthesis
Jade Alglave (University College London, GB), Patrick Cousot (New York University, US), Caterina Urban (ETH Zürich, CH)
November 20–25, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16471

16481 – New Directions for Learning with Kernels and Gaussian Processes
Arthur Gretton (University College London, GB), Philipp Hennig (MPI für Intelligente Systeme – Tübingen, DE), Carl Edward Rasmussen (University of Cambridge, GB), Bernhard Schölkopf (MPI für Intelligente Systeme – Tübingen, DE)
November 27 to December 2, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16481

16482 – Algorithms and Effectivity in Tropical Mathematics and Beyond
Stéphane Gaubert (INRIA Saclay – Île-de-France, FR), Dima Grigoriev (Lille I University, FR), Michael Joswig (TU Berlin, DE), Thorsten Theobald (Goethe-Universität Frankfurt am Main, DE)
November 27 to December 2, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16482

16491 – Symbolic-Numeric Methods for Reliable and Trustworthy Problem Solving in Cyber-Physical Domains
Sergiy Bogomolov (Australian National University – Canberra, AU), Martin Fränzle (Universität Oldenburg, DE), Kyoko Makino (Michigan State University – East Lansing, US), Nacim Ramdani (University of Orléans, FR)
December 4–9, 2016 | Dagstuhl Seminar | http://www.dagstuhl.de/16491

Dagstuhl-Perspektiven-Workshops 14.2 *Dagstuhl Perspectives Workshops*

16151 – Foundations of Data Management
Marcelo Arenas (Pontificia Universidad Catolica de Chile, CL), Richard Hull (IBM TJ Watson Research Center – Yorktown Heights, US), Wim Martens (Universität Bayreuth, DE), Tova Milo (Tel Aviv University, IL), Thomas Schwentick (TU Dortmund, DE)
April 10–15, 2016 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16151

16152 – Tensor Computing for Internet of Things
Evrim Acar (University of Copenhagen, DK), Animashree Anandkumar (University of California – Irvine, US), Lenore Mullin (University of Albany – SUNY, US), Sebnem Rusitschka (Siemens AG – München, DE), Volker Tresp (Siemens AG – München, DE)
April 10–13, 2016 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16152

16252 – Engineering Academic Software
Carole Goble (University of Manchester, GB), James Howison (University of Texas – Austin, US), Claude Kirchner (INRIA – Le Chesnay, FR), Oscar M. Nierstrasz (Universität Bern, CH), Jurgen J. Vinju (CWI – Amsterdam, NL)
June 19–24, 2016 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16252

16472 – QoE Vadis?
Markus Fiedler (Blekinge Institute of Technology – Karlskrona, SE), Sebastian Möller (TU Berlin, DE), Peter Reichl (Universität Wien, AT), Min Xie (Telenor Research – Trondheim, NO)
November 20–25, 2016 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/16472

GI-Dagstuhl-Seminare 14.3 *GI-Dagstuhl Seminars*

16082 – Informatik@Schule 2016 – Das Verhältnis von informatischer Bildung und „Digitaler Bildung“
Torsten Brinda (Universität Duisburg-Essen, DE), Ira Diethelm (Universität Oldenburg, DE), Rainer Gemulla (Universität Mannheim, DE), Ralf Romeike (Universität Erlangen-Nürnberg, DE), Johannes Schöning (Hasselt University – Diepenbeek, BE), Carsten Schulte (FU Berlin, DE)
February 21–24, 2016 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/16082

16353 – Aware Machine-to-Machine Communication
Mayutan Arumaithurai (Universität Göttingen, DE), Stephan Sigg (Aalto University, FI), Xiaoyan Wang (Ibaraki University, JP)
August 28 to September 2, 2016 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/16353

16394 – Software Performance Engineering in the DevOps World
Pooyan Jamshidi (Imperial College London, GB), Philipp Leitner (Universität Zürich, CH), André van Hoorn (Universität Stuttgart, DE), Ingo Weber (Data61 / NICTA – Sydney, AU)
September 25–30, 2016 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/16394

Lehrveranstaltungen 14.4 *Educational Events*

16103 – Spring School "Models, Systems, and Algorithms for Role-based Business Intelligence Applications"
Wolfgang Lehner (TU Dresden, DE), Esteban Zimanyi (Free University of Brussels, BE)
March 6–9, 2016 | Educational Event | http://www.dagstuhl.de/16103

16184 – Workshop Wissenschaftsjournalismus
Roswitha Bardohl (Schloss Dagstuhl – Saarbrücken, DE), Gordon Bolduan (Universität des Saarlandes, DE), Tim Schröder (Oldenburg, DE)
May 1–4, 2016 | Educational Event | http://www.dagstuhl.de/16184

16323 – Sommerschule "Data Management Techniques"
Goetz Graefe (Google – Madison, US)
August 7–12, 2016 | Educational Event | http://www.dagstuhl.de/16323

16393 – de.NBI Sommer Schule 2016 – From Big Data to Big Insights: Computational Methods for the Analysis and Interpretation of Mass-Spectrometric High-Throughput Data
Stefan Albaum (Universität Bielefeld, DE), Martin Eisenacher (Ruhr-Universität Bochum, DE), Oliver Kohlbacher (Universität Tübingen, DE), Knut Reinert (FU Berlin, DE)
September 25–30, 2016 | Educational Event | http://www.dagstuhl.de/16393

16403 – Autumn School 2016 for Information Retrieval and Information Foraging
Ingo Frommholz (University of Bedfordshire – Luton, GB), Norbert Fuhr (Universität Duisburg-Essen, DE), Thomas Mandl (Universität Hildesheim, DE)
October 3–7, 2016 | Educational Event | http://www.dagstuhl.de/16403

16503 – Lehrerfortbildung in Informatik
Manuel Garcia Mateos (LPM Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 14–16, 2016 | Educational Event | http://www.dagstuhl.de/16503

Forschungsgruppentreffen 14.5 *Research Group Meetings*

16013 – Informationsmanagement für öffentliche Mobilitätsanbieter
Karl-Heinz Krempels (RWTH Aachen, DE), Berthold Radermacher (VDV – Köln, DE)
January 7–8, 2016 | Research Group Meeting | http://www.dagstuhl.de/16013

16053 – Lehrstuhltreffen AG Zeller
Andreas Zeller (Universität des Saarlandes, DE)
February 3–5, 2016 | Research Group Meeting | http://www.dagstuhl.de/16053

16084 – Quantum Information: Theory & Implementation
Christoph Becher (Universität des Saarlandes – Saarbrücken, DE), Jürgen Eschner (Universität des Saarlandes – Saarbrücken, DE), Jörg Hettel (HS Kaiserslautern – Zweibrücken, DE), Hans-Jürgen Steffens (HS Kaiserslautern – Zweibrücken, DE)
February 24–26, 2016 | Research Group Meeting | http://www.dagstuhl.de/16084

16089 – Forschungsaufenthalt
Klaus Keimel (TU Darmstadt, DE)
February 21–26, 2016 | Research Group Meeting | http://www.dagstuhl.de/16089

16123 – GIBU 2016: GI-Beirat der Universitätsprofessoren
Gregor Snelting (KIT – Karlsruher Institut für Technologie, DE)
March 20–22, 2016 | Research Group Meeting | http://www.dagstuhl.de/16123

16124 – Klausurtagung des Forschungsprojekts "diagnoseIT – Expert-Guided Automatic Diagnosis of Performance Problems in Enterprise Applications"
Stefan Siegl (NovaTec Holding GmbH – Leinfelden-Echterdingen, DE), André van Hoorn (Universität Stuttgart, DE)
March 20–23, 2016 | Research Group Meeting | http://www.dagstuhl.de/16124

16125 – Arbeitstreffen Text-Technology Lab
Rüdiger Gleim (Goethe-Universität – Frankfurt a. M., DE), Alexander Mehler (Goethe-Universität – Frankfurt a. M., DE)
March 21–23, 2016 | Research Group Meeting | http://www.dagstuhl.de/16125

16133 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)
Michaela Huhn (Ostfalia Hochschule – Wolfenbüttel, DE), Matthias Riebisch (Universität Hamburg, DE), Bernhard Schätz (fortiss GmbH – München, DE)
March 30 to April 1, 2016 | Research Group Meeting | http://www.dagstuhl.de/16133

16153 – Klausurtagung Telematik Karlsruhe
Robert Bauer (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)
April 13–15, 2016 | Research Group Meeting | http://www.dagstuhl.de/16153

16154 – Klausurtagung AG Schneider
Klaus Schneider (TU Kaiserslautern, DE)
April 13–15, 2016 | Research Group Meeting | http://www.dagstuhl.de/16154

16183 – Kolloquium zum GI Dissertationspreis 2015
Steffen Hölldobler (TU Dresden, DE)
May 1–4, 2016 | Research Group Meeting | http://www.dagstuhl.de/16183

16185 – Lehrstuhltreffen "Embedded Intelligence"
Bernhard Sick (Universität Kassel, DE)
May 1–4, 2016 | Research Group Meeting | http://www.dagstuhl.de/16185

16213 – Gemeinsamer Workshop der Graduiertenkollegs: GRK 1780 Crossworlds der TU Chemnitz & GRK 1907 RoSI der TU Dresden
Kai Herrmann (TU Dresden, DE), Wolfgang Lehner (TU Dresden, DE), Hannes Voigt (TU Dresden, DE), Martin Weißbach (TU Dresden, DE)
May 22–25, 2016 | Research Group Meeting | http://www.dagstuhl.de/16213

16239 – Forschungsaufenthalt
Jon McCormack (Monash University – Caulfield, AU)
June 11–17, 2016 | Research Group Meeting | http://www.dagstuhl.de/16239

16244 – Workshop Buchprojekt "Corporate Semantic Web"
Thomas Hoppe (Datenlabor Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE), Anatol Reibold (OntoPort UG – Darmstadt, DE)
June 12–15, 2016 | Research Group Meeting | http://www.dagstuhl.de/16244

16253 – Clusterseminar "Intelligente Systeme zur Entscheidungsunterstützung"
Lars Mönch (FernUniversität in Hagen, DE)
June 22–24, 2016 | Research Group Meeting | http://www.dagstuhl.de/16253

16259 – Forschungsaufenthalt
Tobias Rawald (GFZ – Potsdam, DE)
June 19 to July 1, 2016 | Research Group Meeting | http://www.dagstuhl.de/16259

16273 – Evaluierungskommission WGL
Roswitha Bardohl (Schloss Dagstuhl – Saarbrücken, DE), Raimund Seidel (Universität des Saarlandes, DE), Michael Wagner (Schloss Dagstuhl – Trier, DE)
July 4–6, 2016 | Research Group Meeting | http://www.dagstuhl.de/16273

16274 – Deutsch-Pakistanischer Workshop
Karsten Berns (TU Kaiserslautern, DE)
July 7–8, 2016 | Research Group Meeting | http://www.dagstuhl.de/16274

16283 – Retreat SFB 1102: Information Density and Linguistic Encoding
Elke Teich (Universität des Saarlandes, DE)
July 13–15, 2016 | Research Group Meeting | http://www.dagstuhl.de/16283

16293 – Workshop zur wissenschaftlichen Weiterbildung FAU, Lehrstuhl Informatik 1
Felix Freiling (Universität Erlangen-Nürnberg, DE)
July 17–20, 2016 | Research Group Meeting | http://www.dagstuhl.de/16293

16294 – Implicit and Explicit Guidance in CSCL Environments
Daniel Bodemer (Universität Duisburg-Essen, DE), Armin Weinberger (Universität des Saarlandes, DE)
July 20–22, 2016 | Research Group Meeting | http://www.dagstuhl.de/16294

16333 – Arbeitstreffen AG Prof. Bernd Becker – "Challenges in Computer Aided Engineering"
Bernd Becker (Universität Freiburg, DE)
August 15–18, 2016 | Research Group Meeting | http://www.dagstuhl.de/16333

16334 – Crowdsourcing Research – Transcending Disciplinary Boundaries
Michele Catasta (EPFL – Lausanne, CH), Gianluca Demartini (University of Sheffield, GB), Ujwal Gadiraju (Leibniz Universität Hannover, DE), Cristina Sarasua (Universität Koblenz-Landau, DE)
August 15–18, 2016 | Research Group Meeting | http://www.dagstuhl.de/16334

16335 – Lehrstuhltreffen Rechtsinformatik
Christoph Sorge (Universität des Saarlandes, DE)
August 17–18, 2016 | Research Group Meeting | http://www.dagstuhl.de/16335

16423 – DDI Moving Forward: Facilitating Interoperability and Collaboration with Other Metadata Standards
Arofan Gregory (Open Data Foundation – Tucson, US), Jared Lyle (University of Michigan – Ann Arbor, US), Steven McEachern (Australian National University – Canberra, AU), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE), Benjamin Zapilko (GESIS – Köln, DE)
October 16–21, 2016 | Research Group Meeting | http://www.dagstuhl.de/16423

16433 – DDI Moving Forward: Improvement and Refinement of Selected Areas
Michelle Edwards (Cornell University, US), Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 23–28, 2016 | Research Group Meeting | http://www.dagstuhl.de/16433

16493 – Klausurtagung "LST Schmeck"
Birger Becker (FZI – Karlsruhe, DE), Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)
December 7–9, 2016 | Research Group Meeting | http://www.dagstuhl.de/16493

16494 – Klausurtagung des Forschungsprojekts "diagnoseIT – Expert-Guided Automatic Diagnosis of Performance Problems in Enterprise Applications"
Stefan Siegl (NovaTec Holding GmbH – Leinfelden-Echterdingen, DE), André van Hoorn (Universität Stuttgart, DE)
December 4–7, 2016 | Research Group Meeting | http://www.dagstuhl.de/16494

16495 – Workshop "SimPhon.Net"
Ingmar Steiner (Universität des Saarlandes, DE)
December 4–7, 2016 | Research Group Meeting | http://www.dagstuhl.de/16495

16504 – Secan Lab Meeting
Thomas Engel (University of Luxembourg, LU)
December 12–13, 2016 | Research Group Meeting | http://www.dagstuhl.de/16504

16505 – Redaktions-Sitzung "Informatik in der Schule"
Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 13–14, 2016 | Research Group Meeting | http://www.dagstuhl.de/16505

16513 – FORSEC-Workshop: Security in highly connected IT Systems
Felix Freiling (Universität Erlangen-Nürnberg, DE), Günther Pernul (Universität Regensburg, DE), Guido Schryen (Universität Regensburg, DE)
December 18–21, 2016 | Research Group Meeting | http://www.dagstuhl.de/16513

16514 – Erneuerbare Mobilität
Karl-Heinz Krempels (Fraunhofer Institut FIT – St. Augustin, DE)
December 19–21, 2016 | Research Group Meeting | http://www.dagstuhl.de/16514

16519 – Forschungsaufenthalt
Bernd Becker (Universität Freiburg, DE)
December 19–20, 2016 | Research Group Meeting | http://www.dagstuhl.de/16519

Mitglied der Leibniz-Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2016 | ISSN 2199-1995
http://www.dagstuhl.de