SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2017

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2017

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, Michael Gerke, Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Heike Clemens, Dr. Andreas Dolzmann, Jutka Gasiorowski, Michael Gerke, Shida Kunz, Ph. D., Heike Meißner, Petra Meyer, Thomas Schillo, Raimund Seidel, Ph. D., Dr. Michael Wagner und die Autoren der Zusammenfassungen in Kapitel 6
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
Druck	Kern GmbH, Bexbach
Druckauflage	170
ISSN	2199-1995
Online-Version	https://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	April 2018
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

2017 war für Schloss Dagstuhl – Leibniz-Zentrum für Informatik ein sehr erfolgreiches Jahr.

Der Höhepunkt war sicherlich die im Mai gefällte Entscheidung der GWK (Gemeinsame Wissenschaftskonferenz), aufgrund der positiven Evaluierung die Finanzierung des LZI für weitere sieben Jahre zu gewährleisten.

Mit 77 Dagstuhl Seminaren und Perspektiven Workshops, die von beinahe 2 600 Wissenschaftlern besucht wurden, wurde 2017 ein neuer Rekord aufgestellt. Fast die Hälfte dieser Wissenschaftler war zum ersten Mal bei uns, mehr als ein Drittel waren Nachwuchswissenschaftler, mehr als drei Viertel kam aus dem Ausland und fast ein Fünftel war weiblich.

Auch bei der *dblp* Bibliographiedatenbank wurden 2017 neue Rekorde aufgestellt, indem *dblp* um mehr als 400 000 Einträge auf über vier Millionen Publikationsdatensätze wuchs. Die Nutzung der Datenbank nahm ebenfalls zu und erreichte durchschnittlich sechs Seitenaufrufe pro Sekunde. Unterdessen wurden auch Fortschritte bei der Sicherung zusätzlicher Grundfnanzierung zwecks völliger Integration von *dblp* ins LZI und weiterem Ausbau gemacht. Dabei wurden alle formellen Hürden problemlos genommen, und wir sind zuversichtlich, dass die GWK bald diese Ausweitung der Finanzierung genehmigen wird.

Unsere Abteilung Open-Access-Verlagsdienstleistungen hat 2017 ebenfalls neue Rekorde aufgestellt: Mehr als 1 100 Konferenzpapiere wurden während des Jahres in den Reihen *LIPIcs* und *OASIcs* veröffentlicht.

2017 was a very successful year for Schloss Dagstuhl – Leibniz Center for Informatics.

First and foremost, as a consequence of our positive evaluation, GWK (the Joint Science Conference of the German states) decided in May to continue funding of LZI for another seven year period.

We hosted a record high of 77 Dagstuhl Seminars and Perspective Workshops that were attended by almost 2,600 researchers of which almost half were first time visitors to Dagstuhl, more than one third were junior researchers, more than three quarters came from abroad, and almost a fifth were women.

The *dblp* computer science bibliography also set new records in 2017 by growing by more than 400,000 entries to reach a total of more than four million publication records. The usage of the database has also grown and has reached an average of six page requests per second. At the same time the formal process for receiving additional base funding in order to incorporate *dblp* completely into LZI and expand it has proceeded smoothly, taking all formal hurdles without problems. We are optimistic that GKW will soon give its official nod to this funding expansion.

Also our open access publishing division set new records in 2017 with more than 1,100 papers published during the year in the *LIPIcs* and *OASIcs* series.

Dagstuhl saw some changes during the year, e.g. the construction of a new lecture hall and of a new parking lot. More importantly, it saw changes in personnel: Marc

Es gab während des Jahres einige Veränderungen in Dagstuhl, z. B. die Fertigstellung eines neuen Vortragssaales und eines neuen Parkplatzes. Darüber hinaus gab es personelle Veränderungen im Stab: Marc Herbstritt verlies uns unglücklicherweise zum Jahresende und Rosi Bardohl war das ganze Jahr über aus gesundheitlichen Gründen abwesend und ging zum Jahresende ausser Dienst.

Es ist uns jedoch gelungen, zwei neue Mitarbeiter für den Wissenschaftlichen Stab zu gewinnen, nämlich Michael Gerke und Michael Didas.

Weitere Details können Sie den folgenden Kapiteln entnehmen. Viel Spaß beim Lesen des Berichtes.

Raimund Seidel

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

Herbstritt unfortunately left us with the end of the year, Rosi Bardohl was on sick leave all year moving to an inactive status with the end of the year.

However, we were successful in hiring two new members to the scientific staff, namely Michael Gerke and Michael Didas.

You can find many more details in the following chapters. Enjoy the report.

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 Dezember 2017.
As of December 2017.

Fig. 1.1
Aerial photography of Schloss Dagstuhl.

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 vier 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.

Dagstuhl Publishing

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

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

Kapitel 4 stellt Dagstuhls Verlagswesen ausführlicher dar.

dblp computer science bibliography

Since 2011, Schloss Dagstuhl operates the dblp computer science bibliography in close cooporation with the University of Trier. Listing more than four million articles, dblp is the world's most comprehensive open data collection of computer science research articles. The goal of dblp is to support computer scientists in their daily work, for example when reviewing the literature of a given author or subject area, or when searching for online full-text versions of research articles. The dblp database is often considered to be the reference database for quality-assured and normalized bibliographic metadata in computer science. Additionally, dblp supports funding agencies and decision makers by providing and curating personalized author profiles. By operating dblp, Schloss Dagstuhl furthers its mission of promoting the identification, dissemination and implementation of new computer science developments at an internationally recognized level.

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

Dagstuhl Publishing

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

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

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

1.2 Neuigkeiten in 2017 / *News from 2017*

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. Bereits Mitte 2016 stand diese Evaluierung für Schloss Dagstuhl an. Am 16. Mai 2017 hat sich der Ausschuss der Gemeinsamen Wissenschaftskonferenz (GWK) auf Empfehlung des Leibniz-Senats abschließend für ein weitere Förderung ausgesprochen.

Evaluation

The Leibniz Association evaluates its institutions every seven years at the latest to ensure that they still fulfill the prerequisites for a joint funding by the German federal and state governments. Mid of 2016, Schloss Dagstuhl has been evaluated as well. On May 16, 2017, the Joint Science Conference (Gemeinsame Wissenschaftskonferenz, GWK), on the recommendation of the Leibniz Senate, decided in favor of continuing the funding of Schloss Dagstuhl.

Wie der Senat der Leibniz-Gemeinschaft in seiner Stellungnahme[2] festhält, unterstütze Schloss Dagstuhl mit seinem hoch angesehenen Veranstaltungsprogramm als soziale Forschungsinfrastruktur die Informatik auf einem weltweit herausragenden Niveau. Zusätzlich, so der Senat weiter, habe das Zentrum seit der letzten Evaluierung ausgesprochen erfolgreich mit neuen Angeboten auf wissenschaftliche Bedarfe reagiert. Sowohl die Open-Access-Angebote als auch die Beteiligung an der Bibliographie-Datenbank dblp seien von herausgehobener Bedeutung. Sie bildeten mittlerweile unverzichtbare und breit genutzte Dienstleistungen für die Informatikgemeinschaft.

Der Senat hebt zudem hervor, dass Schloss Dagstuhl in der Gestaltung und Steuerung der Diversität der Seminargruppen bereits gute Erfolge erzielt habe. Mit 13% bis 17% liege der Anteil von Wissenschaftlerinnen dagegen nach wie vor auf dem niedrigen Niveau der letzten Evaluierung. Hier müsse das international herausragende Zentrum eine Vorreiterrolle einnehmen.

Die Antragsauswahl und die Programmsteuerung verantworte ein sehr gut arbeitendes Gremium renommierter Informatikerinnen und Informatiker in enger Zusammenarbeit mit der engagierten LZI-Leitung. Hervorgehoben wird zudem die äußerst kompetente und serviceorientierte Unterstützung von Leitung, Gremien und Gästen von Schloss Dagstuhl durch die Beschäftigten des LZI.

As stated by the Senate of the Leibniz Association[2], Schloss Dagstuhl supports the computer science research community at a worldwide outstanding level with its highly esteemed seminar program. In addition, the Senate added, the Center has been very successful in responding to emerging scientific needs with new and improved services since the last evaluation. Both its open access publishing service and the dblp computer science bibliography are of particular importance. Both services have become indispensable and widely used assets of the international computer science community.

The Senate also emphasizes Schloss Dagstuhl's success in the support of the diversity of the seminar groups. At merely 13% to 17%, however, the proportion of women scientists remains at the low level already seen at the last evaluation. An internationally outstanding center such as Schloss Dagstuhl should play a pioneering role here.

In the view of the Senate, the selection of applications is managed by a very well functioning committee of renowned computer scientists in close cooperation with the committed LZI management. The extremely competent and service-oriented support of management, committees, and guests by the Dagstuhl staff is further highlighted.

Das Team

Ende 2017 beschäftigte Schloss Dagstuhl insgesamt 37 Vollzeitäquivalente bzw. 53 Angestellte. Schloss Dagstuhl beschäftigte in diesem Jahr eine Auszubildende in der Küche. Ein Mitarbeiter des dblp-Teams, Florian Reitz, verteidigte im Dezember erfolgreich seine Dissertation. Während Marc Herbstritt Schloss Dagstuhl verließ, verstärkten Michael Gerke und Michael Didas den Wissenschaftlichen Stab.

Nahezu alle Mitarbeiter von Schloss Dagstuhl wurden 2017 über den Kernhaushalt des Zentrums bezahlt. Eine Ausnahme bilden zwei Mitarbeiter im Rahmen laufenden Projektes „Skalierbare Autoren-Disambiguierung in Literaturdatenbanken", welche aus aus Fördermitteln des Leibniz-Wettbewerbes finanziert wurden. Zudem unterstützte das Heidelberger Institut für Theoretische Studien (HITS) mit einer großzügigen Spende von 42 000 € die Open-Access-Publikationsaktivitäten von Schloss Dagstuhl.

The Team

By the end of 2017, Schloss Dagstuhl had a total of 53 staff members corresponding to 37 full-time positions. Schloss Dagstuhl employed one trainee in its kitchen. In December, Florian Reitz of the dblp team successfully defended his dissertation. Marc Herbstritt left Schloss Dagstuhl, while Michael Gerke and Michael Didas joined the Scientific Staff.

Nearly all staff at Schloss Dagstuhl were funded from the center's core budget in 2017. An exception is the dblp team, where two dblp staff members were financed by a project grant of the Leibniz Competition (project "Scalable Author Disambiguation for Bibliographic Databases"). Additionally, Dagstuhl's open-access publishing service was supported by a generous donation in the amount of 42,000 € from the Heidelberg Institute for Theoretical Studies (HITS).

Seminare und Workshops

In 2017 wurden 102 Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops gestellt. Dies ist zwar weniger als im Vorjahr, entspricht aber der langfristigen Tendenz zu etwa 100 Anträgen pro Jahr. Auch durch die gute Antragslage des Vorjahrs stieg die Anzahl der Seminare und Workshops in 2017 auf 77. Dies sind zwei Veranstaltungen mehr als das bisher erreiche Maximum von 75 Veranstaltungen in 2014.

Seminars and Workshops

In 2017, 102 proposals were submitted for Dagstuhl seminars and Dagstuhl Perspectives workshops. Although this is slightly less than in the previous year, it confirms to the long-term trend of about 100 applications per year. Due to last year's high number of applications, the number of seminars and workshops held in 2017 rose to 77. This is two events more than the formerly highest number of 75 events per year achieved in 2014.

[2] https://www.leibniz-gemeinschaft.de/fileadmin/user_upload/downloads/Evaluierung/Senatsstellungnahmen/9.3.2017/LZI_-_Senatsstellungnahme_vom_09-03-2017_mit_Anlagen.pdf

Von den mehr als 3 300 Gästen, die sich in Dagstuhl trafen, nahmen etwa 2 500 an Seminaren teil. Etwa 43% 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.

Etwa 64% aller in 2017 stattgefundenen Seminare hatten mindestens eine Frau im drei- bis fünfköpfigen Organisatorenteam, und fast 20% aller Organisatoren waren Frauen. Der Frauenanteil unter allen Seminarteilnehmern betrug 19%.

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

Bibliographiedatenbank dblp

In Jahr 2017 konnte die *dblp computer science bibliographie* erneut mehr als 400 000 neue Publikationen aufnehmen. Dies entspricht mehr als 1 500 neuen Publikationen pro Arbeitstag. Ende 2017 konnte dblp so einen neuen Meilenstein erreichen: Die Datenbank indexiert nun über 4 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 etwa 17,9 Millionen Seitenzugriffe von etwa 560 000 verschiedenen Nutzern aus aller Welt. Dies entspricht mehr als sechs Seitenzugriffen pro Sekunde; zu Spitzenzeiten erfolgen über 100 Zugriffe gleichzeitig. Im Durchschnitt beginnt etwa alle 2,5 Sekunden ein neuer Nutzer, mit dblp zu arbeiten.

Im Verlauf des Jahres 2017 hat dblp seine Bemühungen intensiviert, externe Personen-Identifikatoren in dblp zu integrieren, insbesondere ORCIDs und Wikidata-IDs. Ende 2017 listete dblp bereits etwa 750 000 Autoren-Publikations-Paare, für die eine ORCID bekannt ist. Mehr als 11 500 Autoren-Profile sind dabei mit verifizierten ORCIDs verlinkt, und über 4 000 prominente Profile sind nun mit Wikidata verknüpft. Diese Verknüpfungen ermöglichen dblp, externe Identifikatoren zu konsolidieren, widersprüchliche Informationen aufzudecken und weitere offene, biographische Informationen zu erschließen.

Mehr Informationen zu dblp finden sich in Kapitel 3.

Dagstuhl Publishing

Wie in den Vorjahren haben die Open-Access-Publikationsaktivitäten auch in 2017 starken Zuspruch bekommen. So wurden allein in der Konferenzbandreihe LIPIcs erstmals über 1 100 Publikationen innerhalb eines Jahres veröffentlicht. Zudem gab es auch in 2017 wieder viele Bewerbungen von wissenschaftlichen Konferenzen zur Veröffentlichung des Konferenzbandes in der Serie LIPIcs. Außerdem konnte der renommierte Wissenschaftler Luca Aceto als Hauptherausgeber der Serie LIPIcs gewonnen werden.

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

About 2,500 of the 3,300 guests hosted at Dagstuhl participated in seminars. About 43% 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.

About 64% of all seminars had at least one woman organizer in organizer teams of three to five scientists, whereas almost 20% of all organizers were women. The ratio of women to all seminar participants was 19%.

See Chapter 2 for more details.

dblp computer science bibliography

During 2017, the dblp database grew again by more than 400,000 new publication records. This corresponds to more than 1,500 new records for each working day of the year. In doing so, dblp was able to reach a new milestone: the database now indexes more than 4 million publications across all areas of computer science.

Up to 17.9 million web pages are visited each month by about 560,000 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 100 request are made concurrently. On average, every 2.5 seconds, a new user session is started.

During 2017, dblp intensified its efforts in integrating external person identifiers schemes into dblp, in particular ORCIDs and Wikidata IDs. By end of 2017, dblp already listed about 750,000 author-publication pairs for which an ORCID is known. More than 11,500 author profiles do now contain linked and verified ORCIDs, and more than 4,000 prominent dblp profiles have been linked to Wikidata. These links help dblp to consolidate external identifiers, to monitor for contradictory information, and to query for further open biographical information.

More information about dblp can be found in Chapter 3.

Dagstuhl Publishing

As in the previous years, Schloss Dagstuhl's open-access publishing services experienced an on-going strong increase in demand from the community in 2017. For the first time, more than 1,100 articles have been published within one year in the conference proceedings series LIPIcs. Furthermore, LIPIcs again received and accepted proposals from several major scientific conferences. The renowned researcher Luca Aceto has been voted as the chair of the LIPIcs editorial board.

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

Öffentlichkeitsarbeit und Weiterbildung

Auch in 2017 wurde die erfolgreiche Vortragsreihe „Dagstuhler Gespräche“ als gemeinsame Veranstaltung von Schloss Dagstuhl und der Stadt Wadern fortgeführt. Ziel dieser Gespräche ist es, der interessierten Öffentlichkeit die breite Vielfalt der Informatik und deren praktische Anwendungen im Alltag oder in wirtschaftlichen Prozessen nahezubringen und in einen gemeinsamen Dialog einzusteigen.

Den Auftakt machte am 19. Mai 2017 der Präsident der Leibniz-Gemeinschaft, Prof. Matthias Kleiner. Unter dem Titel „Die Vielfalt der Wissenschaften“ gab er einen breiten Einblick in die vielfältigen Zukunftsthemen, denen sich die Leibniz-Forscherinnen und -Forscher widmen, seien es der Erhalt der Artenvielfalt, das Erkunden der Milchstraße und neuer Planeten, oder die Entwicklung individualisierter Medizin. Kleiner betonte dabei die Wichtigkeit, relevante Themen dabei auch öffentlich und konstruktiv-kontrovers zu debattieren, damit gesellschaftliches und politisches Handeln von wissenschaftlichen Erkenntnissen profitieren kann.

Am 24. November 2017 führte Prof. Michael Backes vom CISPA Helmholtz Center i.G. die Reihe mit einem Vortrag zum Thema „Autonome Systeme der Zukunft – Sicherheit und Datenschutz als Standortvorteil“ fort. An greifbaren Beispielen wie selbstfahrenden Autos, dem smarten Eigenheim oder autonomen Industrierobotern zeigt Prof. Backes auf, wie nachhaltgig sich unsere Gesellschaft in den kommenden Jahren verändern wird — und welche Chancen und Risiken in diesem Wandel liegen.

Die Vorträge fanden regen Anklang und lösten angeregte Diskussionen aus, und werden im kommenden Jahr gewiss fortgesetzt werden.

In Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisierte Schloss Dagstuhl 2017 zum 27. 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.

Public Relations and Professional Training

The successful series of “Dagstuhler Gespräche” (Dagstuhl Talks) was continued as a joint event of Schloss Dagstuhl and the city of Wadern. These talks aim at giving the interested public an understanding of the broad range of computer science and its practical applications in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

The first talk of 2017 was given by the President of the Leibniz Association, Prof. Matthias Kleiner, on May 19. Under the title “Die Vielfalt der Wissenschaften” (The Variety of Sciences), the talk gave an insight into the various reasearch topics Leibniz researchers are working on, be it the conservation of species, exploring the Galaxy and new planets, or developing individualized medicine. Kleiner emphasized the importance of discussing relevant topics publicly, constructively, and controversially, so that social and political action can benefit from scientific knowledge and research discoveries.

On November 24, 2017, Prof. Michael Backes of the CISPA Helmholtz Center i.G. continued the series with a talk on the topic of “Autonome Systeme der Zukunft – Sicherheit und Datenschutz als Standortvorteil” (Autonomous Systems of the Future – Security and Data Protection as a Locational Advantage). Backes used tangible examples, such as autonomous driving, smart homes, or autonomous industrial robots, to illustrate the lasting societal changes of the coming years — as well as the opportunities and risks that come with this change.

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

In 2017, Schloss Dagstuhl hosted its teacher training course for the 27th 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.

Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch im Jahr 2017 gab es wieder eine Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum[3] (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. Fünf

Joint Outreach with the Heidelberg Laureate Forum

2017 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum[3] (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. Five participants were selected and invited to participate in the Dagstuhl Seminar “Recent Trends in

ausgewählte Teilnehmer des HLF 2017 erhielten in der Woche vor der fünften Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Recent Trends in Knowledge Compilation“ (17381).

Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2018 zugestimmt.

Spender und Förderer

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

2017 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 647 Bände als Spende, davon 602 Monographien des Springer-Verlags im Wert von 42 119 €.

Dagstuhl Publishing wurde in 2017 mit 42 000 € durch das Heidelberger Institut für Theoretische Studien (HITS) unterstützt.

Das Allen Institute for Artificial Intelligence (AI2) hat Schloss Dagstuhl eine zweckgebundene Spende in Höhe von 80 000 US$ zur Unterstützung der Aktivitäten in den Leistungsbereichen Bibliographiedatenbank dblp und Dagstuhl Publishing zukommen lassen.

NSF Förderung von Nachwuchswissenschaftlern

Seit 2013 stehen Mittel zur Unterstützung von Nachwuchswissenschaftlern aus den USA bei der Teilnahme an Dagstuhl Seminaren zur Verfügung. Diese Fördermöglichkeit wird durch die National Science Foundation (NSF) finanziert[4].

Im Berichtsjahr konnte durch die Förderung 42 Forschern aus den USA eine Teilnahme an insgesamt 31 Seminaren ermöglicht werden. Insgesamt wurden dafür 70 442,58 $ Fördermittel ausgegeben. Die Förderung wurde um ein weiteres Jahr bis September 2018 verlängert.

Baumaßnahmen und Renovierung

Seit Mai 2017 steht der Vortragssaal „Karlsruhe“ nun in modernisierter und vergrößerter Form wieder voll umfänglich zur Verfügung. Der neue Hörsaal bietet Platz für bis zu 36 Seminarteilnehmer. In der Nähe der Tische befinden sich versenkbare Bodensteckdosen. Als Tafeln dienen sechs Whiteboards[5]. Am Pult besteht für die Vortragenden die Möglichkeit, sich per HDMI und VGA mit dem Beamer zu verbinden. Gäste können den installierten Beamer über einen Touchscreen steuern. Als Projektionsfläche dient eine Weißwand, die sich in der Mitte der Whiteboards befindet, wie Fig. 1.2 zeigt. Zum Präsentieren ist ebenfalls eine Mikrofonanlage vorhanden. Der Hörsaal ist mit zwei Access-Points ausgestattet, welche die WLAN Versorgung der Gäste gewährleisten. Der Raum ist zudem mit einer Belüftungs- sowie einer Klimaanlage ausgestattet.

Knowledge Compilation” (17381), taking place during the week before the fifth edition of the forum.

Satisfied with the outstanding success of the initiative, both partners agreed to continue the cooperation in 2018.

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 647, including 602 monographs at the total value of 42,119 € donated by Springer Science+Business Media publishing house.

The Heidelberg Institute for Theoretical Studies (HITS) supported Dagstuhl Publishing with 42,000 €.

The Allen Institute for Artificial Intelligence (AI2) donated 80,000 US$ to Schloss Dagstuhl to support the activities of the dblp database and of Dagstuhl Publishing.

NSF Grant for Junior Researchers

Since 2013, funds have been available to support researchers from the USA in participating in Dagstuhl events. This grant opportunity[4] is financed by the National Science Foundation (NSF).

In 2017, 42 US-based scientists were supported with a total amount of 70,442.58 $ and hence were able to participate in 31 Dagstuhl Seminars overall. The grant was extended for one more year until September 2018.

Construction Work and Renovation

Since May 2017, the new lecture hall “Karlsruhe” is available again in its modernized and expanded form. The room now can accommodate up to 36 seminar participants. Near the tables, there are retractable floor power sockets. The room has six whiteboards[6]. At the lecturer’s desk, speakers can connect to the projector via HDMI and VGA. Guests can control the installed projector via a touch screen. The projection screen is in the middle of the whiteboard, as shown in Fig. 1.2. Also, a microphone system is available for presentations. The lecture hall is equipped with two access points, which guarantee Wi-Fi access for the participants. The room is also equipped with ventilation and air conditioning.

Towards the end of the year, the construction of the new parking lot was finished by planting around its edges, as

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

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

Fig. 1.2
The modernized and expanded lecture hall “Karlsruhe”.

Gegen Ende des Jahres wurde mit der Bepflanzung des Randes auch der neue Parkplatz fertiggestellt, wie in Fig. 1.1 zu erkennen. Zudem wurde die Buchsbaumpflanzen im Schlossgarten wurden durch eine eine widerstandsfähigere Bepflanzung mit ähnlicher Optik ersetzt.

Ausstattung

In 2017 wurde im Neubau fünf Zimmer mit einem zusätzlichen Schlafsofa ausgestattet. Dies soll eine bequemere Nutzung des Zimmers als gegebenenfalls Doppelzimmer oder Familienzimmer erlauben. Zudem wurden insgesamt drei Zimmer mit Betten in Überlänge ausgestattet.

Im April 2017 wurde die gesamte Mechanik des Flügels im Musikzimmer im Schloss general überholt. Eine Neubeschaffung, wie zeitweise auch erwägt, erwies sich als nicht notwendig.

Der Billardtisch im Schloss wurde zudem komplett renoviert.

visible in Fig. 1.1. In addition, the boxwood plants in the courtyard were replaced by a more resistant planting with a similar appearance.

Facilities

In 2017, five rooms in the new guest house were equipped with an additional sofa bed each. These beds should allow a more comfortable use of the rooms as double or family rooms. Furthermore, a total of three rooms were equipped with extra-long beds.

In April 2017, the entire mechanism of the grand piano in the castle’s music room was completely overhauled. A new procurement, as sometimes considered, proved to be unnecessary.

The pool table in the castle has also been completely renovated.

[5] Zusätzlich zu den Whiteboards existiert ebenfalls noch die alte Tafelanlage an der Seitenwand des Vortragsraums.
[6] In addition to the whiteboards, there is also an old blackboard on the side wall of the lecture room.

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.[7]

Kapitel 6 enthält Zusammenfassungen der Dagstuhl-Seminare und Perspektiven-Workshops. Im Kapitel 14 sind alle Veranstaltungen, die 2017 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.[7]

Chapter 6 contains a collection of the summaries of the 2017 Seminars and Perspectives Workshops. Chapter 14 provides a comprehensive list of all events that took place during the year under review, and a seminar program covering the coming 24 months is available on the Dagstuhl website.

[7] https://www.dagstuhl.de/dagrep/

2.2 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 2017 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[8] veröffentlicht.

Eine Liste der vergangenen und kommenden Dagstuhl-Perspektiven-Workshops ist auf der Dagstuhl-Website verfügbar.[9]

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 2017 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[8].

A list of past and upcoming Dagstuhl Perspectives Workshop can be found on our web site.[9]

2.3 Einreichung der Anträge und Begutachtungsverfahren

2.3 *Proposal Submission and Review Process*

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch Auswahl der Anträge gewährleistet, die aus Sicht von Schloss Dagstuhl das größte Potential haben, abseits

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

[8] https://www.dagstuhl.de/dagman
[9] https://www.dagstuhl.de/pw-list

Towards Cross-Domain Performance Modeling and Prediction: IR/RecSys/NLP
http://www.dagstuhl.de/17442

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2017.

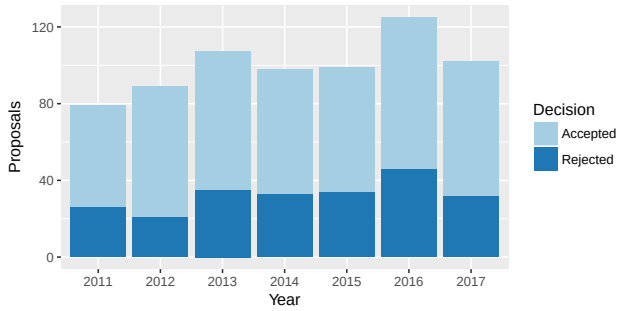

Fig. 2.2
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2011–2017.

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 2017 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 102 Anträgen in den Antragsrunden im Januar und Juni 2017. Der hohen Qualität der Anträge entsprechend, wurden etwa 69 % 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 70 in 2017 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 14 Seminare bereits 2017 ausgerichtet werden, der Großteil wurde jedoch für das Seminar-Programm in 2018 eingeplant (hier und im Folgenden wird, sofern nicht anders angegeben, das Wort “Seminar” sowohl für Dagstuhl-Seminare als auch für Dagstuhl-Perspektiven-Workshops verwendet). Nur fünf der 2017 genehmigten Seminare werden – oft auf ausdrücklichen Wunsch der Organisatoren – in 2019 stattfinden.

disciplines – in scientific discussions on new and important research problems and their most promising solutions, outside of the existing conferences.

The center solicits topics for new seminars and workshops twice a year from leading researchers worldwide, who submit their proposals together with a list of potential scientists to be invited. The proposals and suggested invitee lists are then reviewed by Dagstuhl’s Scientific Directorate (see Section 11.3) and finally discussed and decided during a two-day meeting at Schloss Dagstuhl.

This process ensures that every Dagstuhl Seminar and Dagstuhl Perspectives Workshop is backed by a strong team of organizers, addresses a topic of relevance to the computer science community, presents a coherent and well-structured scientific agenda, and brings together the right group of participants whose collective expertise can lead to a significant breakthrough in the area to be addressed. The balance of research communities and geographical regions, and especially the inclusion of junior and female researchers, are also taken into account during the review process.

The international scientific community expressed a lively interest in organizing seminars and workshops at Schloss Dagstuhl in 2017, submitting 102 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2017 and June 2017 submission rounds. The quality of the proposals was excellent, resulting in a 69 % 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 70 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2017 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 2018 seminar program, although it was possible to schedule 14 of them already in 2017 (here and in the following, the word "seminar" is meant to include both Dagstuhl Seminars and Dagstuhl Perspektives Workshops, if not specified otherwise). Only five seminars approved in 2017 will be held in 2019, often due to an explicit request by the organizers.

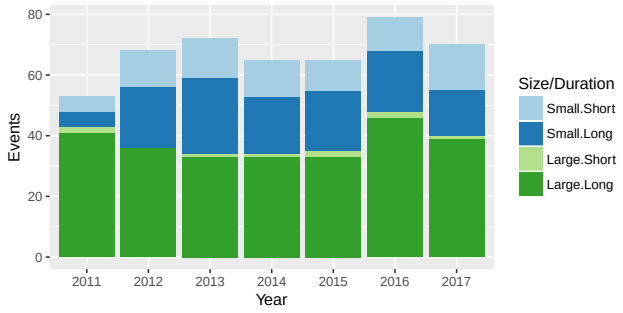

Fig. 2.3
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2011–2017. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

Seminar-Programm 2017

2.4 *The Seminar Program in 2017*

In 45 von 48 Wochen, in denen das Tagungszentrum 2017 geöffnet war, fand mindestens ein Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop statt. In 32 Wochen waren es sogar zwei. In drei Wochen war das Zentrum nur durch andere Veranstaltungen belegt.

Seit 2012 ist es aufgrund des damals fertiggestellten Gästehauses möglich, zwei Seminare parallel in einer Woche zu veranstalten. Dadurch ist, verglichen mit den Jahren zuvor, seit 2012 die Gesamtanzahl an Seminaren pro Jahr gestiegen. 2017 fanden insgesamt 77 Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops statt. In Fig. 2.4 ist die Entwicklung der vergangenen Jahre dargestellt.

At least one Dagstuhl Seminar or Dagstuhl Perspectives Workshop was held in 45 of the 48 weeks the center was open in 2017. In 32 weeks, there were actually two seminars in parallel. In three weeks, there were exclusively other events scheduled.

Since the guest house opened in 2012, it has been possible for the center to schedule two parallel seminars in any given week. Thus, there was an increase of seminars held since 2012 compared with the years before. Altogether, there were 77 Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2017. 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 2017 knapp die Hälfte (46 %, 1 126 von 2 433) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 18 % der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten, weitere 10 % an zweien. Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass etwa 43 % der Antwortenden 2017 das erste Mal, 16 % zum zweiten Mal und weitere 10 % zum dritten Mal (siehe Fig. 2.5a) teilgenommen haben.

Ein beträchtlicher Anteil der Gäste besteht aus jungen Wissenschaftlern, die am Anfang ihrer Karriere stehen, und für die der Aufenthalt in Dagstuhl oftmals prägend ist für den weiteren Verlauf ihres Lebenswegs. Etwa 35 % der Gäste der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops in 2017, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.5b). Diese ausgewogene Verteilung zwischen Nachwuchswis-

Participants in Dagstuhl Seminars come from all over the world, and a significant number of them choose to repeat the experience. Nevertheless, we see many fresh new faces every year, reflecting the changing informatics research across the globe. As in the previous years, in 2017, a bit less than half (1,126 of 2,433, or 46 %) of the researchers were first-time visitors to Dagstuhl. About an additional 18 % of the participating researchers had already attended one previous seminar in the years before, and another 10 % had already attended two. Slightly different numbers are obtained from our guest survey: About 43 % of the responders were first-time visitors, an additional 16 % state their second visit, and yet another 10 % their third (see Figure 2.5a).

A substantial number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 35 % of 2017 Dagstuhl Seminar and Dagstuhl Perspectives Workshop survey respondents 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

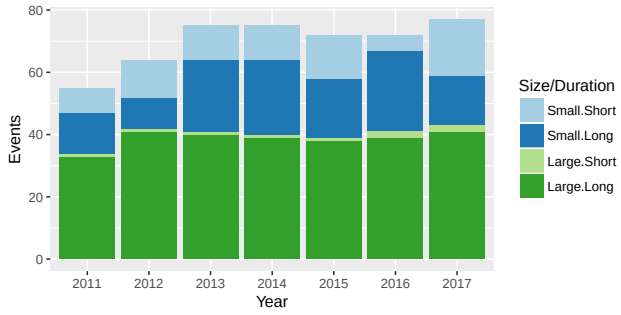

Fig. 2.4
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2011–2017. Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

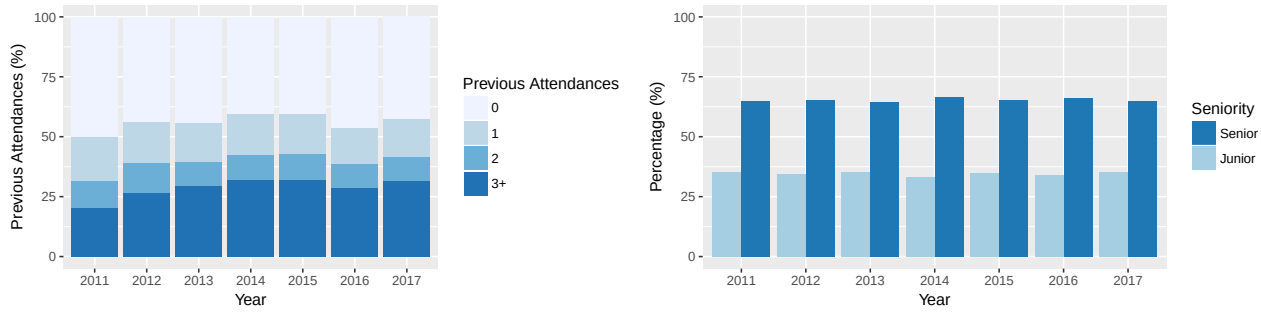

(a) Distribution of the number of previous attendances of participants, according to survey data.

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

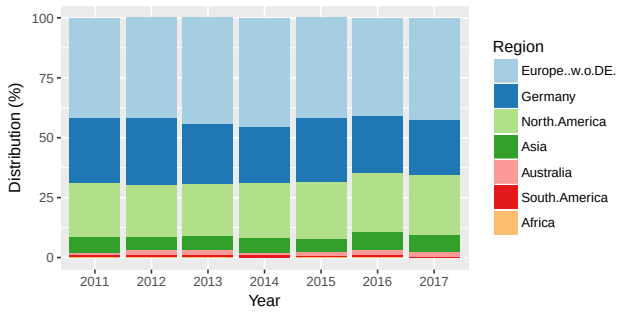

(c) Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2011–2017.

senschaftlern 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 77 % war der Anteil von Seminarteilnehmern aus dem Ausland 2017 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2017 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details können Kapitel 13 entnommen werden.

In 2017 waren etwa 64 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 19 % höher als in den Jahren zuvor, was auch 2016 schon der Fall war (siehe Fig. 2.6b).

and their senior colleagues.

At around 77 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2017. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2017. For a detailed breakdown please refer to Chapter 13.

In 2017, 64 % of all organizer teams in our scientific seminar program were mixed with respect to gender (see Fig. 2.6a). The percentage of female seminar participants was higher than in previous years at 19 %, continuing the trend from 2016 (see Fig. 2.6b).

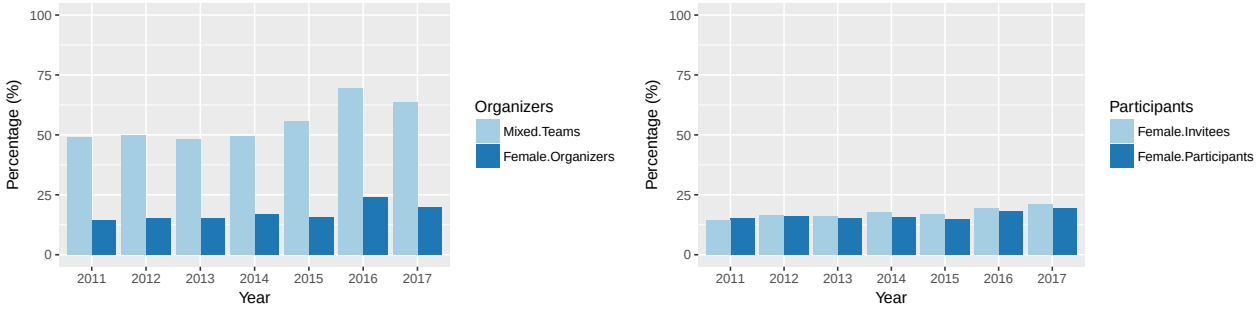

(a) Female organizers and mixed-gender organizer teams.

(b) Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2011–2017.

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 2017.

In den Seminaren, die sich Themen aus der theoretischen Informatik gewidmet haben wurden klassische Themen wie *Computability Theory* (17081) aber auch Themen aus dem Bereich der Künstlichen Intelligenz wie z. B. *Game Theory in AI, Logic, and Algorithms* (17111) diskutiert. Ergänzt wurde dies durch Themen, die Anwendungen in aktuellen Gebieten der Informatik erörterten – wie *Computer-Assisted Engineering for Robotics and Autonomous Systems* (17071), *Formal Synthesis of Cyber-Physical Systems* (17201), *Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions* (17431), oder *Big Stream Processing Systems*(17441).

Die fortschreitende digitale Vernetzung der Welt gab Anlass für Seminare im Bereich *Netzwerke*. Beispiele sind *The Critical Internet Infrastructure Revisited* (17511) oder *Foundations of Wireless Networking* (17271). Mit Seminaren wie *Cybersafety in Modern Online Social Networks* (17372) und *Online Privacy and Web Transparency* (17162) wurde eine Brücke zu den Themen Security, Cryptography und Privacy geschlagen, die bei den Seminaren in 2017 weiterhin gut vertreten waren. Diese Themen wurden von den Grundlagen wie *Quantum Cryptanalysis*

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

The following overview gives some topical focal points and a few respective seminars from 2017. 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 2017. The seminar summaries in Chapter 6 provide a full overview of the 2017 scientific seminar program.

Among the seminars which addressed topics from theoretical computer science, there were classical topics like *Computability Theory* (17081) but also topics from artificial intelligence like *Game Theory in AI, Logic, and Algorithms* (17111). This has been supplemented by topics that discuss current applications in the area of computer science like *Computer-Assisted Engineering for Robotics and Autonomous Systems* (17071), *Formal Synthesis of Cyber-Physical Systems* (17201), *Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions* (17431), or *Big Stream Processing Systems*(17441).

The ever-increasing digital interconnectedness gave motivation for several seminars about *Networking* topics, e.g. *The Critical Internet Infrastructure Revisited* (17511) or *Foundations of Wireless Networking* (17271). Seminars like *Cybersafety in Modern Online Social Networks* (17372) and *Online Privacy and Web Transparency* (17162) bridged the gap between *Networking* and topics like *Security*, *Cryptography* and *Privacy*. These topics were well represented in 2017 seminars and were discussed in a wide range: Starting from *Quantum Cryptanalysis* (17401) over *Opportunities and Risks of Blockchain Technologies* (17132) to applications in *Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition* (17281).

(17401) über Blockchains (*Opportunities and Risks of Blockchain Technologies* (17132)) bis hin zu Anwendungen wie *Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition* (17281) 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 *From Characters to Understanding Natural Language (C2NLU): Robust End-to-End Deep Learning for NLP* (17042) oder *Deep Learning for Computer Vision* (17391) diskutiert.

Erfreulicherweise waren auch Seminare über die verschiedenen Anwendungen der Informatik im täglichen Leben vertreten, wie z. B. *User-Generated Content in Social Media* (17301) oder *Citizen Science: Design and Engagement* (17272). Selbst Wahlsysteme und Abstimmungsmethoden waren durch das Seminar *Voting: Beyond Simple Majorities and Single-Winner Elections* (17261) abgedeckt.

Das Zusammenspiel von und die Kommunikation zwischen Menschen und Computern wurde ebenfalls diskutiert, vom klassischen Mensch-Maschine Interface wie in *Computational Interactivity* (17232) oder *Shape-Changing Interfaces* (17082) bis zu neuen Trends im Bereich am Körper tragbarer Geräte *Body-Centric Computing* (17392), sogar über virtuelle Realitäten hinaus, wie im Seminar *Beyond VR and AR: Reimagining Experience Sharing and Skill Transfer Towards an Internet of Abilities* (17062).

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2017 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.

Also the *Machine Learning* area, 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 *From Characters to Understanding Natural Language (C2NLU): Robust End-to-End Deep Learning for NLP* (17042) or *Deep Learning for Computer Vision* (17391).

Fortunately, also several seminars on applications of computer-science in daily life took place in Dagstuhl during 2017, e.g. *User-Generated Content in Social Media* (17301) or *Citizen Science: Design and Engagement* (17272). Even the discussion of voting processes and election systems was covered in *Voting: Beyond Simple Majorities and Single-Winner Elections* (17261).

Another topic was the interactiopn between human and computers, from traditional HCI in *Computational Interactivity* (17232) or *Shape-Changing Interfaces* (17082) to new trends in the field of wearable devices in *Body-Centric Computing* (17392) or even beyond virtual reality in *Beyond VR and AR: Reimagining Experience Sharing and Skill Transfer Towards an Internet of Abilities* (17062).

This brief selection of seminars should not hide the fact that each of the 2017 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

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.

Further Event Types 2.7

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

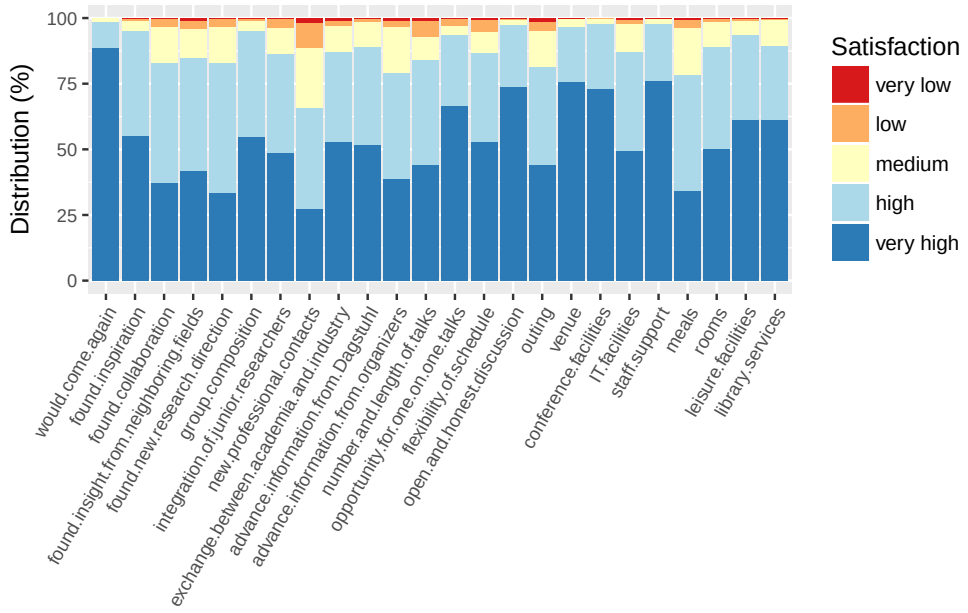

Fig. 2.7
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2017. According to survey results.

Qualitätssicherung 2.8 *Quality Assurance*

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

Fig. 2.7 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2017 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 474 Fragebögen, welche die Meinung von etwa 57 % der 2 572 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 Seminars and Dagstuhl Perspectives Workshops, the questionnaire containing questions about their satisfaction with the content of the event and the organization of their visit. The results of each questionnaire are made available to all of the center's departments every week, thus enabling a quick response to issues and requests. At the same time, anonymized results of the content questions are made available to the seminar participants via e-mail, typically in the week following their stay at the center. This enables the organizers to receive feedback on how the seminar went and tips for organizing future seminars. In 2013, Schloss Dagstuhl began sending the report as a PDF attachment with an enhanced visual layout.

Fig. 2.7 shows the satisfaction of responding participants in 2017 with regard to selected aspects of their stay. The results were compiled from 1,474 questionnaires, representing the responses of about 57 % of all 2,572 participants. These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

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

Auslastung des Zentrums 2.9

Auch 2017 konnte Schloss Dagstuhl die durch das neue Gästehauses ermöglichte hohe Auslastung weitgehend halten. Es gab 2017 insgesamt 12 988 Gasttage, wobei 11 091 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Letztere Zahl bedeutet einen leichten Anstieg verglichen mit 2016. Allerdings gab es insgesamt weniger Gasttage als in 2016. Es fanden im Berichtsjahr 121 Veranstaltungen mit insgesamt 3 308 Gästen statt. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2017 ebenso unbelegt wie eine Woche zum Jahresanfang, zwei Wochen im Juli/August und eine Woche am Jahresende. Diese wurden zu Instandhaltungs- und Verwaltungsarbeiten benötigt.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2017 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[10] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

Utilization of the Center

Thanks to the new guest house, Schloss Dagstuhl was able to uphold the high capacity utilization again in 2017. There were 12,988 overnight stays in total, with 11,091 overnight stays in Dagstuhl Seminars and Dagstuhl Perspectives Workshops. The latter number was a bit higher than in 2016. However, there were fewer overnight stays in total in 2017 compared to stays in 2016. The center hosted a total of 121 events with 3,308 guests in 2017. See Chapter 13 for further details.

Weekends were kept free in 2017, as well as a week at the beginning of the year, two weeks in July/August, and a week at the end of the year, this time being required for maintenance work to building facilities and administrative work.

A comprehensive listing of all events at Schloss Dagstuhl in 2017, 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[10] of upcoming events and further information and materials on all events past, present and future.

[10] https://www.dagstuhl.de/no_cache/programm/kalender/

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.[11] 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.[12] Hence, dblp has become indispensable to the computer science community as both a research tool and a research data set.

[11] Google Scholar liefert zum Suchbegriff „dblp“ über 28 900 Treffer; im Einzelnen weisen SpringerLink ca. 2 750 Artikel, Elsevier ScienceDirect über 650 Artikel, die ACM Digital Library ca. 1 850 Artikel und IEEE Xplore über 1 800 Artikel nach.

[12] The search term “dblp” results in 28,900 hits at Google Scholar; in particular, SpringerLink lists about 2,750 articles, Elsevier ScienceDirect lists more than 650 articles, the ACM Digital Library lists 1,850 articles, and IEEE Xplore lists more than 1,800 articles.

Schloss Dagstuhl und dblp

3.2 Schloss Dagstuhl and dblp

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.

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.

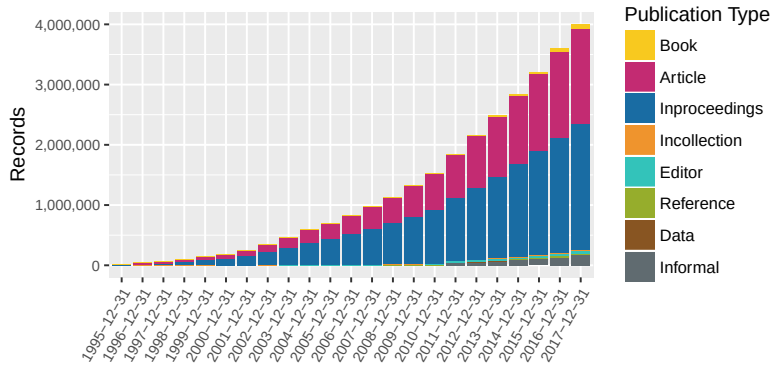

(a) Total number of records by year and type

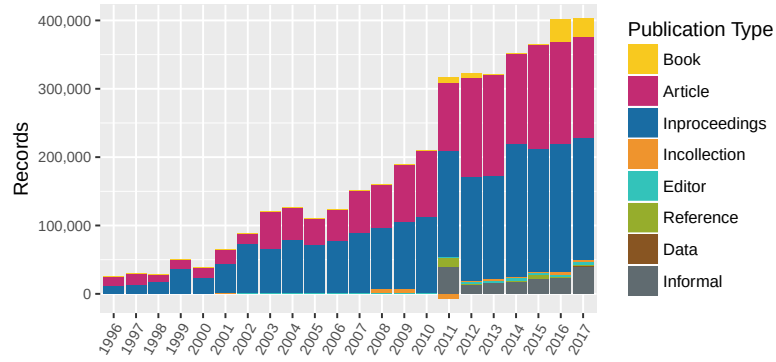

(b) New records by year and type

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.

Am 21. Dezember 2017 wurde dabei ein neuer Meilenstein erreicht: dblp listet nun mehr als 4 Millionen Publikationen. Im Zeitraum von Anfang Januar 2017 bis Ende Dezember 2017 wurden mehr als 400 000 neue Publikationseinträge in dblp aufgenommen. Dies entspricht mehr als 1 500 neuen Publikationen pro Arbeitstag. Somit konnte nun bereits zum vieren mal in Folge die Rekordaufnahmequote des Vorjahres übertroffen werden. Die neu aufgenommenen Einträge verteilen sich zu 44,4% auf Konferenzbeiträge, zu 36,7% auf Journalartikel, zu 6,5% auf Monographien und Dissertationen, sowie zu 12,4% 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.

On December 21, 2017, a new milestone has been reached: dblp now lists more than 4 million publications. Between January 1, 2017, and December 31, 2017, the dblp database grew by more than 400,000 publication records. This corresponds to more than 1,500 new records for each working day of the year. Hence, for four successive years, dblp has been able to surpass the number of new records included in the previous year. These new records consist of 44.4% conference papers, 36.7% journal articles, 6.5% monographs and PhD theses, and 12.4% other publications.

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

	Trier 1		Trier 2		Dagstuhl		Total	
	2016	2017	2016	2017	2016	2017	2016	2017
user sessions (visits) per day	26 911	27 931	1 427	2 836	1 254	5 366	29 593	36 133
page views per day	501 208	466 989	26 354	35 140	35 406	85 537	562 969	587 668
page views per user session	18,6	16,7	18,5	12,4	28,2	15,9	19,0	16,3
distinct users (IPs) per month	393 273	390 886	25 249	58 975	20 416	86 985	438 938	536 847
data served per month	1 187,6 GB	1 082,3 GB	72,7 GB	82,8 GB	120,7 GB	235,0 GB	1 381,0 GB	1 400,1 GB

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

Nutzungsstatistiken

3.4

Usage Statistics

2017 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

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

Die 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 der 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 2017, three official dblp web servers were updated and synchronized on a daily basis:

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

The 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. Server Trier 1 is by far the most widely known due to its high visibility and prominence in the Google search engine. However, server Dagstuhl has become increasingly more visible, and the number of its users quadrupled during the course of 2017.

Figure 3.3 shows the average usage of all three servers in 2017. These figures ignore the traffic caused by known bots and crawlers.

Integration von externen Identifikatoren

3.5

Integration of external identifier schemes

Im Verlauf des Jahres 2017 hat dblp seine Bemühungen intensiviert, externe Personen-Identifikatoren in dblp zu integrieren. Während im Bereich der Publikationen die DOIs[13] bereits als eindeutiger, persistenter Bezeichner etabliert ist, steckt die Verwendung von Personen-Identifikatoren noch weitgehend im den Kinderschuhen. Mit ORCID[14] und Wikidata[15] haben sich jedoch zwei viel versprechende Initiativen etabliert.

ORCID ist ein weit verbreitetes, persistentes Identifizierungssystem für Forscher, welches Ende 2012 eingeführt wurde. Dem erfolgreichen DOI-Vorbild nachempfunden wird ORCID von allen großen Wissenschaftsverlagen unterstützt. Ende 2017 gab es bereits mehr als 4 Millionen registrierte ORCIDs, von denen etwa 1,5 Millionen mindestens einer Publikation im öffentlichen ORCID-Korpus zugeordnet waren. Inzwischen verlangen viele Verlage von ihren Autoren, dass sie bei der Einreichung oder Veröffentlichung eines Beitrags zwingend eine ORCID angeben.

Derzeit gibt es in dblp etwa 750 000 Signaturen (d.h. Autoren-Publikationspaare), für die eine ORCID bekannt

During 2017, dblp intensified its efforts in integrating external person identifiers schemes into dblp. While for publications, the use of DOIs[13] as unique, persitent identifiers has been widely adopted in the past years, in the case of person identifiers, only a few first steps have been made. The two most promising open person identifier scheme initiatives are ORCID[14] and Wikidata[15].

ORCID is a widely used persistent identifier scheme for researchers, introduced in late 2012. It has been modelled similar to the successful DOI scheme and has the support of all major scholarly publishers. As of end of 2017, there are more than 4 million registered ORCIDs, of which about 1.5 million have at least one publication listed in the public ORCID corpus. Many publishers now ask their authors to provide an ORCID when submitting or publishing a paper.

At the moment, there are about 750,000 signatures (i.e., author-publication pairs) in dblp for which an ORCID is known. This is about 6% of all signatures. We expected ORCIDs to help us identify authors and create clear, unambiguous author bibliographies, and we were not

ist. Dies sind etwa 6% aller Signaturen in dblp. Die Hoffnung, dass ORCIDs uns helfen, Autoren zu identifizieren und saubere, eindeutige Autorenbibliographien zu erstellen, wurde nicht enttäuscht (siehe Fig. 3.4). Im Zuge des ersten Datenabgleichs fanden wir in unserem Korpus

- über 600 Fälle, in denen ein Autorenprofil mit mehr als einem ORCID assoziiert war. Dies deutet auf Profile hin, die Publikationen von tatsächlich verschiedenen Autoren auflisten.
- über 5 000 Fälle, in denen dieselbe ORCID zu mehr als einem Autorenprofil assoziiert war. Dies deutet darauf hin, dass Publikationen eines einzelnen Autors über mehrere Profile verteilt wurden.

Die Kuration dieser fehlerhaften Profile ist ein kontinuierlicher und andauernder Prozess. Während wir zwar auch einige Fälle gefunden haben, in denen ORCID-Informationen falsch waren (z.B. Fälle, in denenen Publikationen dem falschen Autor zugeschrieben wurden), sind die Daten im Allgemeinen doch sehr zuverlässig. Entsprechend der Qualitätsstandards von dblp erfolgt jedoch für jede verlinkte ORCIS eine manuelle Nachkontrolle durch ein Mitglied des dblp-Teams.

Wikidata ist eine kollaborative Wissensdatenbank, die ursprünglich als strukturierte Datenbasis für Wikimedia-Projekte erdacht, allerdings inzwischen auch als offen lizensierte Datenquelle für die Allgemeinheit bereitsteht. Für dblp fungiert Wikidata als kuratierte, Linked-Open-Data-Quelle über alle Wissenschaftsdisziplinen hinweg. Die Verknüpfung mit Wikidata ermöglicht es dblp, externe Identifikatoren zu konsolidieren, widersprüchliche Informationen aufzudecken und weitere offene, biographische Informationen zu erschließen.

Ende 2017 waren bereits mehr als 4 000 prominente dblp-Profile mit Wikidata verknüpft. Die Entwicklung der Anreicherung mit externen Personen-Identifikations-Schemata ist in Fig. 3.5 zusammengefasst.

disappointed (see Fig. 3.4). For the initial import, we found in our corpus

- more than 600 cases where an author profile was related to more than one ORCID. This indicates that the profile actually lists publications from different authors.
- more than 5,000 cases where the same ORCID appears in more than one author profile. This indicates that publications of a single author have been incorrectly distributed over multiple profiles.

The curation of the defective profiles is an on-going process. While we found several cases where ORCID information is wrong (e.g., authors accidentally claiming publications that where written by someone else), the data is in general very reliable. However, according to dblp's data quality standards, manual confirmation by a member of the dblp team is required for each linked ORCID.

Wikidata is a collaboratively edited knowledge base intended to provide an open, structured data source to be used by Wikimedia projects such as Wikipedia, as well as by anyone else, under a public domain license. For dblp, Wikidata serves as a curated, linked open data hub across all science disciplines. Linking to Wikidata allows dblp to consolidate external identifiers, to monitor for contradictory information, and to query for further open biographical information.

As of end of 2017, now more than 4,000 prominent dblp profiles have been linked to Wikidata. Fig. 3.5 summarizes the development of external person identification schemes in dblp.

[13] https://doi.org
[14] https://orcid.org
[15] https://wikidata.org

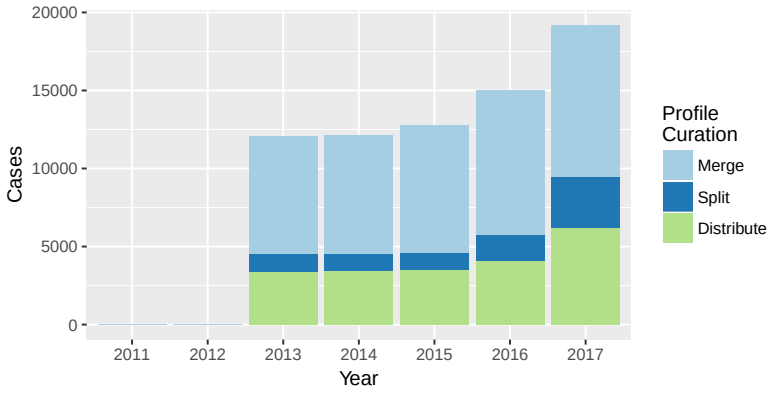

Fig. 3.4

Data curation of existing dblp person profiles. No curation data has been tracked before 2013. We distinguish between three curation cases: *Merge* = Two or more synonymous profiles have been merged into a single profile. *Split* = A single, homonymous profile has been split into two or more profiles. *Distribute* = A mixed case where records from two or more profiles have been redistributed between two or more profiles.

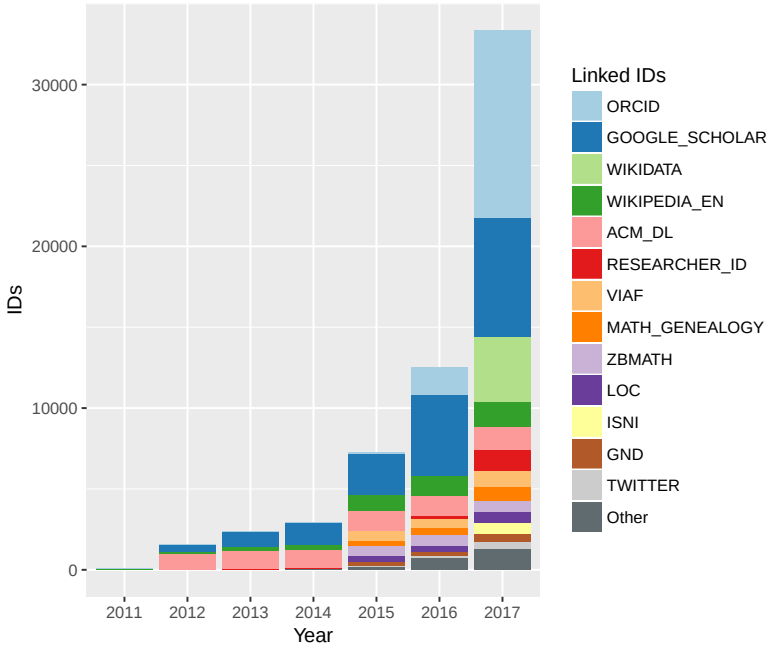

Fig. 3.5

Linked external person IDs in dblp. A single person profile may be linked to multiple external IDs.

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*[16] dokumentiert um eine Zitation der Seminare im wissenschaftlichen Kontext zu ermöglichen. Zudem erlaubt es auch den Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Dagstuhl-Seminaren und -Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 11.4) agiert als 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 2017 wurde für 48 Dagstuhl-Seminare und -Perspektiven-Workshops ein Bericht in der Reihe *Dagstuhl Reports* veröffentlicht. An dieser Stelle bedanken wir uns ganz herzlich bei den Organisatoren und Kollektoren für die erfolgreiche Zusammenarbeit.

Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*[16] which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

The periodical started with the first seminars of January 2011 and publishes, in monthly issues, reports on Dagstuhl Seminars and Perspectives Workshops that took place 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).

In 2017, 48 reports of Dagstuhl Seminars and Dagstuhl Perspectives Workshops have been published. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*[17] die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Dagstuhl-Perspektiven-Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 11.4) fungiert hier ebenfalls als Herausgebergremium. In 2017 wurde keine Ausgabe veröffentlicht. Allerdings wurden mehrere Manifestos durch das Herausgebergremium begutachtet und zur Veröffentlichung aktzepiert. Diese werden zeitnah in 2018 veröffentlicht werden.

Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*[17] in an Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. In 2017 no volume was published. However, several Dagstuhl Manifestos has been reviewed by the editorial board and has been accepted for publication. They will be published in early 2018.

Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*[18] ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspekti-

Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*[18] book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl

[16] http://drops.dagstuhl.de/dagrep
[17] 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 2017 konnte ein Buch veröffentlicht werden, siehe Fig. 4.1.

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[19] veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausgebergremium (Fig. 4.2), diskutiert sorgfältig alle Anträge, um ausschließlich qualitativ hochwertige sowie professionell durchgeführte Veranstaltungen in die Reihe aufzunehmen und um gegebenenfalls Empfehlungen zur Verbesserung der Veranstaltungsstruktur zu geben.

In 2017 wurden 3 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.3.

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[20] veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausgebergremium (siehe Fig. 4.4) besteht aus einschlägig bekannten Wissenschaftlern und wird seit Oktober 2017 von Luca Aceto als Hauptherausgeber geleitet.

Die Amtszeiten von Wolfgang Thomas und Pascal Weil sind 2017 ausgelaufen. Beide haben sowohl als Mitglieder des Herausgebergremiums aber auch als Hauptherausgeber eine herausragende Rolle in der Entwicklung der Serie gespielt. Für diese Verdienste möchten wir uns an dieser Stelle herzlich bedanken.

Die Amtszeit von Catuscia Palamidessi wurde 2017 nach einem anonymen Wahlverfahren innerhalb des Herausgebergremiums bis Mai 2021 verlängert. Zudem wurden Luca Aceto, Anca Muscholl, Thomas Schwentick und

Perspectives Workshop. Each book requires a proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). In 2017, one volumes was published; see Fig. 4.1.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[19] aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 4.2, discusses carefully all submitted proposals to ensure that only significant and professionally organized events are added to the series and that – if applicable – suggestions are given for improving the structure of the event.

In 2017, Dagstuhl published 3 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.3.

LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series[20] publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers (see Fig. 4.4) supervises the conferences that are accepted for LIPIcs and is headed since October 2017 by Luca Aceto.

The terms of Wolfgang Thomas and Pascal Weil ended in 2017. Both served as members of the editorial board as well as chairs of the board and played an important role for the development of the series. We would like to take this opportunity to thank them for the extraordinary dedication.

In 2017, the term of Catuscia Palamidessi was extended until May 2021, based on an anonymous voting within the editorial board. Furthermore, Luca Aceto, Anca Muscholl, Thomas Schwentick, and Reinhard Wilhelm were voted as new members of the editorial board. See also Fig. 4.4.

The series published the proceedings of 25 major

[18] http://drops.dagstuhl.de/dfu
[19] http://drops.dagstuhl.de/oasics

Vol. 7 | The Constraint Satisfaction Problem: Complexity and Approximability | Editors: Andrei Krokhin and Stanislav Zivny
https://www.dagstuhl.de/dagpub/978-3-95977-003-3

Fig. 4.1
Dagstuhl Follow-Ups volumes published in 2017.

Prof. Dr. Daniel Cremers
TU Munich, Germany

Prof. Dr. Barbara Hammer
Bielefeld University, Germany

Prof. Dr. Marc Langheinrich
University of Lugano, Switzerland

Prof. Dr. Dorothea Wagner
Karlsruhe Institute of Technology, Germany | Chair

Fig. 4.2
OASIcs Editorial Board.

Reinhard Wilhelm neu in der Gremium gewählt. Siehe auch Fig. 4.4.

In 2017 wurden Tagungsbände von 25 Konferenzen mit insgesamt mehr als 1100 Artikeln veröffentlicht, so viel wie noch nie zuvor; siehe Fig. 4.5.

In 2017 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. 4.6 sind alle Konferenzen aufgelistet, deren Anträge 2017 bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischweise 5 Jahre) eingegangen wurde.

conferences with more than 1100 articles in total in 2017, marking again a record high since the series was started; see Fig. 4.5.

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2017 has again seen several applications for LIPIcs, continuing the high interest from the previous years. Fig. 4.6 lists all conferences that have been accepted in 2017 for a cooperation covering several years (typically 5 years).

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

Vol. 56 | 6th Symposium on Languages, Applications and Technologies (SLATE 2017) https://www.dagstuhl.de/dagpub/978-3-95977-056-9
Vol. 57 | 17th International Workshop on Worst-Case Execution Time Analysis (WCET 2017) https://www.dagstuhl.de/dagpub/978-3-95977-057-6
Vol. 59 | 17th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2017) https://www.dagstuhl.de/dagpub/978-3-95977-042-2

Fig. 4.3
OASIcs volumes published in 2017.

Prof. Dr. Luca Aceto Gran Sasso Science Institute, Italy and Reykjavik University, Iceland | Chair | tenure started in June 2017
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
Prof. Dr. Anca Muscholl LaBRI and University Bordeaux, France | tenure started in June 2017
Dr. Catuscia Palamidessi INRIA, France | tenure extended in June 2017
Prof. Dr. Thomas Schwentick TU Dortmund, Germany | tenure started in June 2017
Prof. Raimund Seidel, Ph. D. Saarland University, Germany | tenure started in June 2017
Prof. Dr. Wolfgang Thomas RWTH Aachen, Germany | tenure ended in May 2017
Pascal Weil, Ph. D CNRS, France and University Bordeaux, France | tenure ended in May 2017
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Saarland University, Germany | tenure started in June 2017

Fig. 4.4
LIPIcs Editorial Board.

Vol. 63 | 11th International Symposium on Parameterized and Exact Computation (IPEC 2016)
https://www.dagstuhl.de/dagpub/978-3-95977-023-1

Vol. 66 | 34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-028-6

Vol. 67 | 8th Innovations in Theoretical Computer Science Conference (ITCS 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-029-3

Vol. 68 | 20th International Conference on Database Theory (ICDT 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-024-8

Vol. 70 | 20th International Conference on Principles of Distributed Systems (OPODIS 2016)
https://www.dagstuhl.de/dagpub/978-3-95977-031-6

Vol. 71 | 2nd Summit on Advances in Programming Languages (SNAPL 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-032-3

Vol. 72 | 7th Conference on Algebra and Coalgebra in Computer Science (CALCO 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-033-0

Vol. 74 | 31st European Conference on Object-Oriented Programming (ECOOP 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-035-4

Vol. 75 | 16th International Symposium on Experimental Algorithms (SEA 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-036-1

Vol. 76 | 29th Euromicro Conference on Real-Time Systems (ECRTS 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-037-8

Vol. 77 | 33rd International Symposium on Computational Geometry (SoCG 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-038-5

Vol. 78 | 28th Annual Symposium on Combinatorial Pattern Matching (CPM 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-039-2

Vol. 79 | 32nd Computational Complexity Conference (CCC 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-040-8

Vol. 80 | 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-041-5

Vol. 81 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-044-6

Vol. 82 | 26th EACSL Annual Conference on Computer Science Logic (CSL 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-045-3

Vol. 83 | 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-046-0

Vol. 84 | 2nd International Conference on Formal Structures for Computation and Deduction (FSCD 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-047-7

Vol. 85 | 28th International Conference on Concurrency Theory (CONCUR 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-048-4

Vol. 86 | 13th International Conference on Spatial Information Theory (COSIT 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-043-9

Vol. 87 | 25th Annual European Symposium on Algorithms (ESA 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-049-1

Vol. 88 | 17th International Workshop on Algorithms in Bioinformatics (WABI 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-050-7

Vol. 90 | 24th International Symposium on Temporal Representation and Reasoning (TIME 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-052-1

Vol. 91 | 31st International Symposium on Distributed Computing (DISC 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-053-8

Vol. 92 | 28th International Symposium on Algorithms and Computation (ISAAC 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-054-5

Fig. 4.5
LIPIcs volumes published in 2017.

DISC | International Symposium on Distributed Computing accepted for 2017–2021
GIScience | International Conference on Geographic Information Science accepted for 2018–2022
TIME | International Symposium on Temporal Representation and Reasoning accepted for 2017–2021
WABI | Workshop on Algorithms in Bioinformatics accepted for 2017–2021

Fig. 4.6
Conferences accepted in 2017 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. 4.7
LITES Editorial Board.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[21] veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 4.7), begutachtet alle eingereichten Arbeiten. Die Zeitschrift wird gemeinsam mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*[22] der Fachgesellschaft *European Design and Automation Association (EDAA)*[23] 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 2017 wurde eine Ausgabe von *LITES* mit insgesamt 8 Artikeln veröffentlicht.

DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*[24] werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2017 wurde die dritte Ausgabe mit zwei Heften und insgesamt 23 Artefakten 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“[25], der in 2016 durch zwei Seminare ergänzt wurde: „Reproducibility of Data-Oriented Experiments in e-Science“[26] und „Rethinking Experimental Methods in Computing“[27].

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*[21] 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. 4.7), reviews all submitted contributions. The journal is jointly published with the *EMbedded Systems Special Interest Group (EMSIG)*[22] of the *European Design and Automation Association (EDAA)*[23]. 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 2017, one issue of *LITES* containing 8 articles in total was published.

DARTS: Dagstuhl Artifacts Series

The *DARTS* series[24] publishes evaluated research data and artifacts. It is organized as a periodical. In 2017, one volume containing two issues with 23 artifacts in total was published.

The publishing of research data and artifacts is currently in the general focus of the scientific community and funding agencies. In the area of computer science, this topic is also under discussion. For example, in 2015 a Perspectives Workshop on “Artifact Evaluation for Publications”[25] took place which was complemented with two seminars in 2016: “Reproducibility of Data-Oriented Experiments in e-Science”[26] and “Rethinking Experimental Methods in Computing”[27].

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.

[21] http://drops.dagstuhl.de/lites
[22] http://www.emsig.net/
[23] https://www.edaa.com/
[24] https://www.dagstuhl.de/darts
[25] https://www.dagstuhl.de/15452
[26] https://www.dagstuhl.de/16041
[27] https://www.dagstuhl.de/16111

Infrastruktur 4.2 *Infrastructure*

Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 4.8. Die Bände aus den Reihen *LIPIcs* und *OASIcs* werden zudem bei Scopus[28] 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. 4.8.

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. 4.8. The LIPIcs and OASIcs volumes are submitted to Scopus[28] 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. 4.8.

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

[28] https://www.scopus.com

dblp
Dagstuhl Reports https://dblp.org/db/journals/dagstuhl-reports/
Dagstuhl Manifestos https://dblp.org/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups https://dblp.org/db/series/dfu/
OASIcs https://dblp.org/db/series/oasics/
LIPIcs https://dblp.org/db/series/lipics/
LITES https://dblp.org/db/journals/lites/
DARTS https://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. 4.8
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

Dagstuhl Publishing

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[29] 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 multidisziplinären Repositoriums.

AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Transformation und Open Science“[30] mit organisiert, welcher bereits der vierte Workshop in Folge seit 2013 ist. Der Workshop fand am 19. und 20. Januar 2017 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt.

AG Open Access der Schwerpunktinitiative „Digitale Information“

Die Allianz der deutschen Wissenschaftsorganisationen, zu der neben der Max-Planck-Gesellschaft, der Helmholtz-Gemeinschaft, sowie weiteren Organisationen auch die Leibniz-Gemeinschaft gehört, hat eine Schwerpunktinitiative „Digitale Information“ ins Leben gerufen, bei der auch das Thema *Open Access* als Handlungsfeld vertreten ist. Mit Dr. Marc Herbstritt wurde seitens der Leibniz-Gemeinschaft ab Juli 2013 ein Mitglied des wissenschaftlichen Stabs von Schloss Dagstuhl in die Arbeitsgruppe „Open Access“[31] 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. Mit dem Jahr 2017 endet auch die zweite Phase der Schwerpunktinitiative „Digitale Information“ und damit dann auch die Arbeitsgruppe.

Technisches Back-end: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)[32] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[33] berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer grossen

LeibnizOpen

The Leibniz Association has established the *LeibnizOpen*[29] 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”[30] was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop took place at the Leibniz Association headquarters in Berlin 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.[31]

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. The second phase of the priority initiative “Digital Information” ends with the year 2017 and so also the working group.

Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)[32]. The general guidelines of the Dublin Core initiative[33] 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

[29] http://www.leibnizopen.de/
[30] https://www.dagstuhl.de/dagpub/journalmanagement-leibniz/2017-01-19-workshop/
[31] https://www.allianzinitiative.de/archiv/open-access.html

Dagstuhl Publishing

Leserschaft zur Verfügung. Als technische Grundlage dient eine adaptierte Version des OPUS-Systems.[34]

Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)[35] 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:

- emis.de: Das unter Leitung des FIZ Karlsruhe, Leibniz-Institut für Informationsinfrastruktur, organisierte Mathematik-Publikations-Portal European Mathematical Information Service (EMIS) spiegelt alle Bände der LIPIcs-Reihe.[36]
- 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.[37]

for this is an adapted version of the OPUS system.[34]

Long-term Archiving

All publications are submitted to the German National Library (D-NB)[35] for (digital) long-term archiving.

Mirroring

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

- emis.de: The portal for electronic math resources European Mathematical Information Service (EMIS), organized under the auspices of FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, mirrors all volumes of the LIPIcs series[36].
- 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.[37]

[32] https://www.dagstuhl.de/drops
[33] http://dublincore.org/
[34] https://opus4.kobv.de/
[35] http://www.dnb.de/DE/Netzpublikationen/Langzeitarchivierung/langzeitarchivierung_node.html
[36] https://subs.emis.de/LIPIcs/
[37] http://vesta.informatik.rwth-aachen.de/Dagstuhl/

5 Resonanz
Feedback

Resonanz zu Seminaren und Workshops

5.1

Feedback on Seminars and Workshops

Langzeit-Feedback

Schloss Dagstuhl bekommt viel Feedback. Besonders erfreulich ist es, wenn uns lange nach einer Veranstaltung eine Rückmeldung erreicht, welchen positiven Einfluss Dagstuhl manchmal auf die Karriere seiner Gäste hat.

Long Term Feedback

Schloss Dagstuhl receives a lot of feedback. We are especially happy to get word of how Dagstuhl positively impacts the careers of our guests, often a long time after an event at Dagstuhl.

Carolina Dania, from "Mapping OCL as a Query and Constraint Language", Ph.D. Thesis

13211 – Automated Reasoning on Conceptual Schemas | Dagstuhl Seminar | https://www.dagstuhl.de/13211

This doctoral dissertation owes a great deal of its initial motivation and final focus to the very lively and insightful discussion that took place during the Dagstuhl Seminar "Automated Reasoning on Conceptual Schemas" (19- 24 May, 2013), which we have the fortune to participate in.

Nataliia Bielova (cited with appreciated permission)

12401 – Web Application Security | Dagstuhl Seminar | https://www.dagstuhl.de/12401

At the Dagstuhl-Seminar 12401 on Web Application Security, I've met a researcher at Inria working in the same field as me, and later on I was hired in her team. Now I am a permanent researcher at Inria.

Sebastian Böcker (cited with appreciated permission)

15492 – Computational Metabolomics | Dagstuhl Seminar | https://www.dagstuhl.de/15492

Hallo, da wir demnächst wieder zusammen kommen zu einem verwandten Thema, wollte ich kurz ein paar Resultate vermelden für das Seminar 15492 (Computational Metabolomics):

- Forschungssemester von mir bei Pieter Dorrestein (UCSD) 2016
- gemeinsames Paper mit Pieter Dorrestein und Nuno Bandeira (UCSD) bei Nature Communications
- gemeinsamer NIH Antrag mit Pieter Dorrestein und Nuno Bandeira
- gemeinsamer NIH Antrag mit David Wishart (University of Alberta – Edmonton, CA) und zwei weiteren
- gemeinsamer DFG Normalantrag mit Michael Anton Witting (Helmholtz Zentrum – München, DE), Einreichung Dez 2017

Alles zumindest indirekte Folgen des Dagstuhl-Seminars.
Schönen Gruß,[38]
Sebastian Böcker.

Resonanz von Teilnehmern

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

Feedback from Participants

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

[38] engl.: Hello, as we will meet again soon for a related topic, I wanted to give you a heads up about some results for the seminar 15492: -research sabbatical of mine with Pieter Dorrestein (UCSD) 2016 | -joint paper with Pieter Dorrestein and Nuno Bandeira (UCSD) in Nature Communications | -joint NIH Proposal with Pieter Dorrestein and Nuno Bandeira | - joint NIH Proposal with David Wishart (UAlberta) and two others | -joint DFG proposal with Michael Anton Witting (HMGU), to be submitted 12/2017 | All consequences, at least indirectly, of the Dagstuhl Seminar. Greetings,

m.c. schraefel

17392 – Body-Centric Computing | Dagstuhl Seminar | https://www.dagstuhl.de/17392

Please share my thanks and appreciation with the Dagstuhl organisation for once again providing excellent support to hold this seminar. It was a very productive meeting, in no small part because of the infrastructure and human support you provide. The room was a great space - opening into a useable courtyard space. Also the staff, as always, were excellent - in the kitchen they kindly catered to a vegan diet and remained friendly and professional throughout our stay. The rooms were alway so tidy and clean. Every piece contributes to an excellent experience. I was able to take some photographs[39]uring the event just so you can see how great it was to have such flexible spaces. Will look forward to putting together proposals for the future, and to encourage people to participate and acknowledge Dagstuhl as a key enabler for work that comes from this process.
With best regards,
m.c.

Thomas Neumann

17222 – Robust Performance in Database Query Processing | Dagstuhl Seminar | https://www.dagstuhl.de/17222

Sehr geehrte Damen und Herren,
ich wollte mich nur noch einmal für das wunderbare Seminar bedanken. Das war wirklich eine tolle Gelegenheit, mit Kollegen aus Wissenschaft und Industrie in einem Workshop zusammenzuarbeiten, und hat bei uns auch weitere Zusammenarbeit über das eigentliche Seminar hinaus ausgelöst. Eine solche Gelegenheit gibt es in unserem Bereich außerhalb von Dagstuhl eigentlich gar nicht. Und das Ambiente und die Infrastruktur von Dagstuhl waren natürlich hervorragend, und haben das Seminar sehr angenehm und produktiv gemacht.
Mit freundlichen Grüßen[40]
Thomas Neumann.

Resonanz unserer Organisatoren

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

Feedback from Organizers

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

Organizers of Dagstuhl Seminar 17142

17142 – Formal Methods of Tranformations | Dagstuhl Seminar | https://www.dagstuhl.de/17142

We warmly thank Schloss Dagstuhl for making this fruitful event possible, and for their help in the organization. It is highly appreciated as organizers, and allowed us to focus only on the scientific aspects of the seminar.

Organizers of Dagstuhl Seminar 17162

17162 – Online Privacy and Web Transparency | Dagstuhl Seminar | https://www.dagstuhl.de/17162

Finally, several new collaborations have been created as a result of this seminar, and at least one person has received a job offer due to the discussions that took place in Dagstuhl. A group of participants have organised a Slack community to exchange news and ideas in the area.

[39] https://www.flickr.com/photos/mcphoo/sets/72157685607129152

[40] engl.: Dear Sirs or Madams, I just wanted to thank you again for the wonderful Seminar. It was really a great opportunity to collaborate with colleagues from science and industry in a workshop and has started further collaborations beyond the seminar itself. Such an opportunity is, Dagstuhl aside, virtually nonexistent in our field. And, of course, the ambience and infrastructure in Dagstuhl is excellent and made the seminar very enjoyable and productive. Kind regards

Organizers of Dagstuhl Seminar 17171

17171 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/17171

No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar. We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions.

Organizers of Dagstuhl Seminar 17332

17332 – Scalable Set Visualizationss | Dagstuhl Seminar | https://www.dagstuhl.de/17332

Schloss Dagstuhl was the perfect place for hosting a seminar like this. The unique scientific atmosphere and the historic building provided not only all the room we needed for our program and the working groups, but also plenty of opportunities for continued discussions and socializing outside the official program. On behalf of all participants the organizers want to express their deep gratitude to the entire Dagstuhl staff for their outstanding support and service accompanying this seminar.

Resonanz in Sozialen Netzwerken

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

Feedback in Social Media

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

@Blouchtika

Dagstuhl Seminar | https://twitter.com/Blouchtika/status/834777746266058752

Les invitations à (Leibniz) Schloß @dagstuhl c'est un peu comme les invitations à Hogwarts (enveloppe et tout). Manque le hibou.[41]

Lucas Bechberger (University of Osnabrück, Germany)

17192 – Human-Like Neural-Symbolic Computing | Dagstuhl Seminar | http://lucas-bechberger.de/2017/05/17/dagstuhl-seminar

To summarize, I am really glad that I was able to participate in this Dagstuhl seminar: I got a better overview of neural-symbolic approaches, I was able to present my research topic and get feedback on it, I met many bright researchers with very interesting ideas, and I familiarized myself with the LTN framework which might be useful for my own future research.

Juan Sequeda (Capsenta Inc., TX, USA)

17262 – Federated Semantic Data Management | Dagstuhl Seminar | https://twitter.com/juansequeda/status/880771866364915712

Amazing week at @dagstuhl Arrived with ideas, skepticism. Work hard play hard. Left with even more ideas and different skepticism.

Mark J. Nelson (Falmouth University, UK)

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design | Dagstuhl Seminar | https://twitter.com/mjntendency/status/934088476101873665

There's also a quite good library there (albeit not a very interdisciplinary one), which I could visit at any hour of day or night, and which helped remind me that there are more interesting things to read than blog posts, news articles, and tweets...

[41] engl.: The invitations to (Leibniz) Schloß @dagstuhl are a bit like the invitations to Hogwarts (envelope and everything). Just missing the owl.

Tommy Thompson (Anglia Ruskin University, UK)

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design | Dagstuhl Seminar | https://twitter.com/GET_TUDA_CHOPPA/status/934214586202542080

Back home after my visit to the AI for game design seminar at Dagstuhl. These events always leave me with much to think about. Thanks to everyone for the insightful discussion long long (LONG) into the night.

Gillian Smith (Worcester Polytechnic Institute, MA, USA)

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design | Dagstuhl Seminar | https://twitter.com/gillianmsmith/status/934884345868845057

Just landed in Boston after an amazing, intense, and reflective week at Dagstuhl. Thanks to everyone for a wonderful time! Looking forward to the many new projects we started together.

Alex J. Champandard (AiGameDev.com KG, Austria)

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design | Dagstuhl Seminar | https://twitter.com/alexjc/status/935140590911705088

At Dagstuhl, we found hope together in ways we could never have alone. We devised solutions of the scale that we never imagined in isolation. Pieces of a big puzzle clicked together in serendipitous ways that still now gives me goose bumps...

Antonios Liapis (University of Malta, Malta)

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design | Dagstuhl Seminar | https://twitter.com/SentientDesigns/status/935573078930788356

Dagstuhl was invigorating! When you feel alone in the dark, shout and you may find that others are there in the tunnel with you. Community is important; ideas are found in discussion and willingness to see things from another perspective.

César González-Pérez (CSIC Santiago de Compostela, Spain)

17492 – Multi-Level Modelling | Dagstuhl Seminar | https://twitter.com/verdewek/status/937818356656300032

Estoy aquí www.dagstuhl.de, mirando la biblioteca de 17.000 volúmenes abierta 24 horas. Tengo que mirarla entera, ¡pero solo me quedan cuatro días![42]

Jessica Paquette (Apple Computer Inc., USA)

17502 – Testing and Verification of Compilers | Dagstuhl Seminar | https://twitter.com/barrelshifter/status/940210769471377409

dagstuhl's library is rad as heck

Ramakrishnan Durairajan (University of Oregon, OR, USA)

17511 – The Critical Internet Infrastructure Revisited | Dagstuhl Seminar | https://twitter.com/52_41_4d/status/943397489583026176

Hand-written abstracts at @dagstuhl. Two words: so cool!

Nikolaus Forgó (Leibniz University Hannover, Germany)

18262 – 10 Years of Web Science: Closing The Loop | Dagstuhl Seminar | http://lawandit.tumblr.com/post/165612332667/dagstuhl-is-one-of-the-best-academic

Dagstuhl is one of the best academic ideas/concepts I know. Participants with (very) divers backgrounds are expected to stay there for a week, live and work together at a rather remote place and distraction is reduced to a minimum. So happy to be on board again!

[42] engl.: I'm here www.dagstuhl.de, looking at the 17,000 volume library open 24 hours a day. I have to look at it all, but I only have four days left!

Resonanz im Fragebogen

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

Survey Feedback

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

17041 – Randomization in Parameterized Complexity | Dagstuhl Seminar | https://www.dagstuhl.de/17041

The staff was GREAT across the board. Make sure to keep them.

17072 – Applications of Topology to the Analysis of 1-Dimensional Objects | Dagstuhl Seminar | https://www.dagstuhl.de/17072

Altogether this is a great and inspiring place with wonderful people running it!

17101 – Databases on Future Hardware | Dagstuhl Seminar | https://www.dagstuhl.de/17101

This was my first Dagstuhl seminar and I thoroughly enjoyed the experience. While it is quite difficult for me to spend a week away from my main job, I think the investment in time was more than worth it. If another seminar with a similar degree of relevance to my work were to come along I would certainly make an effort to come.

17161 – Ambient Notification Environments | Dagstuhl Seminar | https://www.dagstuhl.de/17161

Dagstuhl is one of the most remarkable things we have in Computer Science! Long may it last.

17162 – Online Privacy and Web Transparency | Dagstuhl Seminar | https://www.dagstuhl.de/17162

Thank you for the wonderful venue and hospitality. Dagstuhl flls an important need for academia to exchange ideas and to collaborate.

17192 – Human-Like Neural-Symbolic Computing | Dagstuhl Seminar | https://www.dagstuhl.de/17192

Dagstuhl is ahead of other venues as it is extremely well-managed and has kept the high standards of its seminars!

17202 – Challenges and Opportunities of User-Level File Systems for HPC | Dagstuhl Seminar | https://www.dagstuhl.de/17202

Dagstuhl is a great venue – it genuinely provides an environment more conducive to collaboration and creativity than any other workshop I have attended.

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

It is really a step beyond regular conferences for getting inspired.

17281 – Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition | Dagstuhl Seminar | https://www.dagstuhl.de/17281

Dagstuhl provides a unique opportunity for in depth interaction with peers in the field.

17332 – Scalable Set Visualizations | Dagstuhl Seminar | https://www.dagstuhl.de/17332

Thanks for hosting this seminar. As with my previous experiences, Dagstuhl provided an opportunity for some of the best scientists to focus on compelling research questions. It is a unique venue in this regard and is a massive asset to German Computer Science Research in particular and to international researchers generally.

17381 – Recent Trends in Knowledge Compilation | Dagstuhl Seminar | https://www.dagstuhl.de/17381

Great idea to make the lower room in new building into another seminar room!

17392 – Body-Centric Computing | Dagstuhl Seminar | https://www.dagstuhl.de/17392

It was fantastic how nicely nicely my family was welcomed!

17421 – Computational Proteomics | Dagstuhl Seminar | https://www.dagstuhl.de/17421

I believe this is a unique and exceptional meeting location and conference. I will most certainly recommend to any and all colleagues that are invited to attend.

17461 – Connecting Visualization and Data Management Research | Dagstuhl Seminar | https://www.dagstuhl.de/17461

Great seminar. I am especially grateful for the NSF travel grant. This was an immense help.

17481 – Reliable Computation and Complexity on the Reals | Dagstuhl Seminar | https://www.dagstuhl.de/17481

Please keep up the funding for this fantastic place!

17491 – Computational Metabolomics: Identification, Interpretation, Imaging | Dagstuhl Seminar | https://www.dagstuhl.de/17491

It is a remarkable facility and a real treasure. I am honoured to be able to attend these seminars.

17492 – Multi-Level Modelling | Dagstuhl Seminar | https://www.dagstuhl.de/17492

Dagstuhl is such an outstanding institution. Just continue on your path. You are doing a tremendous service to the community.

17022 – Automated Program Repair | Dagstuhl Seminar | https://www.dagstuhl.de/17022

Dagstuhl is what academia should be. Please don't change.

17031 – Planning and Robotics | Dagstuhl Seminar | https://www.dagstuhl.de/17031

Dagstuhl supports "old style" discourse-oriented science. That is great! Keep it!

17162 – Online Privacy and Web Transparency | Dagstuhl Seminar | https://www.dagstuhl.de/17162

Dagstuhl reminds me a rarity among places of intellectual retreats – please do not change it in any bad way. It might be interesting to broaden the fields a bit.

17351 – Machine Learning and Formal Methods | Dagstuhl Seminar | https://www.dagstuhl.de/17351

Dagstuhl is an oasis in the harsh landscape of life. Please preserve its unique beauty, and I am very appreciate towards all that has been done so far to keep it as it is.

17351 – Machine Learning and Formal Methods | Dagstuhl Seminar | https://www.dagstuhl.de/17351

No. The Dagstuhl experience was always extraordinary, but it has improved even further in the last few years.

17372 – Cybersafety in Modern Online Social Networks | Dagstuhl Seminar | https://www.dagstuhl.de/17372

Dagstuhl remains a one of a kind and key institute for computer scientists and their collaborators.

17372 – Cybersafety in Modern Online Social Networks | Dagstuhl Seminar | https://www.dagstuhl.de/17372

This was my first visit to Schloss Dagstuhl. I had heard about others' experiences before but the attendance surpassed my expectations. This is a unique venue and facility. Please resist any efforts to change it.

17371 – Deduction Beyond First-Order Logic | Dagstuhl Seminar | https://www.dagstuhl.de/17371

None. Dagstuhl is still the gold standard for in-depth seminars.

17111 – Game Theory in AI, Logic, and Algorithms | Dagstuhl Seminar | https://www.dagstuhl.de/17111

The integration of art into Dagstuhl is a wonderful trend that must continue!!! Perhaps even encourage scientific participants to actively engage in doing art, in addition to observing it.

17272 – Citizen Science: Design and Engagement | Dagstuhl Seminar | https://www.dagstuhl.de/17272

Continue to partner for NSF support of junior scholars traveling from the US. Sadly I was not eligible by a matter of just a few months; it would be preferable if eligibility criterion was "pre tenure" for academics.

17361 – Finite and Algorithmic Model Theory | Dagstuhl Seminar | https://www.dagstuhl.de/17361

Schloss Dagstuhl is a unique place. Its value for computer science is priceless.

17452 – Algorithmic Cheminformatics | Dagstuhl Seminar | https://www.dagstuhl.de/17452

My above comments are all about things you should definitely retain. You are creating an atmosphere similar to that of a college campus, or one of the classical science meetings in Europe in the earlier 20th century. The directness and project-focus of those formats is being lost in many placed. Please do not lose it here.

17121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | https://www.dagstuhl.de/17121

The fact that the place is isolated and that people spend their spare time in the cafetria is what separates this from other conferences/workshops as participants talk much more with one another in a non-formal way, which is great.

17171 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/17171

I got to know a lot of people that I only know from papers before.

17191 – Theory of Randomized Optimization Heuristics | Dagstuhl Seminar | https://www.dagstuhl.de/17191

Dagstuhl seminars are always a pleasure. They provide for interactions among participants that cannot be had in this form anywhere else.

17191 – Theory of Randomized Optimization Heuristics | Dagstuhl Seminar | https://www.dagstuhl.de/17191

The best aspect is the uniquely creative atmosphere produced by both the outstanding scientific content and the brilliant organization of the seminar. The seminar effectively promoted my ideas on future research directions, I was able to look at my research from a different perspective and place it in a broad picture of the current state of the art. The seminar inspired me so much that it is hard to think of any bad aspects, despite that closer to the end I felt somewhat tired and was not able to be as active as at the beginning.

17202 – Challenges and Opportunities of User-Level File Systems for HPC | Dagstuhl Seminar | https://www.dagstuhl.de/17202

I think the best aspect was the relatively small number of participants, since this allows proper discussion of issues, while still having a large enough group to provide different viewpoints on the topics being discussed. The facility is very nice, and free from distractions that would otherwise cause participants to disperse after the scheduled talks. We've had several interesting discussions in the evening after dinner, which rarely happens at large conferences.

17202 – Challenges and Opportunities of User-Level File Systems for HPC | Dagstuhl Seminar | https://www.dagstuhl.de/17202

Met many new collaborators. Only knew 1 person at the start.

17281 – Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition | Dagstuhl Seminar | https://www.dagstuhl.de/17281

I got to know other people with similar interests much better than I would have at a "normal conference".

17371 – Deduction Beyond First-Order Logic | Dagstuhl Seminar | https://www.dagstuhl.de/17371

Everything was excellent, the most useful were informal evening and coffee break discussions.

17431 – Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions | Dagstuhl Seminar | https://www.dagstuhl.de/17431

The seminar brought people from many countries together in one place. It was interesting to learn about all the different research directions at numerous institutions and labs. The seminar was also quite informal and there was plenty of time devoted to discussions which were great. The randomized lunch/dinner placements that forced participants to interact was one of my favorite things about this seminar. Overall, I had a great time and would happily visit again.

17092 – Search as Learning | Dagstuhl Seminar | https://www.dagstuhl.de/17092

I'm really impressed by the facilities here.

17352 – Analysis and Synthesis of Floating-point Programs | Dagstuhl Seminar | https://www.dagstuhl.de/17352

New meeting room on ground floor of new wing is very nice. A great addition!

17372 – Cybersafety in Modern Online Social Networks | Dagstuhl Seminar | https://www.dagstuhl.de/17372

The atmosphere is so conducive to research and it is so impressive to be able to meet colleagues from all over the world in this setting.

17461 – Connecting Visualization and Data Management Research | Dagstuhl Seminar | https://www.dagstuhl.de/17461

This is an awesome venue. There is nothing like this in USA.

17452 – Algorithmic Cheminformatics | Dagstuhl Seminar | https://www.dagstuhl.de/17452

Cleanliness and order are outstanding. I am happy that the facilities are simple and direct..

17491 – Computational Metabolomics: Identification, Interpretation, Imaging | Dagstuhl Seminar | https://www.dagstuhl.de/17491

Super friendly and helpful staff! Everything was really well organised and efficient.

Resonanz zur Bibliographiedatenbank dblp

5.2 *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.

Rainer Unland (Poznan University of Economics and Business, Poland)
dblp | https://dblp.org

Ich muss wirklich gestehen, dass ich mehr als beeindruckt bin, wie vorbildlich, umfangreich und angenehm zu lesen DBLP ist. Ich finde hier mehr über mich heraus, als ich in meinem CV stehen habe - unglaublich. Noch einmal Herzlichen Dank und großes Lob für diese Arbeiten![43]

Gianmaria Silvello (University of Padua, Italy)
dblp | https://twitter.com/giansilv/status/877280435054800903

@dblp_org guys at #JCDL2017 . To me the best computer science biblio db out there. Add citations and conquer the world! ;)

Ashish Sureka (Ashoka University, India)
dblp | https://twitter.com/ashish_sureka/status/910531063977361409

A very nice feature - My ORCID profile is already integrated with my DBLP profile - great work by @dblp_org & @ORCID_Org

Guillaume Cabanac (University of Toulouse, France)
dblp | https://twitter.com/gcabanac/status/914384731231473665

Thank you @dblp_org for indexing French «habilitations à diriger des recherches» (HDR) in Comp. Sci. from @hal_fr

Theodore Vasiloudis (Swedish Institute for Computer Science)
dblp | https://twitter.com/thvasilo/status/918776219432341505

Impressive that @dblp_org remains the best search engine for CS(+)literature. Better filtering, faster, easier to dig through.

Johannes Hölzl (Carnegie Mellon University, USA)
dblp | https://twitter.com/johannes2007/status/926016720636129280

At least for me, @dblp_org is *the* site to find out what an author published (besides homepage). Google Scholar is not enough for me.

Alexander Serebrenik (Eindhoven University of Technology, The Netherlands)
dblp | https://twitter.com/aserebrenik/status/933228583522254848

DBLP tells that this year I've published my 100 conference paper... @dblp_org thank you for tracking our publications!

Andreas Halkjær From (Technical University of Denmark)
dblp | https://twitter.com/andreasfrom/status/971447624455901185

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

[43] engl.: I really have to admit that I am more than impressed how exemplary, comprehensive, and enjoyable to read DBLP is. I find out more about myself than I have in my CV - unbelievable. Once again, many thanks and praise for this work!

Resonanz zu Dagstuhl Publishing

5.3

Feedback on Dagstuhl Publishing

Im Prozess der Veröffentlichung von Konferenz-Proceedings, Zeitschriften-Artikeln und Büchern stehen wir in engem Kontakt mit den Herausgebern und Autoren. Rückmeldungen zu unseren Veröffentlichungsangeboten erhalten wir aber auch im Rahmen unserer regelmäßigen Befragungen der Teilnehmer von Dagstuhl-Seminaren oder Dagstuhl-Perspektiven-Workshops.

We are in close contact with editors and authors as part of the publishing procedures for conference proceedings, journal articles, and books. Additionally, we receive feedback regarding our publishing services in the questionnaires filled out by the participants of Dagstuhl Seminars or Dagstuhl Perspectives Workshops.

17142 – Formal Methods of Transformations | Dagstuhl Seminar | https://www.dagstuhl.de/17142

Please continue to support LIPIcs!

17361 – Finite and Algorithmic Model Theory | Dagstuhl Seminar | https://www.dagstuhl.de/17361

Scientific publication changes (see library announcement on cancellation of Elsevier journals). Dagstuhl should continue to play an active role (as with LIPIcs).

17401 – Quantum Cryptanalysis | Dagstuhl Seminar | https://www.dagstuhl.de/17401

Great that you took a stance for open access

17401 – Quantum Cryptanalysis | Dagstuhl Seminar | https://www.dagstuhl.de/17401

Thank you Dagstuhl for having expanded your own publication efforts, all with Open Access. This is such a valuable community resource.

17401 – Quantum Cryptanalysis | Dagstuhl Seminar | https://www.dagstuhl.de/17401

Keep and expand your Open Access publications.

6 Die Seminare in 2017
The 2017 Seminars

Applications, Interdisciplinary Work

- Algorithmic Cheminformatics (17452)
- Ambient Notification Environments (17161)
- Body-Centric Computing (17392)
- Citizen Science: Design and Engagement (17272)
- Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction (17252)
- Computational Interactivity (17232)
- Computational Metabolomics: Identification, Interpretation, Imaging (17491)
- Computational Proteomics (17421)
- Computer Science Meets Ecology (17091)
- Connecting Visualization and Data Management Research (17461)
- Hyperspectral, Multispectral, and Multimodal (HMM) Imaging: Acquisition, Algorithms, and Applications (17411)
- Internet of People (17412)
- Machine Learning and Formal Methods (17351)
- Shape-Changing Interfaces (17082)
- Search as Learning (17092)
- User-Generated Content in Social Media (17301)

Artificial Intelligence, Computational Linguistics

- Approaches and Applications of Inductive Programming (17382)
- Artificial and Computational Intelligence in Games: AI-Driven Game Design (17471)
- Computer-Assisted Engineering for Robotics and Autonomous Systems (17071)
- Epistemic Planning (17231)
- Planning and Robotics (17031)
- Recent Trends in Knowledge Compilation (17381)

Cryptography, Security, Privacy

- Cybersafety in Modern Online Social Networks (17372)
- Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition (17281)
- Online Privacy and Web Transparency (17162)
- Opportunities and Risks of Blockchain Technologies (17132)
- Quantum Cryptanalysis (17401)

Databases, Information Retrieval, Machine Learning, Data Mining

- Big Stream Processing Systems (17441)
- Challenges and Opportunities of User-Level File Systems for HPC (17202)
- Databases on Future Hardware (17101)
- Deep Learning for Computer Vision (17391)
- Federated Semantic Data Management (17262)
- From Characters to Understanding Natural Language (C2NLU): Robust End-to-End Deep Learning for NLP (17042)
- Human-Like Neural-Symbolic Computing (17192)
- Robust Performance in Database Query Processing (17222)
- Towards Cross-Domain Performance Modeling and Prediction: IR/RecSys/NLP (17442)

Data Structures, Algorithms, Complexity

- Addressing the Computational Challenges of Personalized Medicine (17472)
- Applications of Topology to the Analysis of 1-Dimensional Objects (17072)
- Computability Theory (17081)
- Computational Counting (17341)
- Computational Geometry (17171)
- Finite and Algorithmic Model Theory (17361)
- From Observations to Prediction of Movement (17282)
- Game Theory in AI, Logic, and Algorithms (17111)
- Game Theory Meets Computational Learning Theory (17251)
- Geometric Modelling, Interoperability and New Challenges (17221)
- Randomization in Parameterized Complexity (17041)
- Reliable Computation and Complexity on the Reals (17481)
- Probabilistic Methods in the Design and Analysis of Algorithms (17141)
- Theory and Applications of Hashing (17181)
- Theory of Randomized Optimization Heuristics (17191)
- Topology, Computation and Data Analysis (17292)
- Voting: Beyond Simple Majorities and Single-Winner Elections (17261)

Distributed Computation, Networks, Architecture, Systems

- Computational Complexity of Discrete Problems (17121)
- Foundations of Wireless Networking (17271)
- Mixed Criticality on Multicore / Manycore Platforms (17131)
- Network Function Virtualization in Software Defined Infrastructures (17032)
- The Critical Internet Infrastructure Revisited (17511)
- Using Networks to Teach About Networks (17112)
- Wildly Heterogeneous Post-CMOS Technologies Meet Software (17061)

Geometry, Image Processing, Graphics, Visualization

- Beyond VR and AR: Reimagining Experience Sharing and Skill Transfer Towards an Internet of Abilities (17062)
- Functoriality in Geometric Data (17021)
- Scalable Set Visualizations (17332)

Software Technology, Programming Languages

- Multi-Level Modelling (17492)
- New Challenges in Parallelism (17451)
- Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions (17431)
- Rethinking Productivity in Software Engineering (17102)
- SLEBOK: The Software Language Engineering Body of Knowledge (17342)
- Testing and Verification of Compilers (17502)

Verification, Logic, Formal Methods, Semantics

- Analysis and Synthesis of Floating-point Programs (17352)
- A Shared Challenge in Behavioural Specification (17462)
- Automated Program Repair (17022)
- Deduction Beyond First-Order Logic (17371)
- Formal Methods of Transformations (17142)
- Formal Synthesis of Cyber-Physical Systems (17201)
- Resource Bound Analysis (17291)
- Theory and Applications of Behavioural Types (17051)

6.1 Functoriality in Geometric Data

Organizers: Mirela Ben-Chen, Frédéric Chazal, Leonidas J. Guibas, and Maks Ovsjanikov

Seminar No. 17021

Date: January 8–13, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.1

Participants: Dror Aiger, Annamaria Amenta, Mirela Ben-Chen, Benjamin Berkels, Alex M. Bronstein, Michael M. Bronstein, Antonin Chambolle, Frédéric Chazal, Marco Cuturi, Jie Gao, Franck Hétroy-Wheeler, Klaus Hildebrandt, Vladimir G. Kim, Ron Kimmel, Zorah Lähner, Wilmot Li, Or Litany, Mauro Maggioni, Konstantin Mischaikow, Maks Ovsjanikov, Helmut Pottmann, Emanuele Rodolà, Martin Rumpf, Frank R. Schmidt, Ariel Shamir, Primoz Skraba, William Smith, Justin Solomon, Ronen Talmon, Boris Thibert, Matthias Vestner, Michael Wand, Benedikt Wirth, Hao Zhang

Across science, engineering, medicine and business we face a deluge of data coming from sensors, from simulations, or from the activities of myriads of individuals on the Internet. The data often has a geometric character, as is the case with 1D GPS traces, 2D images, 3D scans, and so on. Furthermore, the data sets we collect are frequently highly correlated, reflecting information about the same or similar entities in the world, or echoing semantically important repetitions/symmetries or hierarchical structures common to both man-made and natural objects.

A recent trend, emerging independently in multiple theoretical and applied communities is to understand geometric data sets through their relations and interconnections, a point of view that can be broadly described as exploiting the *functoriality* of data, which has a long tradition associated with it in mathematics. Functoriality, in its broadest form, is the notion that in dealing with any kind of mathematical object, it is at least as important to understand the transformations or symmetries possessed by the object or the family of objects to which it belongs, as it is to study the object itself. This general idea been successfully applied in a large variety of fields, both theoretical and practical, often leading to deep insights into the structure of various objects as well as to elegant and efficient methods in various application domains, including computational geometry, computer vision and computer graphics.

This seminar brought together researchers and practitioners interested in notions of *similarity*, *correspondence* and, more generally, *relations* across geometric data sets. Mathematical and computational tools for the construction, analysis, and exploitation of such relational networks were the central focus of this seminar.

6.2 Automated Program Repair

Organizers: Sunghun Kim, Claire Le Goues, Michael Pradel, and Abhik Roychoudhury

Seminar No. 17022

Date: January 8–13, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.19

Participants: Yuriy Brun, Celso G. Camilo-Junior, Jonathan Dorn, Lars Grunske, Ciera Jaspan, Aditya Kanade, Sarfraz Khurshid, Dongsun Kim, Sunghun Kim, Julia Lawall, Claire Le Goues, Fan Long, Matías Sebastían Martínez, Sergey Mechtaev, Martin Monperrus, Adrian Nistor, Alessandro Orso, Justyna Petke, Michael Pradel, Mukul Prasad, Abhik Roychoudhury, Marija Selakovic, Kathryn T. Stolee, David R. White, Yingfei Xiong, Andreas Zeller, Tianyi Zhang

Software engineering targets the creation of software for myriad platforms, deployed over the internet, the cloud, mobile devices and conventional desktops. Software now controls cyber-physical systems, industrial control systems, and "Internet of Things" devices, and is directly responsible for humanity's economic well-being and safety in numerous contexts. It is therefore especially important that engineers are able to easily write error-free software, and to quickly find and correct errors that do appear. Future generation programming environments must not only employ sophisticated strategies for localizing software errors, but also strategies for automatically patching them.

Recent years have seen an explosive growth in research on automated program repair, with proposed techniques ranging from pure stochastic search to pure semantic analysis. The Dagstuhl Seminar in January 2017 studies the problem of automated repair in a holistic fashion. This will involve a review of foundational techniques supporting program repair, perspectives on current challenges and future techniques, and emerging applications. The aim is to broadly discuss and revisit underlying assumptions and methods towards the integration of automated patch synthesis into futuristic programming environments.

Conceptually, applications of program repair step far beyond the general goal of ensuring software quality, and the subject is relevant to a broad range of research areas. It is of obvious importance in software testing and analysis, because repair goes hand in hand with traditional testing and debugging activities. It is relevant to researchers in programming languages and systems, e.g., to study language and system-level techniques that integrate patch suggestions during development. The topic is relevant to researchers in systems security, as repair approaches may be customizable to patching vulnerabilities in both application and systems software. Researchers in formal methods may provide insight for provably correct repair, given appropriate correctness specifications. Finally, the topic is connected to human computer interaction in software engineering, since the repair process, if not fully automated, may involve eliciting appropriate developer feedback and acting on it accordingly.

At a technical level, one of the key discussion topics has been the correctness specifications driving the repair process. Most previous work in this domain has relied on test suites as partial correctness specifications. While often available, test suites are typically inadequate for fully assessing patch correctness. Alternative quality specifications, such as "minimality", could be explored. In addition, heavier-weight specifications, such as assertions, may provide stronger functional guarantees, leaving open the research challenge both in how to use them and how they may be derived to guide a repair process. Given the appropriate correctness specification, the task of repair usually involves three steps: localizing an error to a small set of potentially-faulty locations, deriving values/constraints for the computation desired at the faulty locations, and constructing "fix" expressions/statements that satisfy these values/constraints. Each of the three steps can be accomplished by a variety of methods, including heuristic search, symbolic analysis and/or constraint solving techniques. This allows for an interesting interplay for an entire design space of repair techniques involving ingenuous combinations of search-based techniques and semantic analysis being employed at the different steps of the repair process.

The Dagstuhl Seminar has attracted researchers and practitioners from all over the world, comprising participants active in the fields of software engineering, programming languages, machine learning, formal methods, and security. As a result of the seminar, several participants plan to launch various follow-up activities, such as a program repair competition, which would help to further establish and guide this young field of research, and a journal article that summarizes the state of the art in automated program repair.

6.3 Planning and Robotics

Organizers: Malik Ghallab, Nick Hawes, Daniele Magazzeni, Brian C. Williams, and Andrea Orlandini

Seminar No. 17031

Date: January 15–20, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.32

Participants: Rachid Alami, Iman Awaad, Roman Bartak, Michael Beetz, Ronen I. Brafman, Michael Cashmore, Martin Davies, Minh Do, Susan L. Epstein, Alberto Finzi, Hector Geffner, Malik Ghallab, Nick Hawes, Malte Helmert, Andreas Hertle, Joachim Hertzberg, Laura M. Hiatt, Michael W. Hofbaur, Jörg Hoffmann, Felix Ingrand, Luca Iocchi, Gal A. Kaminka, Erez Karpas, Oliver E. Kim, Sven Koenig, Lars Kunze, Bruno Lacerda, Gerhard Lakemeyer, Daniele Magazzeni, Lenka Mudrova, Daniele Nardi, Tim Niemüller, Andrea Orlandini, Amit Kumar Pandey, Simon Parsons, Ron Petrick, Mark Roberts, Enrico Scala, Matthias Scheutz, Reid Simmons, Elizabeth Sklar, Stephen Smith, Siddharth Srivastava, Manuela Veloso, Brian C. Williams

Automated Planning and Scheduling (P&S) and Robotics were strongly connected in the early days of A.I., but became mostly disconnected later on. Indeed, Robotics is one of the most appealing and natural application area for the P&S research community, however such a natural interest seems to not be reflected by advances beyond the state-of-the-art in P&S research in Robotics applications. In light of the accelerated progress and the growth of economic importance of advanced robotics technology, it is essential for the P&S community to respond to the challenges that these applications pose and contribute to the advance of intelligent robotics.

In this perspective, a Planning and Robotics (PlanRob) initiative within the P&S research community has been recently started with a twofold aim. On the one hand, this initiative would constitute a fresh impulse for the P&S community to develop its interests and efforts towards the Robotics research area. On the other hand, it aims at attracting representatives from the Robotics community to discuss their challenges related to planning for autonomous robots (deliberative, reactive, continuous planning and execution etc.) as well as their expectations from the P&S community. The PlanRob initiative was initiated as a workshop series (http://pst.istc.cnr.it/planrob/) started at the International Conference on Automated Planning and Scheduling (ICAPS) in 2013. The PlanRob workshop editions gathered very good feedback from both the P&S and Robotics communities. And this resulted also in the organisation of a specific Robotics Track at ICAPS since 2014.

The aim of this Dagstuhl Seminar was to reinforce such initiative and increase the synergy between these two research communities. Then, most of the attendees contributed with position statements (whose abstracts are available in this report) to present their major challenges and approaches for addressing them. In general, this involved sharing views, thoughts and contributions across the following main topics:

- **Long-term autonomy / Open world planning**, providing an overview on issues related to continuous planning for robots with partial information or even incomplete models;
- **Knowledge Representation and Reasoning in Planning**, with presentations on cognitive features and robot planning;
- **Challenges in Industrial, Logistics & Consumer Robotics**, providing relevant insights related to deployment of robots in real world scenarios;
- **Human-Robot Planning**, with a wide overview on planning solutions for dealing with interactions between humans and robots;
- **Planning and Execution**, discussing issues and challenges related to robust planning and execution for robot control;
- **Task & Motion Planning / Hybrid planners**, with presentations on integrated solutions for robot control at different levels;
- **Reliable and Safe Planning for Robotics**, providing an overview of ISO standards for robots and, more in general, investigating the exploitation of formal methods to guarantee reliability in robotic applications;
- **Technological Issues in Robot planning/Multi-robot Planning**, with statements on technological issues in (multi-) robot solutions.

Each session was animated by (i) an opponent, whose role was to be critical about the position statements and (ii) a moderator, to organise the discussion. Therefore, opponents and moderators have provided a short summary of the session ideas and discussion in dedicated Synthesis Sessions to further foster the discussion.

In addition, two panel sessions have been organised on (i) **Evaluation, Benchmarking and Competitions**, discussing the experience in RoboCup@Home and the organisation of the new Planning and Execution competition (that will be held in 2017), and (ii) **Outreach & Training**, discussing about the

possible organisation of summer schools and the opening of new scientific networking initiatives (e.g., a COST action).

During the seminar, discussions focused on different issues, challenges, possible solutions and new promising trends over a very wide variety of relevant topics: knowledge representation, modelling issues, the need of incomplete models; cognitive features such as, for instance, learning and goal reasoning; human-aware solutions for flexible human-robot interaction; adaptive solutions for human-robot collaboration; robust execution capable of effectively dealing with failure; integration issues in robotic architecture that, e.g., exploit different kind of models and then perform hybrid reasoning; application of formal methods to provide verification and validation functionalities to guarantee reliable robotic systems; etc. Indeed, addressing the integration of P&S and Robotics for development of intelligent robots entails covering a heterogeneous spectrum of problems, often requiring complex solutions that require a vast set of knowledge and technologies.

During the seminar, there was a very high level of engagement and interaction between the participants, enabling a lively and productive week. The main outcome of the seminar was to share a common understanding of issues and solutions with thorough discussions. And the workshop ended with an open discussion on possible follow ups and possible actions to create further opportunities for fostering synergies and interactions between the two communities.

6.4 Network Function Virtualization in Software Defined Infrastructures

Organizers: David Hausheer, Oliver Hohlfeld, Diego R. López, Bruce MacDowell Maggs, and Costin Raiciu

Seminar No. 17032

Date: January 15–18, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.74

Participants: Theo Benson, Christian Esteve Rothenberg, Aaron Gember-Jacobson, Gabriela Gheorghe, David Hausheer, Oliver Hohlfeld, Felipe Huici, Andreas Kassler, Wolfgang Kellerer, James Kempf, Dirk Kutscher, Diego R. Lopez, Bruce MacDowell Maggs, Oliver Michel, Leonhard Nobach, Costin Raiciu, Fernando M. V. Ramos, Jan Rüth, Michael Scharf, Fabian Schneider, Laurent Vanbever, Timothy Wood, Andreas Wundsam, Thomas Zinner

Network management currently undergoes massive changes towards realizing more flexible management of complex networks. Recent efforts include slicing data plane resources by using network (link) virtualization and applying operating system design principles in Software Defined Networking (SDN). Driven by network operators, network management principles are currently envisioned to be even further improved by virtualizing network functions which are currently realized in dedicated hardware appliances. The resulting Network Function Virtualization (NFV) paradigm abstracts network functions from dedicated hardware to virtual machines running on commodity hardware and enables a Cloud-like network management. All of these efforts contribute to a softwarization of communication networks. This softwarization represents a significant change to network design and management by allowing the application of operating system design and software engineering principles to make network management more efficient, e.g., by enabling flexible and dynamic service provisioning.

Since the NFV efforts are currently mainly driven by carriers and standardization bodies, academic research is decoupled from the industry driven NFV attempts in redesigning network management. Due to this missing link to academic research, opportunities for groundbreaking research and broad impact in academia are currently missing out. This Dagstuhl Seminar thus gathered researchers from academia, industry, and standardization bodies to establish this missing link by fostering collaborations and joint research initiatives. Thus, a particular focus of the seminar was on identifying the diverse connections between industry driven NFV efforts and current academic networking research.

The seminar brought together 24 participants in January 2017 to discuss a potential NFV research agenda within 2.5 days. The program included different invited talks that provided an overview on selected aspects of NFV and lightning talks by each participants to provide first research questions and to sketch research directions. We summarize each talk in the remainder of this report. The main focus of the seminar was then the in-depth discussion the research areas identified in the lightning talks in several breakout sessions, which we also summarize. We closed the seminar by collecting and discussing several opinions from each participant: (i) lessons learned and surprises on NFV during the seminar and (ii) open research questions. We further collected controversial statements on NFV research and asked the seminar participants on whether they agree or disagree to each presented statement. We summarize the outcomes at the end of the report.

6.5 Randomization in Parameterized Complexity

Organizers: Marek Cygan, Fedor V. Fomin, Danny Hermelin, and Magnus Wahlström

Seminar No. 17041

Date: January 22–27, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.103

Participants: Amir Abboud, Arturs Backurs, Andreas Björklund, Édouard Bonnet, Cornelius Brand, Karl Bringmann, Yixin Cao, Radu Curticapean, Marek Cygan, Holger Dell, Fedor V. Fomin, Tobias Friedrich, Petr A. Golovach, Gregory Z. Gutin, Danny Hermelin, Petr Hlineny, Petteri Kaski, Eun Jung Kim, Yiannis Koutis, Łukasz Kowalik, Stefan Kratsch, Bingkai Lin, Daniel Lokshtanov, Dániel Marx, Jesper Nederlof, Fahad Panolan, Christophe Paul, Geevarghese Philip, Marcin Pilipczuk, Michal Pilipczuk, M. S. Ramanujan, Peter Rossmanith, Marc Roth, Saket Saurabh, Ildiko Schlotter, Stefan Szeider, Dimitrios M. Thilikos, Magnus Wahlström, Gerhard J. Woeginger, David P. Woodruff, Meirav Zehavi

Randomization plays a prominent role in many subfields of theoretical computer science. Typically, this role is twofold: On the one hand, randomized methods can be used to solve essentially classical problems easier or more efficiently. In many cases, this allows for simpler, faster, and more appealing solutions for problems that have rather elaborate deterministic algorithms; in other cases, randomization provides the only known way to cope with the problem (e.g. polynomial identity testing, or deciding whether there exists a perfect matching with exactly b red edges in an edge-colored bipartite graph). On the other hand, there are also cases where randomness is intrinsic to the question being asked, such as the study of properties of random objects, and the search for algorithms which are efficient on average for various input distributions.

Parameterized complexity is an approach of handling computational intractability, where the main idea is to analyze the complexity of problems in finer detail by considering additional problem parameters beyond the input size. This area has enjoyed much success in recent years, and has yielded several new algorithmic approaches that help us tackle computationally challenging problems. While randomization already has an important role in parameterized complexity, for instance in techniques such as color-coding or randomized contractions, there is a common opinion within researchers of the field that the full potential of randomization has yet to be fully tapped.

The goal of this seminar was to help bridge this gap, by bringing together experts in the areas of randomized algorithms and parameterized complexity. In doing so, we hope to:

- Establish domains for simpler and/or more efficient FPT algorithms via randomization.
- Identify problems which intrinsically need randomization.
- Study parameterized problems whose instances are generated by some underlying distribution.
- Stimulate the development of a broadened role of randomness within parameterized complexity.

6.6 From Characters to Understanding Natural Language (C2NLU): Robust End-to-End Deep Learning for NLP

Organizers: Phil Blunsom, Kyunghyun Cho, Chris Dyer, and Hinrich Schütze

Seminar No. 17042

Date: January 22–27, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.129

Participants: Heike Adel, Parnia Bahar, Phil Blunsom, Ondřej Bojar, Fabienne Cap, Ryan Cotterell, Vera Demberg, Kevin Duh, Chris Dyer, Desmond Elliott, Manaal Faruqui, Orhan Firat, Alexander M. Fraser, Vladimir Golkov, Jan Hajič, Georg Heigold, Karl Moritz Hermann, Thomas Hofmann, Hang Li, Adam Lopez, Marie-Francine Moens, Hermann Ney, Jan Niehues, Laura Rimell, Helmut Schmid, Martin Schmitt, Hinrich Schütze, Kristina Toutanova, Ngoc Thang Vu, Yadollah Yaghoobzadeh, François Yvon

Deep learning is currently one of most active areas of research in machine learning and its applications, including natural language processing (NLP). One hallmark of deep learning is *end-to-end learning*: all parameters of a deep learning model are optimized *directly for the learning objective*; e.g., for the objective of accuracy on the binary classification task: is the input image the image of a cat? Crucially, the set of parameters that are optimized includes "first-layer" parameters that connect the raw input representation (e.g., pixels) to the first layer of internal representations of the network (e.g., edge detectors). In contrast, many other machine learning models employ hand-engineered features to take the role of these first-layer parameters.

Even though deep learning has had a number of successes in NLP, research on true end-to-end learning is just beginning to emerge. Most NLP deep learning models still start with a hand-engineered layer of representation, the level of tokens or words, i.e., the input is broken up into units by manually designed tokenization rules. Such rules often fail to capture structure both within tokens (e.g., morphology) and across multiple tokens (e.g., multi-word expressions). Given the success of end-to-end learning in other domains, it is likely that it will also be widely used in NLP to alleviate these issues and lead to great advances.

The seminar brought together researchers from deep learning, general machine learning, natural language processing and computational linguistics to develop a research agenda for the coming years. The goal was to combine recent *advances in deep learning architectures and algorithms* with *extensive domain knowledge about language* to make *true end-to-end learning for NLP* possible.

Our goals were to make progress on answering the following research questions.

- C2NLU approaches so far fall short of the state of the art in cases where token structures can easily be exploited (e.g., in well-edited newspaper text) compared to word-level approaches. What are promising avenues for developing C2NLU to match the state of the art even in these cases of text with well-defined token structures?
- Character-level models are computationally more expensive than word-level models because detecting syntactic and semantic relationships at the character-level is more expensive (even though it is potentially more robust) than at the word-level. How can we address the resulting challenges in scalability for character-level models?
- Part of the mantra of deep learning is that domain expertise is no longer necessary. Is this really true or is knowledge about the fundamental properties of language necessary for C2NLU? Even if that expertise is not needed for feature engineering, is it needed to design model architectures, tasks and training regimes?
- NLP tasks are diverse, ranging from part-of-speech tagging over sentiment analysis to question answering. For which of these problems is C2NLU a promising approach, for which not?
- More generally, what characteristics make an NLP problem amenable to be addressed using tokenization-based approaches vs. C2NLU approaches?
- What specifically can each of the two communities involved – natural language processing and deep learning – contribute to C2NLU?
- Create an NLP/deep learning roadmap for research in C2NLU over the next 5–10 years.

6.7 Theory and Applications of Behavioural Types

Organizers: Simon Gay, Vasco T. Vasconcelos, Philip Wadler, and Nobuko Yoshida

Seminar No. 17051

Date: January 29 to February 3, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.1.158

Participants: Gul Agha, Nada Amin, Robert Atkey, Giovanni Tito Bernardi, Laura Bocchi, Edwin Brady, Luis Caires, Marco Carbone, Ilaria Castellani, Tzu-chun Chen, Mariangiola Dezani, Patrick Thomas Eugster, Adrian Francalanza, Ronald Garcia, Simon Gay, Philipp Haller, Thomas Hildebrandt, Hans Hüttel, Keigo Imai, Dimitrios Kouzapas, Roland Kuhn, Ivan Lanese, Hugo-Andrés López, Francisco Martins, Conor McBride, Hernán Melgratti, Fabrizio Montesi, J. Garrett Morris, Nicholas Ng, Dominic Orchard, Luca Padovani, Jovanka Pantovic, Frank Pfenning, Antonio Ravara, Konstantinos Sagonas, Alceste Scalas, Nicolas Tabareau, Peter Thiemann, Hugo Torres Vieira, Emilio Tuosto, Vasco T. Vasconcelos, Philip Wadler, Nobuko Yoshida, Shoji Yuen

Behavioural types describe dynamic aspects of a program, in contrast to data types, which describe the fixed structure of data. Behavioural types include session types, typestate, choreographies, and behavioural contracts. Recent years have seen a substantial increase in research activity, including theoretical foundations, design and implementation of programming languages and tools, studies of the relationships between different forms of behavioural types, and studies of the relationships between behavioural types and more general type-theoretic ideas such as gradual typing and dependent typing. The aim of this seminar was to bring together researchers on behavioural types and related topics, in order to understand and advance the state of the art.

Many of the participants have been active in COST Action IC1201: Behavioural Types for Reliable Large-Scale Software Systems (BETTY), a European research network on behavioural types. Other participants were invited from related research areas and from outside Europe, in order to broaden the scope of the seminar and to make connections between communities.

The programme for the first half of the week was planned in advance, with priority given to two kinds of presentation: (1) demonstrations of programming language implementations and tools, and (2) presentations by participants from outside the BETTY community. The programme for the second half of the week evolved during the seminar, with more emphasis on group discussion sessions.

The seminar was judged to be a success by all the participants. At least one conference submission resulted from collaboration started during the week, other existing collaborations made substantial progress, and several participants planned a submission to the EU RISE funding scheme. We intend to propose a follow-on seminar on a similar topic in the future.

This report contains the abstracts of the talks and software demonstrations, and summaries of the group discussion sessions.

6.8 Wildly Heterogeneous Post-CMOS Technologies Meet Software

Organizers: Jerónimo Castrillón-Mazo, Tei-Wei Kuo, Heike E. Riel, and Matthias Lieber

Seminar No. 17061

Date: February 5–10, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.1

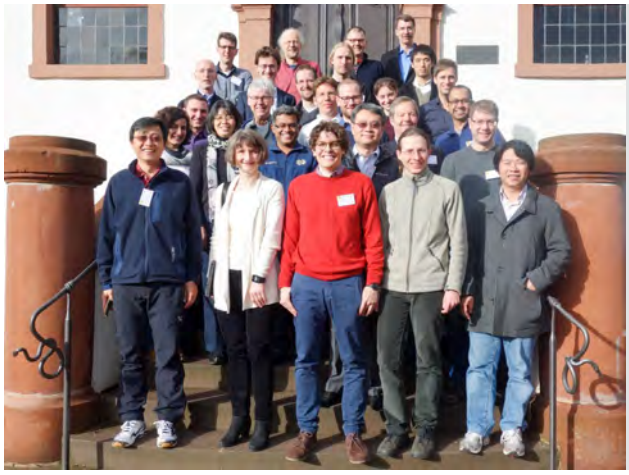

Participants: Rehan Ahmed, Stefano Ambrogio, Nils Asmussen, Tal Ben-Nun, Jerónimo Castrillón-Mazo, Jian-Jia Chen, Yiran Chen, Martin Claus, Ayse Coskun, Suman Datta, Erik P. DeBenedictis, Gianluca Fiori, Shunsuke Fukami, Balazs Gerofi, Sebastian Hack, Xiaobo Sharon Hu, Paul H. J. Kelly, Christoph W. Kessler, Tei-Wei Kuo, Roland Leißa, Max Lemme, Matthias Lieber, Michael Niemier, Stuart Parkin, Andy D. Pimentel, Sebastian Reineke, Heike E. Riel, Mark Rodwell, Eleonora Testa, Marcus Völp

Topic and Structure

The end of exponential scaling in conventional CMOS technologies has been forecasted for many years by now. While advances in fabrication made it possible to reach limits beyond those predicted, the so anticipated end seems to be imminent today. An indication of this is the research boom, both in academia and industry, in emerging technologies that could complement or even replace CMOS devices. Examples for such emerging technologies include tunnel FETs, nonvolatile memories such as magnetoresistive RAM, 3D integration, carbon nanotube transistors, and graphene.

The main goal of this seminar was to discuss bridges between material research, hardware components and, ultimately, software for information processing systems. Given a new class of wildly heterogeneous systems that integrate different technologies, we want to reason about enabling hardware and software abstractions, from languages and system-software down to hardware mechanisms. The challenge of realizing an efficient wildly heterogeneous system can only be tackled by employing holistic and synergistic approaches in an interdisciplinary environment. By bringing together experts from the individual fields and also researchers working interdisciplinarily across fields, the seminar helped to foster a mutual understanding about the challenges of advancing computing beyond current CMOS technology and to create long-term visions about a future hardware/software stack.

The seminar was structured around four partially overlapping areas, namely: (i) far-fetched materials and physics such as spin, nanomagnets, phase transition, and correlated phenomena, (ii) near future materials (and software) such as phase-change memory, nanowires, nanotubes, and neuromorphic devices, (iii) low-level software layers for new technologies such as operating systems, runtime support, middleware, and HW/SW-co-designed firmware, and (iv) upper software layers such as new programming/specification languages, models, and software synthesis.

Important questions addressed by the seminar included:

- Materials/Devices: What are the current status and the roadmap of post-CMOS materials and technologies? What will be the expected characteristics of the new devices? Will new technologies enable a fundamentally different computing paradigm, e. g., beyond von Neumann? What are the challenges for proper benchmarking of different technologies?
- Hardware/Software Stack: How much of the hardware's heterogeneity and its characteristics should be exposed to programmers? How general may be a programming model/language for future (yet unknown) hardware? How to make software adapt itself to hardware with fluctuating resources? Which new applications can be enabled by emerging materials and technologies and what needs to be done at the software layers to make them viable?
- Analysis: How can we model the interactions across the layers of the hardware/software stack? What kind of formal operational models and analysis methods are needed for evaluating heterogeneous systems? Can system-level analysis of new technologies give insights to material scientists, disrupting the otherwise incremental innovation paradigm?

Main Conclusions

Summary

There will probably be no CMOS replacement for chips with billions of transistors in the next 20 years, but architectural advances at various levels (such as 3D transistors, 3D integration of memory and logic, specialization, and reconfigurability) will lead to performance improvements despite the scaling limitations of planar CMOS technology. New non-volatile memories (e. g., spin-based) bear the potential to radically change various areas of computing, such as data-intensive processing and neuromorphic computing. New hardware architectures will need rethinking today's software stack and our widely used programming models. Finally, even though some post-CMOS technologies will not replace high-end CMOS transistors, there is great potential in

new, yet unknown, applications. Applications, backed by a strong commercial demand, will give some technologies the push to become viable. Examples are radio-frequency for carbon nanotubes, graphene based sensors, organic low-cost transistors for wearables, and memristors for neuromorphic computing.

Post-CMOS logic for compute-intensive applications

Currently, there is no alternative to CMOS on the horizon to realize logic for large von Neumann computing, due to lower projected performance and/or yield challenges. Candidates discussed on the seminar have been: tunnel FETs, III-V, 2D materials such as graphene, CNTs, or spintronics. This means that general purpose and high-performance computing will most probably be based on CMOS in the medium term. To workaround the CMOS scaling problem, architectural specialization will gain more and more importance leading to general purpose computing systems with (various) specialized accelerators. We already find them today in, e. g., mobile devices or GPU high-performance computing accelerators. Additionally, reconfigurable logic, such as FPGAs, and application-specific circuits have a high potential for performance gains. However, it is a big challenge to program such heterogeneous systems. Work towards solutions based on dataflow programming, memory access patterns, skeletons, and domain-specific languages have been discussed at the seminar. Additionally, operating system might need to adapt to allow, for example, accelerators to perform system calls.

Emerging memory technologies

In near future, new non-volatile memories will be available that could unify RAM and permanent storage, including MRAM and RRAM. While these could provide huge benefits for memory-intensive applications, the implications on architecture and software stack are not yet clear. For example, what will be the role of the file system in such an architecture? And how to deal with security aspects when every bit in RAM is permanent? Looking further into the future, the spin-based, non-volatile racetrack memory has the potential to compete with SRAM in terms of performance, while consuming considerably less energy. High-performance and energy-efficient non-volatile memories will also be important for neuromorphic devices.

Going 3D

3D integration enables the integration of heterogeneous technologies for logic, memory, communication, and sensing on a single chip. At the transistor level, 3D corrugated transistors were discussed as a promising direction to keep reducing the footprint while avoiding short-channel effects. Advancing today's die stacking technology through fine-grained vias linking the layers, will provide a substantial improvement for latencies and bandwidths in the systems. Bringing memory closer to logic will lower the memory wall (or even lead to a breakdown?). This means that many existing applications could be compute bound (again) and the processor architecture could be potentially simplified by removing the overhead that was added to workaround the memory wall, such as big caches and prefetchers, making place for additional compute units. In this optimistic scenario, general-purpose computing would receive a great (one-time) performance boost. For compilers and applications, we would have to rethink our way of optimizing code.

Computing beyond von Neumann

Architectural approaches beyond von Neumann were also discussed to speedup specific applications. Examples were neuromorphic computers, analog circuits, and dataflow machines. Of course these approaches cannot replace general purpose processors completely. A possible future architecture would combine classical von Neumann processors with non-von Neumann accelerators (on the same chip) to enable mixed programming. The recent industry adoption of machine learning drives the need for neuromorphic computers. While these systems already outperform general purpose processors today, new technologies such as non-volatile memories and analog spintronics promise even greater gains. Promising analog circuits were shown to perform well for concrete NP-complete problems such as SAT and graph coloring. Along these lines, a theoretical framework was introduced that may serve to abstractly compare the asymptotic energy efficiency between the analog and the digital realizations of a system. Finally, dataflow machines where discussed that stream data directly between computational units without the overhead of registers and caches, thereby removing the "Turing tax".

Special applications

Some of the materials considered in the seminar are very likely not able to compete with CMOS for logic, but have strengths in other electronic application areas such as sensors, radio frequency, and displays. Carbon nanotubes and graphene are promising materials for high-frequency antennas required for upcoming wireless communication systems. In the particular case of Graphene, it seems that the initial technological hype has passed, and engineering has taken over to produce new clever devices (e. g., nano-membranes for sensing). Organic electronics are already commercially available in displays and OLEDs. Their distinct features of flexibility, low production cost (printed electronics), and biodegradability could potentially open completely new application areas for logic, but not at comparable speed and efficiency to CMOS. These devices have also been deemed important for bio-compatibility. However, there is a long road ahead for testing and certifying actual devices in living tissue, which is not a trivial task, considering the wealth of molecules being investigated in this domain.

Co-design and design space exploration

Proper hardware/software co-design will be very important to achieve performance gains given the limits of CMOS and the prospective wildly heterogeneous and/or application-specific computing systems. Given a specific application problem, numerous implementation alternatives, from the algorithm down to the hardware architecture and technologies, might be feasible. Tools that help developers navigating the huge design space (e. g., using modeling and benchmarking techniques) and automate an efficient implementation as much as possible are needed. It appears to be that the large part of the software is less flexible than the hardware and much work has to be done to make software future-proof.

6.9 Beyond VR and AR: Reimagining Experience Sharing and Skill Transfer Towards an Internet of Abilities

Organizers: Anind K. Dey, Jonna Häkkilä, Niels Henze, Kai Kunze

Seminar No. 17062

Date: February 5–10, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.23

Participants: Susanne Boll, Cedric Carêmel, Ashley Colley, Anind K. Dey, N. Chloe Eghtebas, Hans-Werner Gellersen, Nitesh Goyal, Scott Greenwald, Jonna Häkkilä, Niels Henze, Masahiko Inami, Shunichi Kasahara, Gudrun Klinker, Kai Kunze, Pedro Lopes, Stephan Lukosch, Joachim Meyer, Florian Michahelles, Yun Suen Pai, Thies Pfeiffer, Yannick Prié, Enrico Rukzio, Thad Starner, Yuji Uema, Kaisa Väänänen, Dan Witzner Hansen, Katrin Wolf, Eva Wolfangel

Sharing experiences and knowledge have always been essential for human development. They enable skill transfers and empathy. Over history, mankind developed from oral traditions to cultures of writing. With the ongoing digital revolution, the hurdles to share knowledge and experiences vanish. Already today it is, for example, technically feasible to take and store 24/7 video recordings of one's life. While this example creates massive collections of data, it makes it even more challenging to share experiences and knowledge with others in meaningful ways. Facilitating the third wave of VR and AR technologies we are currently witnessing, researchers started to broadly look at VR and AR again. Topics ranging from using AR to mitigate skills gaps [1] and understanding user interaction with commercial AR games [4], to using focus depth as an input modality for VR [2], and understanding the effect of gender in VR [3]. The goal of the seminar was to take a step back from the technical research to look at the fundamental aspects of interactive media.

A recurring theme in science fiction literature is the act of downloading another human's abilities to one's mind. Although current cognitive science and neuroscience strongly suggest that this is impossible, as our minds are embodied; we believe that skill transfer and effective learning will accelerate tremendously given recent technological trends; just to name a few of the enabling technologies, human augmentation using virtual/augmented reality, new sensing modalities (e.g. affective computing) and actuation (e.g. haptics), advances in immersive storytelling (increasing empathy, immersion, communication) etc.

Ultimately, we believe this will lead to "downloadable" experiences and abilities. The effects will definitely not be instant and it will most likely be very different from the Sci-Fi theme. Yet, these differences are exactly what we want to explore in this seminar. Computer scientists in wearable computing, ubiquitous computing, human computer interaction, affective computing, virtual reality and augmented reality have been working on related topics and enabling technologies for years. However, these developments are disjointed from each other. With this seminar we want to bring them together working in the virtual/augmented/mixed reality, ubiquitous computing, sensing and HCI fields discussing also with experts in cognitive science, psychology and education.

While sharing experiences and knowledge through communication and socializing are a long time focus of various research efforts, we believe it is necessary to rethink and redefine experience sharing and skill transfer in light of the following current technological advances like the following:

1. Affordable Virtual Reality and Augmented Reality systems will become available to consumers in the near future (or already are available).
2. Advances in new sense sharing technologies (e.g. eye gaze, haptics, odors).
3. Advances in real-life tracking of physical and cognitive activities and emotional states.
4. Educators, cognitive scientists and psychologists have now a better understanding of individual and group behaviors, empathy and fundamentals of learning.

The seminar was structured around lightning talks by the participants, two hands-on workshops and three thematic sessions. In the lightning talks, the participants introduced themselves and shared their vision with the group. The first hands-on workshop by Shunichi Kasahara introduced the term Superception and showcased prototypes in this domain. The second workshop organized by Pedro Lopes enabled participants to experiment with electrical muscle stimulation by connecting off-the-shelf devices to embedded systems. Three days of the seminar started with thematic sessions run by one of the organizers. The sessions explored the future of human-computer symbiosis, human augmentation, and enabling technologies.

References

1 Yamashita, S., Matsuda, A., Hamanishi, N., Suwa, S., & Rekimoto, J. (2017, March). Demulti Display: A Multiplayer Gaming Environment for Mitigating the Skills Gap. In Proceedings of the Tenth International Conference on Tangible, Embedded, and Embodied Interaction (pp. 457–463). ACM.

2 Pai, Y. S., Outram, B., Vontin, N., & Kunze, K. (2016, October). Transparent Reality: Using Eye Gaze Focus Depth as Interaction Modality. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology (pp. 171–172). ACM.

3 Schwind, Valentin; Knierim, Pascal; Tasci, Cagri; Franczak, Patrick; Haas, Nico; Henze, Niels. "These are not my hands!": Effect of Gender on the Perception of Avatar Hands in Virtual Reality Inproceedings CHI '17 Proceedings of the 2017 Annual Symposium on Computer-Human Interaction, ACM Press, New York, NY, USA, 2017, ISBN: 9781450346559,

4 Colley, Ashley, Jacob Thebault-Spieker, Allen Yilun Lin, Donald Degraen, Benjamin Fischman, Jonna Häkkilä, Kate Kuehl, Valentina Nisi, Nuno Jardim Nunes, Nina Wenig, Dirk Wenig, Brent Hecht and Johannes Schöning. 2017. The Geography of Pokémon GO: Beneficial and Problematic Effects on Places and Movement. In Proc. CHI'17.

6.10 Computer-Assisted Engineering for Robotics and Autonomous Systems

Organizers: Erika Abraham, Hadas Kress-Gazit, Lorenzo Natale, and Armando Tacchella

Seminar No. 17071

Date: February 12–17, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.48

Participants: Erika Abraham, Jan Broenink, Simon Burton, Pablo Bustos, Alessandro Cimatti, Manuel Alcino Cunha, Kerstin I. Eder, Rüdiger Ehlers, Samira Farahani, Stefano Ghidoni, Christoffer R, Heckman, Christian Heinzemann, Holger Hermanns, Felix Ingrand, Nils Jansen, Joost-Pieter Katoen, Hadas Kress-Gazit, Jan Kretinsky, Morteza Lahijanian, Gerhard Lakemeyer, Ratan Lal, Martin Leucker, Ingo Lütkebohle, Daniele Magazzeni, Björn Matthias, Lorenzo Natale, Sandeep Neema, Petter Nilsson, Shashank Pathak, Subramanian Ramamoorthy, Vasumathi Raman, Kristin Yvonne Rozier, Christian Schlegel, Maria Svorenova, Armando Tacchella, Sebastian Wrede

This seminar focused on *autonomous systems*, and more specifically robots, that operate without, or with little, external supervision. For these systems to be integrated into society, it is highly important to make sure that they are functionally safe. *Formal Methods* are techniques adopted in engineering for the verification of software and hardware systems. As models are a basic requirement for the formal analysis of systems, *Model-driven Software Engineering* plays an important role to enable the application of *Formal Methods*. Though autonomous systems are increasingly involved in our everyday life, both exact formalizations of safe functionality (standards, what we want to be confident in) and methods to achieve confidence (methodologies, how we get confident in the properties we want to assure) are still scarce.

This seminar brought together experts in *Artificial Intelligence* and *Robotics*, *Model-driven Software Engineering*, and *Formal Methods*. It included researchers from academia as well as from industry. The following list summarizes high-level themes that emerged from the seminar:

- Dealing with highly complex systems, it is difficult to verify or even model all aspects of the system, therefore focusing effort on efficient falsification rather than costly verification can be highly impactful for industrial applications.
- The community can and should leverage results and systems built for different robotic competitions to reason about possible requirements and techniques to verify/falsify them. These competitions include the DARPA robotics challenge, the Amazon picking challenge, different leagues in Robocup, etc. Creating benchmarks based on these competitions will enable progress in verification of autonomous systems.
- Creating small interdisciplinary teams that include people from formal methods, robotics and model based design that tackle small yet realistic problems, possibly inspired by industrial applications, will help formalize the language of requirements, models and verification techniques that will have an impact on autonomous systems.

Fig. 6.1

"Feeling like a knight in the gorgeous @dagstuhl castle. Here for a exciting computational proteomics seminar!" Twitter post by 17421 Dagstuhl Seminar participant Pieter-Jan Volders.
https://twitter.com/pjvolders/status/919895848397492227. Photo courtesy of Pieter-Jan Volders.

6.11 Applications of Topology to the Analysis of 1-Dimensional Objects

Organizers: Benjamin Burton, Maarten Löffler, Carola Wenk, and Erin Moriarty Wolf Chambers

Seminar No. 17072

Date: February 12–17, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.64

Participants: Maike Buchin, Benjamin Burton, Gregory R. Chambers, Hsien-Chih Chang, Arnaud de Mesmay, Anne Driemel, Parker Evans, Brittany Terese Fasy, Philipp Kindermann, Linda Kleist, Boris Klemz, Stephen G. Kobourov, Irina Kostitsyna, David Letscher, Maarten Löffler, Anna Lubiw, Wouter Meulemans, Martin Nöllenburg, Tim Ophelders, Florian T. Pokorny, Jessica S. Purcell, Regina Rotman, Saul Schleimer, Lena Schlipf, André Schulz, Eric Sedgwick, Bettina Speckmann, Jonathan Spreer, Frank Staals, Darren Strash, Dylan Thurston, Stephan Tillmann, Marc van Kreveld, Mikael Vejdemo-Johansson, Birgit Vogtenhuber, Carola Wenk, Erin Moriarty Wolf Chambers

Description of the seminar. One-dimensional objects embedded in higher-dimensional spaces are one of the most natural phenomena we encounter: ranging from DNA strands to roads to planetary orbits, they occur at all granularities throughout the sciences. Computer-assisted analysis of one-dimensional data is now standard procedure in many sciences; yet the underlying mathematics are not always well understood, preventing the most powerful analytical tools from being used.

Adding to the confusion, one-dimensional objects are studied under different names in different areas of mathematics and computer science (knots, curves, paths, traces, trajectories). In mathematics, 1-dimensional objects are well-understood, and research endeavors have moved on to higher dimensions. On the other hand, many fundamental applications demand solutions that deal with 1-dimensional objects, and these computational problems have largely been studied in separate communities by those unaware of all of the mathematical foundations.

The main goal of the proposed seminar was to identify connections and seed new research collaborations along the spectrum from knot theory and topology, to computational topology and computational geometry, all the way to graph drawing. Each of the invited speakers explored synergies in algorithms concerning 1-dimensional objects embedded in 2- and 3-dimensional spaces, as this is both the most fundamental setting in many applications, as well as the setting where the discrepancy in computational complexity between generic mathematical theory and potential algorithmic solutions is most apparent. In addition, each talk proposed a set of open questions from their research area that could benefit from attention from the other communities, and participants of the seminar were invited to propose their own research questions.

Below, we (the organizers) briefly describe the three main areas bridged; the abstracts of talks in the seminar and preliminary results from the working groups are also outlined later in this report.

Curves in Trajectory Analysis. Applications of computational topology are on the rise; examples include the analysis of GIS data, medical image analysis, graphics and image modeling, and many others. Despite how fundamental the question of topological equivalence is in mathematics, many of the relatively simple settings needed in computational settings (such as the plane or a 2-manifold) have been less examined in mathematics, where computability is known but optimizing algorithms in such "easy" settings has not been of interest until relatively recently.

Homotopy is one of the most fundamental problems to consider in a topological space, as this measure captures continuous deformation between objects. However, homotopy is notoriously difficult, as even deciding if two curves are homotopic is undecidable in a generic 2-complex. Nonetheless, many application settings provide restrictions that make computation more accessible. For example, most GIS applications return trajectories in a planar setting, at which point finding optimal homotopies (for some definition of optimal) becomes tractable.

Homology has been more recently pursued, as finding good homologies reduces to a linear algebra problem which can be solved efficiently. An example of this in the 1-dimensional setting is the recent work by Pokorny on clustering trajectories based on relative persistent homology. However, it is not always clear that optimal homologies provide as intuitive a notion for similarity measures compared with homotopy, and further investigations into applications settings is necessary.

Curves in Knot Theory. A fundamental question in 3-manifold topology is the problem of isotopy. Testing if two curves are ambiently isotopic is a foundational problem of *knot theory*: essentially, this asks whether two knots in 3-space are topologically equivalent. Problems in knot theory are tightly

related to problems in 3-manifold topology, a field that has seen major breakthroughs in recent years, including Perelman's 2002 solution to the geometrisation and Poincaré conjectures, and Agol's recent 2012 proof of the virtual Haken conjecture. Algorithms and computation in these fields are now receiving significant attention from both mathematicians and computer scientists.

Complexity results are surprisingly difficult to come by. For example, one of the most fundamental and best-known problems is detecting whether a curve is knotted. This is known to be in both NP and co-NP; the former result was shown by Hass, Lagarias and Pippenger in 1999, but the latter was proven unconditionally by Lackenby just this year. Finding a polynomial time algorithm remains a major open problem. Hardness results are known for some knot invariants, but (despite being widely expected) no hardness result is known for the general problem of testing two knots for equivalence. Techniques such as randomisation and parameterised complexity are now emerging as fruitful methods for understanding the inherent difficulty of these problems at a deeper level.

Algorithmically, many fundamental problems in knot theory are solved by translating to 3-manifold topology. Here there have been great strides in practical software in recent years: software packages such as *SnapPy* and *Regina* are now extremely effective in practice for moderate-sized problems, and have become core tools in the mathematical research process. Nevertheless, their underlying algorithms have significant limitations: *SnapPy* is based on numerical methods that can lead to numerical instability, and *Regina* is based on polytope algorithms that can suffer from combinatorial explosions. It is now a major question as to how to design algorithms for knots and 3-manifolds that are exact, implementable, and have provably viable worst-case analyses.

Curves in Graph Drawing. On the computer science end of the spectrum, the study of one-dimensional objects is closely related to Graph Drawing.

Graph Drawing studies the embedding of zero- and one-dimensional features (vertices and edges of graphs) into higher-dimensional spaces; both from an analytic (given an embedding, what can we say about it) and synthetic (come up with a good embedding) point of view. Computational questions (how can we embed a given graph such that it satisfies certain properties / optimises certain criteria) and fundamental questions (which classes of graphs admit which styles of embeddings) have been studied extensively, and a large body of algorithmic results is readily available.

Planarity (non-crossing edges) is a central theme in graph drawing. There is a rich literature discussing which graphs can be drawn planarly, when, and how, as well as how to avoid crossings or other undesirable features in a drawing, such as non-rational vertices. Traditionally, edges have always been embedded as straight line segments; however, there is a recent trend to consider different shapes and curves, drastically increasing the space of possible drawings of a graph. The potential benefits of this broader spectrum are obvious, but the effects (both computational and fundamental) are still ill understood.

Connections between graph drawing and knot theory have long been recognised, yet are still being actively explored. Already in 1983, Conway and Gordon showed that every spatial representation of K_7 contains at least one knotted Hamiltonian cycle. Based on this, in 2013, Politano and Rowland characterised which knots appear as Hamiltonian cycles in canonical book embeddings of complete graphs (as defined by Otsuki in 1996).

Goals and Results of this Seminar. Now is an exciting time for computational and algorithmic knot theory: practical algorithms are showing their potential through experimentation and computer-assisted proofs, and we are now seeing key breakthroughs in our understanding of the complex relationships between knot theory and computability and complexity theory. Early interactions between mathematicians and computer scientists in these areas have proven extremely fruitful, and as these interactions deepen it is hoped that major unsolved problems in the field will come within reach.

Similarly, applications for graph drawing and trajectory analysis are in great demand, especially given the rise of massive amounts of data through GIS systems, map analysis, and many other application areas. However, despite the fact that many problems on curves are seen as mathematically trivial, there are few CS researchers who are truly familiar with the deeper topological results from mathematics. It is likely that many algorithmically interesting questions can benefit from an understanding of this rich history and toolset.

This seminar brought together a group of researchers from computer science and mathematics that study algorithms and mathematical properties of curves in various settings, as the interplay between these two groups is recent. In addition, we invited researchers in applications domains, who often do heuristic analysis of 1-dimensional objects in a variety of settings. Working groups were formed organically, but often allowed participants from various subfields to swap both open problems and favorite tools, and the overview talks discussed favorite tools and techniques from subdomains that may be useful to those in other areas. Concretely, we hope that in addition to the work begun in the working groups, many of these new collaborations will have positive long-term effects on all areas.

6.12 Computability Theory

Organizers: Klaus Ambos-Spies, Vasco Brattka, Rodney Downey, and Steffen Lempp

Seminar No. 17081

Date: February 19–24, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.89

Participants: Klaus Ambos-Spies, George Barmpalias, Laurent Bienvenu, Vasco Brattka, Peter Cholak, Chitat Chong, Barbara Csima, Rodney Downey, Willem L. Fouché, Sergey Goncharov, Noam Greenberg, Rupert Hölzl, Iskander Shagitovich Kalimullin, Takayuki Kihara, Julia Knight, Ulrich Kohlenbach, Steffen Lempp, Alberto Marcone, Alexander Melnikov, Wolfgang Merkle, Joseph S. Miller, Russell G. Miller, Kenshi Miyabe, Antonio Montalbán, Keng Meng Ng, André Otfrid Nies, Arno Pauly, Alexander Shen, Richard A. Shore, Theodore A. Slaman, Andrea Sorbi, Mariya I. Soskova, Sebastiaan A. Terwijn, Henry Towsner, Daniel Turetsky, Nikolay K. Vereshchagin, Wei Wang, Klaus Weihrauch, Linda Brown Westrick, Guohua Wu, Yue Yang, Liang Yu

Computability theory grew from work to understand effectiveness in mathematics. Sophisticated tools have been developed towards this task. For a while, the area tended to be concerned with internal considerations such as the structure of the various hierarchies developed for the tasks of calibrations. More recently, this machinery has seen modern applications into areas such as model theory, algorithmic randomness, analysis, ergodic theory, number theory and the like; and the tools have been used to answer several classical questions. Seminar 17081 was an opportunity for researchers in several areas of modern computability theory to get together and interact.

The format was for 2–3 lectures in the morning with at least one being an overview, and a similar number of 3–4 in the afternoon, with Wednesday afternoon and Friday afternoon free. The weather being miserable, participants opted to stay at the Schloss Wednesday afternoon, and quite a bit of work was done in pairs in the time left free, on the Wednesday afternoon in particular. At least one problem seen as significant in the area was solved (one concerning the strength of Ramsey's Theorem for Pairs in reverse mathematics), and the organizers know of several other papers in preparation resulting from the seminar.

The lectures were from various areas, but the greatest concentration was around

- classification tools in computable analysis (the Weihrauch Lattice) and Reverse Mathematics (on what proof-theoretic strength is needed to establish results in mathematics), and these areas' relationship with generating algorithms, such as in proof mining;
- computable model theory (looking at structures such as groups, rings, or abstract algebraic structures, endowing them with effectivity and seeing what else is algorithmic). Notable was the new work on effective uncountable structures such as uncountable free groups, and on topological groups;
- algorithmic randomness: Here one seeks to give meaning to randomness for individual strings and infinite sequences. Talks given explored the relationship of calibrations of randomness to computational power, and online notions of randomness.

Of course, these are not separate areas but are inter-related, and the talks reflected these inter-relationships.

Currently, computability theory is quite vibrant with many new applications being found, and a number of young and talented researchers entering the field. This was reflected in the age of the presenters of many of the lectures, as well as the significant number of people we could have invited in addition.

All in all, the meeting was a great success and should have significant impact on the development of the area.

6

6.13 Shape-Changing Interfaces

Organizers: Jason Alexander, Sean Follmer, Kasper Hornbæk, and Anne Roudaut

Seminar No. 17082

Date: February 19–22, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.102

Participants: Jason Alexander, Miguel Bruns Alonso, Céline Coutrix, Pierre Dragicevic, Aluna Everitt, Sean Follmer, Audrey Girouard, Marianne Graves Petersen, Fabian Hemmert, Trevor Hogan, Kasper Hornbaek, Yvonne Jansen, Matt Jones, Hyunyoung Kim, Majken Kirkegaard Rasmussen, Panos Markopoulos, Diego Martinez Plasencia, Timothy Merritt, Isabel Qamar, Anne Roudaut, David Schmelzeisen, Jürgen Steimle, Sri Subramanian, Roel Vertegaal, Helge A. Wurdemann

The Shape-Changing Interfaces Dagstuhl seminar aimed to bring together researchers from the disciplines of Computer Science, Design, Engineering, Robotics and Material Science to strenghten this new communinity, discuss grand challenges, form a research agenda, and to create opportunities for active collaborations.

Shape-changing interfaces use changes in physical geometry to convey input and output and are emerging as an alternative interaction method for communicating with computers. Discussions at the seminar were based around three key themes: (1) The technologies involved in shape-change, including soft and modular robotics, smart materials, and mechanical actuation. (2) The design of shape-changing interfaces, including their key application areas, and their industrial and interaction design. (3) The user experience of shape-changing interfaces, including evaluations of such interfaces and psycho-physical evaluation results.

To encourage active discussion, the seminar had no keynote speakers, but instead used brainstorming activities and small working-groups to understand challenges, explore the literature, and plan an agenda. Specifically, the following sessions were run:

Benefits and Applications of Shape-Change: A whole-group brainstorming session developed categories of benefits and potential application areas for shape-changing interfaces.

Related Work: Small working-groups focused on one of five related-work areas (materials, hardware, experience and interaction, design, or applications), researched, and then presented summaries of the five ground-breaking and five most over-looked works in that sub-field.

Grand Challenges: A whole-group brainstorming session generated ideas and themes of grand challenges, small working groups then took a theme and dug deeper into the challenge, generating avenues of work and research agendas.

Worst Case Scenarios: To understand why this field could fail, a brainstorming exercise asked participants to develop a series of 'failure' situations—these were used as a method of creating awareness of the reasons progress in this field could stall.

Personal Reflections: To conclude the seminar, four participants were asked to provide their personal reflections on the experience, and their key take-home messages.

The seminar was attended by 25 researchers from around the world; all of who found the experience invaluable. This report outlines the key findings of these sessions.

6.14 Computer Science Meets Ecology

Organizers: Gustau Camps-Valls, Thomas Hickler, and Birgitta König-Ries

Seminar No. 17091

Date: February 26 to March 3, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.109

Participants: Martin Bücker, Tilo Burghardt, Gustau Camps-Valls, Yun-Heh Jessica Chen-Burger, Joachim Denzler, Matthew Evans, Florian Hartig, Thomas Hickler, Donald Hobern, Forrest Hoffman, Kazuhito Ichii, Martin Jung, Birgitta König-Ries, Ivaylo Kostadinov, Bertram Ludäscher, Miguel Mahecha, Laetitia Navarro, Shawn Newsam, Frank Pennekamp, Natalia Petrovskaya, Markus Reichstein, Andrew Richardson, Ribana Roscher, Brody Sandel, Bernhard Seeger, Johann Wolfgang Wägele, Jakob Zscheischler

Ecology is a discipline that shows clearly the potential but also the challenges of computer supported research described as the 4th scientific paradigm by Jim Gray. It is increasingly data driven, yet suffers from hurdles in data collection, quality assurance, provenance, integration, and analysis.

We believe that ecology could profit from modern computer science methods to overcome these hurdles. However, usually, scientists in ecology are not completely aware of current trends and new techniques in computer science that can support their daily work. Such support could consist in the management, integration, and (semi-)automatic analysis of resources, like experimental data, images, measurements, in the generation of useful metadata, cloud computing, distributed processing, etc. Ecoinformatics is regarded as an important supporting discipline by many ecologists. However, up to now, very few computer scientists are involved in this discipline; mostly ecoinformatics (or biodiversity informatics) is done by people with a strong background in e.g. ecology and a long (mostly self-taught) experience in data management. It lacks a strong connection to cutting-edge computer science research in order to profit from the results of this area. On the other hand, computer scientists know too little about the domain to be able to offer solutions to relevant problems and to identify potential research avenues.

Motivated by our belief that a stronger bond between the disciplines that goes beyond viewing computer science as a "service provider" is of vital importance, we proposed this Dagstuhl seminar. The aim of the Dagstuhl seminar was to establish such links between (geo-)ecologists, ecoinformaticians and computer scientists.

The seminar: perspective and self-evaluation

Before the seminar. It turned out that it was not an easy task to motivate non-computer scientists to attend the seminar. For many, travel costs were a hurdle ultimately preventing attendance. This resulted in an unusually large number of declined invitations (often accompanied by "I would love to attend, but…" emails.

Despite these initial problems, we believe that the aim to start building links among the communities was reached at the seminar: We had fruitful discussions in numerous working groups resulting in some very concrete plans for future work.

Organization of the seminar. A total of 27 attendees gathered at the seminar. The wide variety of expertise and backgrounds constituted an initial challenge for the organization. The agenda considered a first round of presentations of the individuals and their research groups with a clear outline and items to treat (personal background, Research Areas/Interests, prospective links to „Computer Science meets Ecology" seminar). After this, the main topics of interest for a wide audience were designed: essentially, three breakout groups were set up in the very first day of the meeting. Over the course of the seminar, these groups were adjusted, split up, or merged, several times. This resulted in quite a number of topics being touched upon with concrete results ranging from a working example for the application of a new method to a modeling problem to concrete plans for publications, a proposal and follow-up activities. Reports on these groups were given in the plenary session, and can be found in this report.

Broad results of the seminar. Results from the seminar can be categorized in three types: (i) collaborative and networking, as new joint works on specific topics came out of the meeting; (ii) knowledge transfer between fields, as computer scientists learned about the main problems in ecology involving data, while ecologists became aware of what kind of problems data scientists can solve nowadays; and (iii) educational, as several young PhD students and postdocs attended and participated in high level discussions.

Conclusions. The seminar brought together top scientists in the fields of ecology and computer science. The group of individuals was largely interdisciplinary, with a wide range of interests and expertises in each community too: from botany and animal science, to machine learning and computer vision. The seminar was organized in two main types of sessions: plenary and working group sessions to better focus on particular topics. Interesting developments and discussions took place in both, and a high level of cross-fertilization and future collaborations was initiated. On top of this, there was a broad consensus among the participants that the seminar should be the start of a series of yearly or bi-yearly meetings. We hope that the success of this first seminar will encourage broader participation in follow-up activities.

6.15 Search as Learning

Organizers: Kevyn Collins-Thompson, Preben Hansen, and Claudia Hauff

Seminar No. 17092

Date: February 26 to March 1, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.2.135

Participants: Dirk Ahlers, Leif Azzopardi, Robert Capra, Kevyn Collins-Thompson, Gwen Frishkoff, Ujwal Gadiraju, Tim Gollub, Preben Hansen, Claudia Hauff, Jiyin He, Gabriele Irle, Yvonne Kammerer, Noriko Kando, Yiqun Liu, Andreas Nürnberger, Heather O'Brien, Rebecca B. Reynolds, Soo Young Rieh, Marc Rittberger, Daniel Russell, Marcus Specht, Michael Twidale, Alyssa Wise, Christa Womser-Hacker, Grace Hui Yang, Ran Yu

Search is everywhere – it penetrates every aspect of our daily lives and most of us can hardly manage a few hours without resorting to a search engine for one task or another. Despite the success of existing (Web) search technology, there are still many challenges and problems that need to be addressed. Today's Web search engines (often also powering domain-specific and site-specific search) are engineered and optimized to fulfil individual users' lookup tasks. This efficiency, however, also means that we largely view search systems as tools to satisfy immediate information needs, instead of rich environments in which humans heavily interact with information content, and search engines act as intelligent dialogue systems, facilitating the communication between users and content. Web search engines are not designed for complex search tasks that require exploration and learning, user collaborations and involve different information seeking stages and search strategies, despite the fact that more than a quarter of Web searches are complex. In recent years, there has been a growing recognition of the importance of studying and designing search systems to foster discovery and enhance the learning experience during the search process outside of formal educational settings. Searches that lead to learning, are naturally complex. Research progress in this area, however, is slow, with many more open questions than answers. Several critical bottlenecks and major impediments to advancements in the search as learning area exist, including (i) the reliance on small-scale lab studies to evaluate novel approaches which severely limit the diversity of investigable factors as well as the ecological validity and generalizability of the findings; (ii) the lack of awareness among researchers' initiatives in this very multidisciplinary area of work; and (iii) the lack of a shared research infrastructure. The 3-day seminar gathered 26 prominent researchers from the fields of information retrieval, psychology and the learning sciences in order to address the critical bottlenecks around search as learning. The seminar sessions alternated between tutorial-style presentations to learn from each other's disciplines and interactive breakout sessions to find a common ground and address the most pressing issues related to the four big research themes of (i) understanding search as a human learning process; (ii) the measurement of learning performance and learning outcomes during search; (iii) the relationship between the learning process and the search context; and (iv) the design of functionalities and search system interventions to promote learning.

Fig. 6.2
dagstuhl - bodyCentric - wedAM Picture on flickr by 17392 Dagstuhl Seminar participant m. c. schraefel. https://flic.kr/p/YV5qaa. Photo courtesy of m. c. schraefel.

6.16 Databases on Future Hardware

Organizers: Gustavo Alonso, Michaela Blott, and Jens Teubner

Seminar No. 17101

Date: March 5–10, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.1

Participants: Anastasia Ailamaki, Gustavo Alonso, Carsten Binnig, Spyros Blanas, Michaela Blott, Alexander Böhm, Peter A. Boncz, Sebastian Breß, Markus Dreseler, Ken Eguro, Babak Falsafi, Henning Funke, Goetz Graefe, Christoph Hagleitner, Peter Hofstee, Stratos Idreos, Zsolt Istvan, Viktor Leis, Eliezer Levy, Stefan Manegold, Andrew W. Moore, Ingo Müller, Onur Mutlu, Thomas Neumann, Gilles Pokam, Kenneth Ross, Kai-Uwe Sattler, Eric Sedlar, Margo Seltzer, Jürgen Teich, Jens Teubner, Pinar Tözün, Annett Ungethüm, Stratis D. Viglas, Thomas Willhalm, Ce Zhang, Daniel Ziener

Computing hardware is undergoing radical changes. Forced by physical limitations (mainly *heat dissipation* problems), systems trend toward *massively parallel* and *heterogeneous* designs. New technologies, *e.g.*, for *high-speed networking* or *persistent storage* emerge and open up new opportunities for the design of database systems. This push by technology was the main motivation to bring top researchers from different communities – particularly hard- *and* software – together to a Dagstuhl seminar and have them discuss about "Databases on Future Hardware." This report briefly summarizes the discussions that took place during the seminar.

With regards to the mentioned *technology push*, during the seminar *bandwidth*; *memory and storage technologies*; and *accelerators* (or other forms of specialized computing functionality or instruction sets) were considered the most pressing topic areas in database design.

But it turned out that the field is influenced also by a strong push from *economy/market*. New types of *applications* – in particular Machine Learning – as well as the emergence of "compute" as an independent type of resources – *e.g.*, in the form of *cloud computing* or *appliances* – can have a strong impact on the viability of a given system design.

Bandwidth; Memory and Storage Technologies

During the seminar, probably the most often stated issue in the field was *bandwidth* – at various places in the overall system stack, such as CPU ↔ memory; machine ↔ machine (network); access to secondary storage (*e.g.*, disk, SSD, NVM). But very interestingly, the issue was not only brought up as a key limitation to database performance by the seminar attendees with a software background. Rather, it also became clear that the hardware side, too, is very actively looking at bandwidth. The networking community is working at ways to provide more bandwidth, but also to provide hooks that allow the software side to make better use of the available bandwidth. On the system architecture side, new interface technologies (*e.g.*, NVlink, available in IBM's POWER8) aim to ease the bandwidth bottleneck.

Bandwidth usually is a problem only *between* system components. To illustrate, HMC memories ("hybrid memory cube") provide only 320 GB/s of external bandwidth, but internally run at 512 GB/s per cube ("vault"); in a 16-vault configuration, this corresponds to 8 TB/s of internal bandwidth. This may open up opportunities to build heterogeneous system designs with *near-data processing* capabilities. HMC memory units could, for instance, contain (limited) processing functionality associated with every storage vault. This way, simple tasks, such as data movement, re-organization, or scanning could be off-loaded and performed right where the data resides. Similar concepts have been used, *e.g.*, to filter data in the network, pre-process data near secondary storage, etc.

In breakout sessions during the seminar, attendees discussed the implications that such system designs may have. Most importantly, the designs will require to re-think the existing *(programming) interfaces*. How does the programmer express the off-loaded task? Which types of tasks can be off-loaded? What are the limitations of the near-data processing unit (*e.g.*, which memory areas can it access)? How do host processor and processing unit exchange tasks, data, and results? Clearly, a much closer collaboration will be needed between the hard- and software sides to make this route viable.

But new designs may also shake up the commercial market. The traditional hardware market is strongly separated between the memory and logic worlds, with different manufacturers and processes. Breaking up the separation may be a challenge both from a technological and from a business/market point of view.

The group found only little time during the seminar to discuss another potential game-changer in the memory/storage space. Companies are about to bring their first *non-volatile memory*

(NVM) components to the market (and, in fact, Intel released its first round of "3D XPoint" products shortly after the seminar). The availability of cheap, high-capacity, byte-addressable, persistent storage technologies will have profound impact on database software. Discussions during the seminar revolved around the question whether classical persistent (disk-based) mechanisms or in-memory mechanisms are more appropriate to deal with the new technology.

Accelerators A way of dealing with the technology trend toward heterogeneity is to enrich general-purpose systems with more specialized processing units, *accelerators*. Popular incarnations of this idea are *graphics processors (GPUs)* or *field-programmable gate arrays (FPGAs)*; but there are also co-processing units for floating-point arithmetics, multimedia processing, or network acceleration.

Accelerators may fit well with what was said above. *E.g.*, they could be used as near-data processing units. But also the challenges mentioned above apply to many accelerator integration strategies. Specifically, the proper *programming interface*, but also the role of an accelerator in the software system stack – *e.g.*, sharing it between processes – seem to be yet-unsolved challenges.

During the seminar, also the role of accelerators specifically for database systems was discussed. It was mentioned, on the one hand, that accelerators should be used to accelerate functionality outside the database's core tasks, because existing hard- and software is actually quite good at handling typical database tasks. On the other hand, attendees reported that many of the non-core-database tasks, Machine Learning in particular, demand a very high flexibility that is very hard to provide with specialized hardware.

New Applications / Machine Learning Databases are the classical device to deal with high volumes of data. With the success of Machine Learning in many fields of computing, the question arises how databases and Machine Learning applications should relate to one another, and to which extent the database community should embrace ML functionality in their system designs.

Some of the seminar attendees have, in fact, given examples of very impressive and successful systems that apply ideas from database co-processing to Machine Learning scenarios. In a breakout session on the topic, it was concluded that the two worlds should still be treated separately also in the future.

A key challenge around Machine Learning seems to be the very high expectations with regard to the flexibility of the system. ML tasks are often described in high-level languages (such as R or Python) and demand expressiveness that goes far beyond the capabilities of efficient database execution engines. Attempts to extend these engines with tailor-made ML operators were not very well received, because even the new operators were too restrictive for ML users.

Economic/Market Considerations Somewhat unexpectedly, during the seminar it became clear that the interplay of databases and hardware is not only a question of technology. Rather, examples from the past and present demonstrate that even a technologically superior database solution cannot survive today without a clear business case.

The concept of *cloud computing* plays a particularly important role in these considerations. From a business perspective, compute resources – including database functionality – have become a commodity. Companies move their workloads increasingly toward cloud-based systems, raising the question whether the future of databases is also in the cloud.

A similar line of arguments leads to the concept of *database appliances*. Appliances package database functionality in a closed box, allowing *(a)* to treat the service as a commodity (business aspect) and *(b)* to tailor hard- and software of the appliance specifically to the task at hand, with the promise of maximum performance (technology aspect).

And, in fact, both concepts – cloud computing and appliances – may go well together. Cloud setups enable to control the entire hard- and software stack; large installations may provide the critical mass to include tailor-made (database) functionality also within the cloud.

6.17 Rethinking Productivity in Software Engineering

Organizers: Thomas Fritz, Gloria Mark, Gail C. Murphy, and Thomas Zimmermann

Seminar No. 17102

Date: March 5–8, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.19

Participants: Christian Bird, Pernille Bjørn, Marcos Borges, Duncan Brumby, Robert Feldt, Thomas Fritz, James D. Herbsleb, Christian Janssen, Ciera Jaspan, Katja Kevic, Andrew J. Ko, Gloria Mark, André Meyer, Gail C. Murphy, Emerson Murphy-Hill, Brad A. Myers, Christopher J. Parnin, Martin Robillard, Caitlin Sadowski, Federica Sarro, David C. Shepherd, Janet Siegmund, Margaret-Anne Storey, Christoph Treude, Marieke van Vugt, Stefan Wagner, Thomas Zimmermann

There is an ever-growing demand of software being built and a shortage of software developers to satisfy this demand, despite the immense growth in the number of professional software developers. To address this demand, industry and research are looking into understanding and improving the productivity of individual software developers as well as teams. A substantial amount of research has examined the meaning of software productivity over the past four decades. Much of this research introduces particular definitions of productivity, considers organizational issues associated with productivity, or is focused on specific tools and approaches for improving productivity. In fact, many of the seminal work on software productivity is from the 80s and 90s (Peopleware, Mythical Man-Month, Personal Software Process).

At the same time, software development has changed significantly over the past decades with the rise of agile development, distributed development, more rapid release cycles and the high fragmentation of today's work. Simultaneously the technology available to software engineers has improved with social coding tools like GitHub[44] and StackOverflow[45] and better IDEs. Furthermore, research communities, in particular the HCI and CSCW communities, have made significant advances in supporting knowledge workers to become more productive that one might be able to also transfer to software engineers.

The goal of this seminar was to rethink, discuss, and address open issues of productivity in software development and how to measure and foster productive behavior of software developers. Specifically, we focused on the following questions:

- What does productivity mean for an individual and teams/organizations and how is it measured?
- What are the dimensions and factors of productivity?
- What are the purposes and implications of measuring productivity?
- What are the grand challenges in research on productivity?

[44] http://www.github.com
[45] http://www.stackoverflow.com

6

6.18 Game Theory in AI, Logic, and Algorithms

Organizers: Swarat Chaudhuri, Sampath Kannan, Rupak Majumdar, and Michael J. Wooldridge

Seminar No. 17111

Date: March 12–17, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.27

Participants: Natasha Alechina, Eric Balkanski, Suguman Bansal, Patricia Bouyer-Decitre, Tomás Brázdil, Véronique Bruyère, Swarat Chaudhuri, Yiling Chen, Yuan Deng, Rayna Dimitrova, Rüdiger Ehlers, Kousha Etessami, Bernd Finkbeiner, Valentin Goranko, Julian Gutierrez, Paul Harrenstein, Hanno Hildmann, Justin Hsu, Rasmus Ibsen-Jensen, Nicole Immorlica, Wojtek Jamroga, Sampath Kannan, Jan Kretinsky, Antonin Kucera, Orna Kupferman, Martin Lange, Kim Guldstrand Larsen, Rupak Majumdar, Nicolas Markey, Daniel Neider, Evdokia Nikolova, Doron A. Peled, Sophie Pinchinat, Maria Polukarov, Jean-Francois Raskin, Sanjit A. Seshia, Arunesh Sinha, Tami Tamir, Moshe Y. Vardi, Igor Walukiewicz, Michael J. Wooldridge, Mihalis Yannakakis

The Dagstuhl Seminar *17111: Game Theory in AI, Logic, and Algorithms* was held from March 12–17, 2017. The seminar explored research challenges at the interface of computing and game theory. This area has seen fervent research activity in recent times. Specifically, game theoretic ideas have found currency in three key areas of computer science: in the *algorithms* community, algorithmic game theory is now a well-established sub-field; in *formal methods*, model checking and synthesis problems have been studied using game-theoretic concepts; and in *artificial intelligence*, game theory has come to provide the fundamental conceptual vocabulary for the field of multi-agent systems. Despite this manifest common interest, there is surprisingly little trade between game-theoretic approaches in these different subfields of computer science. Our aim in this seminar was to start to build some bridges between these three areas.

6.19 Using Networks to Teach About Networks

Organizers: Timur Friedman, Aiko Pras, and Jürgen Schönwälder

Seminar No. 17112

Date: March 12–15, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.33

Participants: Olivier Bonaventure, Alessio Botta, Stefan Burschka, Georg Carle, Marinos Charalambides, Pieter-Tjerk de Boer, Wouter de Vries, Jordi Domingo-Pascual, John Domingue, Timur Friedman, Gunnar Karlsson, Jean-Yves Le Boudec, Simone Opel, Marc-Oliver Pahl, Cristel Pelsser, Aiko Pras, Bruno Quoitin, Fabio Ricciato, Jürgen Schönwälder, Anna Sperotto, James P. G. Sterbenz, Burkhard Stiller, Iyad Tumar, Klaus Wehrle, Harald Weikert, Lisa Yan

Computer networks have become a common utility and the Internet provides new opportunities for education. In addition, we see an increasing deviation of the deployed Internet from the basic principles driving the design of computer networks. All this has an impact on how we educate young minds in computer networking and hence it is required to rethink how education in computer networking should be organized, which topics are essential to cover and which ones are merely nice illustrations of core concepts. Furthermore, it seems necessary to think about using the Internet itself more intensively to develop new educational materials. In order to start a discussion of such educational aspects, a Dagstuhl seminar titled Using Networks to Teach About Networks has been organized. Some questions discussed during the seminar were:

- Which topics should be taught in a typical undergraduate course? What are the essential basic principles that need to be understood? Which topics should be covered in a typical graduate course? How to deal with the fact that architectural concepts are often violated in real networks?
- How should topics be taught? How to best use the Internet for teaching how the Internet works? How can we more easily 'mesh' teaching materials? Can we better organize the sharing of video content, assignments, or experimental setups? Do we need an open source platform for teaching materials? What about open source books on computer networks replacing traditional textbooks?
- What is the experience with modern teaching styles, such as pure online courses like MOOCs or flipped classrooms? Which role should project work play? How can novel teaching ideas best be leveraged and integrated into existing educational concepts?

6

6.20 Computational Complexity of Discrete Problems

Organizers: Anna Gál, Michal Koucký, Oded Regev, and Till Tantau

Seminar No. 17121

Date: March 19–24, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.45

Participants: Eric Allender, Nikhil Bansal, Harry Buhrman, Igor Carboni Oliveira, Sourav Chakraborty, Gil Cohen, Anindya De, Pavel Dvorak, Lance Fortnow, Anna Gál, Alexander Golovnev, Rohit Gurjar, Kristoffer Arnsfelt Hansen, Prahladh Harsha, Johan Hastad, Pavel Hrubes, Valentine Kabanets, Swastik Kopparty, Michal Koucký, Matthias Krause, Bruno Loff, Shachar Lovett, Meena Mahajan, Or Meir, Shay Moran, Jakob Nordström, Ramamohan Paturi, Pavel Pudlák, Oded Regev, Rüdiger Reischuk, Michael E. Saks, Rahul Santhanam, Ronen Shaltiel, Avishay Tal, Till Tantau, Thomas Thierauf, Jacobo Torán, Dieter van Melkebeek, Ben Lee Volk, Heribert Vollmer

Introduction and goals

Computational complexity studies the amount of resources (such as time, space, randomness, or communication) that are necessary to solve computational problems in various models of computation. Finding efficient algorithms for solving computational tasks is crucial for practical applications and becomes even more important with the use of computers becoming part of everyday life. Despite a long line of research, for many problems that arise in practice it is not known if they can be solved efficiently – in particular in polynomial time.

Beside questions about the existence of polynomial time algorithms for problems like Satisfiability or Factoring where the best known algorithms run in exponential time, there is a huge class of practical problems where algorithms with polynomial running time (e.g. cubic or even quadratic time) are known, but it would be important to establish whether these running times are best possible, or to what extent they can be improved.

These fundamental questions motivate developments in various areas from algorithm design to circuit complexity, communication complexity and coding theory. During this Dagstuhl Seminar, we discussed some of the most exciting recent developments in those areas related to computational complexity.

The seminar "Computational Complexity of Discrete Problems" has evolved out of the series of seminars entitled "Complexity of Boolean Functions," a topic that has been covered at Dagstuhl on a regular basis since the foundation of this research center. An important feature of the current research in computational complexity is the integration of ideas from different subareas of computational complexity and from other fields in computer science and mathematics. We have aimed to attract researchers from various subareas connected to core questions in boolean function complexity and foster further fruitful interactions.

6.21 Mixed Criticality on Multicore / Manycore Platforms

Organizers: Liliana Cucu-Grosjean, Robert Davis, Sanjoy K. Baruah and Zoë Stephenson

Seminar No. 17131

Date: March 26–31, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.70

Participants: Kunal Agrawal, Sebastian Altmeyer, James H. Anderson, Sanjoy K. Baruah, Iain Bate, Enrico Bini, Björn B. Brandenburg, Alan Burns, Thidapat Chantem, Jian-Jia Chen, Liliana Cucu-Grosjean, Robert Davis, Arvind Easwaran, Pontus Ekberg, Sébastien Faucou, Madeleine Faugère, Christian Ferdinand, Laurent George, Adriana Gogonel, Sathish Gopalakrishnan, Emmanuel Grolleau, Zhishan Guo, Leandro Soares Indrusiak, Jaewoo Lee, Jing Li, Martina Maggio, Alberto Marchetti-Spaccamela, Cristian Maxim, Dorin Maxim, Mitra Nasri, Claire Pagetti, Kirk Pruhs, Gurulingesh Raravi, Jan Reineke, Stefan Resch, Philippa Ryan, Luca Santinelli, Zoë Stephenson, Sascha Uhrig, Wang Yi, Dirk Ziegenbein

Real-time applications are characterized by the need for both functional correctness and temporal correctness (appropriate timing behaviour). Real-time systems are present in many diverse areas such as avionics, automotive, space, robotics, and medical applications to cite only a few. Mixed Criticality Systems (MCS) have become an important topic for the real-time systems community. The first cluster of the European collaborative projects on MCS has been completed in September 2016, indicating a maturing of the related concepts within both industry and academia. Nevertheless many of the challenges brought about by the integration of mixed criticality applications onto multicore and manycore architectures remain to be solved. In reality mixed criticality problems have inherited the difficulty of real-time systems: being at the frontier of several domains including real-time scheduling, real-time operating systems / runtime environments, and timing analysis, as well as hardware architectures. This seminar promoted lively interaction, cross fertilization of ideas, synergies, and closer collaboration across different sub-communities of academics and industrialists from aerospace, automotive, and railway industries with specific interests in MCS, as well as with experts in certification.

In common with the first Dagstuhl Seminar on Mixed Criticality Systems, this seminar also focused on the two key conflicting requirements of MCS: separation between criticality levels for assurance and sharing for resource efficiency, along with the related requirement of time composability. An important aspect of this seminar was the presentation of different industry perspectives on the key problems. These perspectives formed the starting point of our seminar, with the first day mainly dedicated to industry statements on current practice and their perception of current work on MCS. The academic participants benefited from substantial and detailed arguments from the industry speakers. There were lively interactive discussions during the talks which led to much improved understanding of current industry practice, as well as helping to build a common vocabulary between academic and industry participants. The first day concluded with presentations by academic speakers presenting their thoughts on more practical mixed criticality models.

The next three days each included sessions devoted to an invited tutorial from a academic speaker. These covered the one-out-of-m multicore problem, Networks-on-Chip and mixed criticality, resource management, and statistical approaches to worst-case execution time estimation. The remaining sessions covered a range of fascinating open problems. In addition, a number of ad-hoc small working groups formed to collaborate on specific topics. We are pleased to report that a significant number of these initial collaborations have gained traction resulting in further work after the seminar, and in some cases the development and submission of papers.

Organization of the seminar report. Section 3 of the full report is an overview of the industry talks and Section 4 of the full report provides an overview of the academic talks. Section 5 of the full report presents working group discussions. Section 6 of the full report summarizes open problems discussed during the seminar. Finally outcomes from the seminar are listed in Section 7 of the full report.

As organizers, we would like to thank Prof. Reinhard Wilhelm for joining us, Dagstuhl's Scientific Directorate for allowing us to run a second seminar on mixed criticality systems, and to the staff at Schloss Dagstuhl for their superb support during the seminar itself.

Finally, we would like to thank all of the participants for the very lively and open discussions. As organizers, we appreciated the feedback and enthusiasm which made running the seminar a great pleasure.

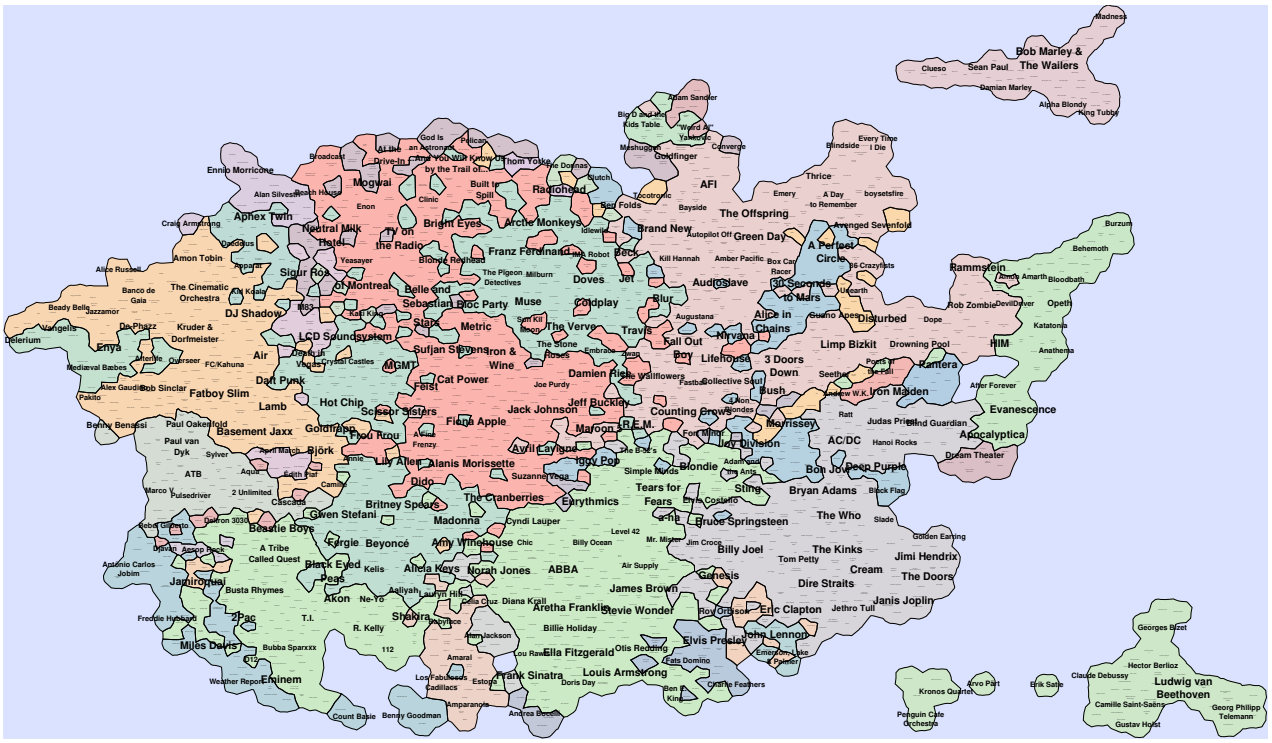

Fig. 6.3

"A virtual map of the music highlighting sets artists and how they relate to each other, using a technique called gvmap."

Press release about Dagstuhl Seminar 17332 – "Scalable Set Visualizations".

https://www.dagstuhl.de/no_cache/en/about-dagstuhl/press/pressemitteilungen/detail/meldung/664/. Photo courtesy of Yifan Hu.

6.22 Opportunities and Risks of Blockchain Technologies

Organizers: Roman Beck, Christian Becker, Juho Lindman, and Matti Rossi

Seminar No. 17132

Date: March 26–29, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.3.99

Participants: Pär Ågerfalk, Michel Avital, Roman Beck, Christian Becker, Joseph Bonneau, Marcus Dapp, Peter Eklund, Fritz Henglein, John Leslie King, Christoph Kreiterling, Juho Lindman, Alberto Montresor, Christoph Müller-Bloch, Matti Rossi, Joachim Schrey, Gerhard Schwabe, Peter Sestoft, Virpi Tuunainen, Marella Venkata, Roger Wattenhofer, Jesse Yli-Huumo

Introduction

The Dagstuhl seminar "Opportunities and Risks of Blockchain Technologies" had 21 participants from universities, public institutions, and enterprises. Blockchain is both an information technology as well as an economic innovation. As a technical innovation it is a new version of a distributed transactional database technology, especially suited for decentralized environments of limited or imperfect trust. As an economic innovation it offers novel tools to any problem domain where there exists a need for a reliable record of transactions in a decentralized environment where not all parties, whether humans or machines, can be fully trusted.

Computer scientists have researched key issues of blockchain technologies such as technical availability, tools, standards, and applications that enable these networks. Our seminar aimed to bridge the gap between this research stream and research perspectives from Service Science, Wirtschaftsinformatik, and Information Systems. We brought together a multi-disciplinary group of academic and industry researchers; specifically those working in fields such as open platforms, open source, distributed trust platforms, cryptocurrency tools, as well as the related social and legal challenges.

We set to analyze and synthesize the current body of knowledge on the emerging landscape of blockchain technologies. We linked the emerging phenomenon of cryptographic economic systems to already established research streams around trust-related issues in payment systems, online currencies, and supply chain management through group work and keynotes. We worked on four theme groups:

1. Research centers
2. Blockchain and Fintech
3. Essence and future of blockchain technologies
4. Impact/changing institutions

In the following we look at each of these shortly. The full report contains a number of position papers that explore these issues in further detail.

Research centers

The research center work group sought ways of strengthening the European and global research on blockchain. A starting point was a proposal to form a network of similar minded Blockchain experts and research groups across Europe. Several groups from countries such as in Denmark, Ireland, UK, and Switzerland could start as a loosely coupled interest group to work on potential research agendas and teaching curricula. The Blockchain seminar at Dagstuhl can be regarded as the starting point for the formation of the research network. Based on this network, the next step would be to convince funding agencies and industry to write research proposals for the DRAO (Distributed Research Autonomous Organization). The idea of DRAO is that the Blockchain research center should not be just another research center, doing research on Blockchain, but actually should be based on Blockchain, as a distributed autonomous research organization.

Furthermore, the group discussed a proposal for Blockchain teaching and education. This would result in a suite of courses on various areas, possibly as follows:

Computer Science Foundations

- Cryptography, authentication and signature methods.
- Distributed computing, distributed algorithms, understanding of the tradeoffs, consensus protocols
- Distributed systems
- Domain specific languages for contracts and for protocols.
- Large scale software engineering for distributed ledger development, software engineering

- Program analysis and software quality, objective way of verifying properties

Information Systems Economics
- Economic theories on incentive models, auctions and mechanism design (basically game theory insights)
- Inter-organizational, distributed governance and management theories
- Collective economies, reputation and trust management
- Digital Mindset and management of digital personas
- Ethics and critical reflections of Blockchain and societal implications

Information Systems Management and Organization
- Innovation Design: from Blockchain idea to prototype
- Digital Entrepreneurship: from Blockchain prototypes to markets
- Taxation, auditing and integration of Blockchain in organizations

IT Law
- Legal implications of Blockchain, property rights, ownership, responsibilities

Blockchain and Fintech

Our second worgroup dealt with the relationship between Blockchain and the Fintech industry. The group produced a matrix of different financial and legal functions, tools required to handle those and the potential of Blockchain to replace solutions to these functions.

Essence and Future of Blockchain Technologies

The workgroup on the essence of blockchain technologies set out to understand what forms the core of the technology. Its preliminary definition is that a blockchain implementation should contain First, the data storage that implements a distributed ledger system (DLS), the actual Blockchain, which is the data structure used in DLS is a hash-linked chain of blocks. A block is a collection of transactions that form the ledger. Furthermore, a Consensus Mechanism allowing for (de-)Centralization of power to decide which transactions are valid in the network. The innovative combination of the above mentioned three components give DLS interesting characteristics that we describe in the next section.

Impacts/changing institutions

This workgroup set out to understand the relationship between technical change and social change. It tried chart the relationship of blockchain and institutions. The discussion centered around the resilience of institutions and the need for stop gap measures, which are often provided by quite traditional public infrastructures and legal frameworks. In some cases, institutions will have to put conditions in place to allow blockchain to work (in particular to avoid harm). The idea of Blockchain being able to replace or eliminate trust was a central topic and it was noted that this can be an issue in cases of fraud (e.g. Ethereum fork as an example). The group also discussed who provides stable identifiers and who decides what can be stored in a given ledger. Similarly, the assignment of value and ownership and their control remain important issues that are now seen as tertiary to the technology. Key questions arising from this were:
- How will blockchain solutions that work well in theory or as prototypes function when used as large-scale solutions?
- Are certain groups or communities better suited to adopt blockchain?

Final comments

We believe that despite some hiccups (e.g., DAO fork) blockchain will emerge as an important technological and economical phenomenon. Its key properties and impacts should be studied intensively to allow for new innovations in the financial sector and other areas, where the technology's affordances promise to create value. The work continues through a manifesto in Business & Information Systems Engineering, a viewpoint in Communications of the ACM, and a special issue in the Journal of the Association for Information Systems.

6.23 Probabilistic Methods in the Design and Analysis of Algorithms

Organizers: Bodo Manthey, Claire Mathieu, Heiko Röglin, and Eli Upfal

Seminar No. 17141

Date: April 2–7, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.4.1

Participants: Dimitris Achlioptas, Aris Anagnostopoulos, Luca Becchetti, Petra Berenbrink, Amin Coja-Oghlan, Artur Czumaj, Daniel Dadush, Martin Dietzfelbinger, Benjamin Doerr, Devdatt Dubhashi, Matthias Englert, Leslie Ann Goldberg, Paul W. Goldberg, Ross Kang, Thomas Kesselheim, Stefan Klootwijk, Robert Krauthgamer, Marvin Künnemann, Stefano Leonardi, Frederik Mallmann-Trenn, Bodo Manthey, Claire Mathieu, Nicole Megow, Morteza Monemizadeh, Seffi Naor, Alessandro Panconesi, Andrea Pietracaprina, Kirk Pruhs, Geppino Pucci, Matteo Riondato, Heiko Röglin, Thomas Sauerwald, Melanie Schmidt, Christian Sohler, Aravind Srinivasan, Marc Uetz, Eli Upfal, Tjark Vredeveld, Philipp Woelfel

Probabilistic methods play a central role in theoretical computer science. They are a powerful and widely applied tool used, for example, for designing efficient randomized algorithms and for establishing various lower bounds in complexity theory. They also form the basis of frameworks like average-case and smoothed analysis, in which algorithms are analyzed beyond the classical worst-case perspective. The seminar was on probabilistic methods with a focus on the design and analysis of algorithms.

Probabilistic methods are often used in algorithm analysis when worst-case analysis does not provide useful or empirically accurate results. For example, worst-case analysis suggests that the simplex method is an exponential-time algorithm for linear programming, while in fact it runs in near-linear time on almost all inputs of interest. For the simplex method and many other algorithms worst-case inputs are often rather contrived and occur hardly ever in practical applications. The last decade has seen much interest in the development of a more realistic and robust algorithmic theory that is not entirely based on worst-case performance. One very successful line of research studies the performance of algorithms on inputs that are to some extent random. Besides average-case analysis, in which inputs are generated randomly according to some fixed distribution, also more sophisticated semi-random models have gained momentum.

Another area in which probabilistic methods play a central role is stochastic optimization. Here uncertainty in the data is modeled by probability distributions and the actual data is only revealed over time. For example, in a scheduling problem one might know the probability distribution of a job's length but one learns its actual length only by executing it.

Probabilistic methods are also central in algorithm design. For many optimization problems, the most efficient known algorithms rely essentially on randomization. In other areas, like sublinear algorithms and hashing, one can even prove that randomization is necessary to obtain good algorithms.

The seminar covered recent progress in the context of probabilistic methods.

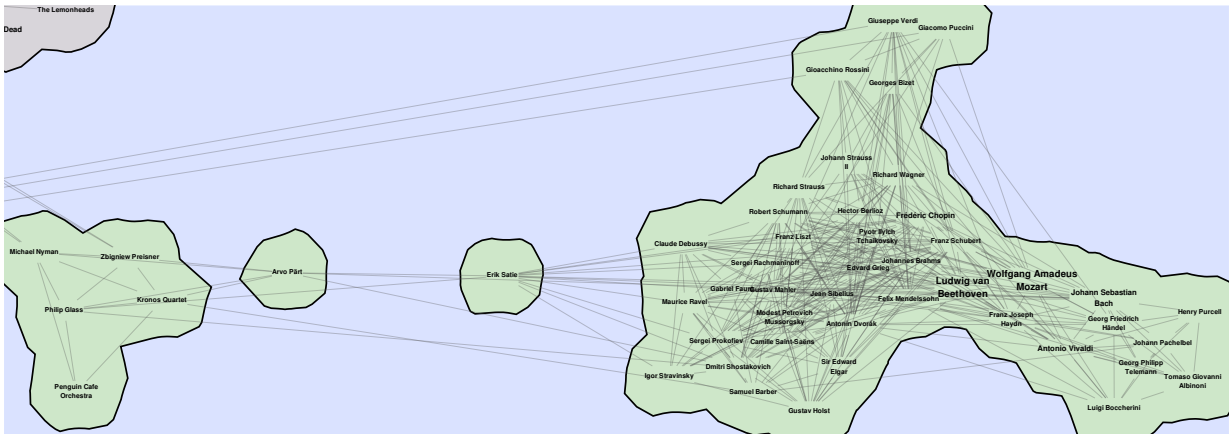

Fig. 6.4
"Zoom in to the classic music area."
Press release about Dagstuhl Seminar 17332 – "Scalable Set Visualizations".
https://www.dagstuhl.de/no_cache/en/about-dagstuhl/press/pressemitteilungen/detail/meldung/664/. Photo courtesy of Yifan Hu.

6.24 Formal Methods of Transformations

Organizers: Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl

Seminar No. 17142

Date: April 2–5, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.4.23

Participants: Johanna Björklund, Adrien Boiret, Mikolaj Bojanczyk, Michaël Cadilhac, Olivier Carton, Loris d'Antoni, Luc Dartois, Frank Drewes, Emmanuel Filiot, Olivier Gauwin, Bruno Guillon, Florent Jacquemard, Ismaël Jecker, Aurélien Lemay, Nathan Lhote, Christof Löding, Andreas Maletti, Sebastian Maneth, Anca Muscholl, Keisuke Nakano, Joachim Niehren, Chih-Hao Luke Ong, Raphaela Palenta, Rogério Reis, Pierre-Alain Reynier, Jacques Sakarovitch, Helmut Seidl, Frédéric Servais, Krishna Shankaranarayanan, Jean-Marc Talbot, Margus Veanes

The Dagstuhl seminar 17142 "Formal Methods of Transformations" was a short two and half day seminar that took place from April 3rd to 5th, 2017. The aim of this seminar was to bring together researchers working on theory and applications of formal models of transformations (also known as transductions) of strings and trees. A model of transformation which has been central in this seminar is that of a transducer, i.e., an automaton extended with output. Transducers were introduced in the 1960s as formal models for linguistics and syntax-directed translation in compilers. Today, research on string and tree transducers is an active field with various new applications in databases, document processing, natural language processing, software engineering, and verification. To make the seminar more focused, we had identified six research directions as key topics for the seminar:

- Canonical Normal Forms: it is well-known that regular languages admit a unique minimal (and canonical) deterministic finite automaton. Similarly, is it possible to characterize classes of transformations by canonical normal form for transducers?
- Transducer-Logic Relationships: there are well-known automata-logic correspondences in the theory of languages. What is known for transformations and can we obtain similar connections?
- Subclass Definability: a fundamental question, which requires a deep understanding of the manipulated objects, is that of subclass definability: given (an effective description of) an object in some class C, does this object belong to a given subclass C' of C? What are the recent results and open problems with respect to subclass definability for transformations?
- New Decidability Results: what are the recent decidability breakthroughs in the theory of transformations and what are the important open problems?
- New Transducer Models: is there a need for new transducer models tailored to new applications?
- Model-Checking Data-Centric Systems: what are the potential applications of transducer theory to the verification of systems that transform data?

The seminar gave a large overview of recent results and open problems with respect to these research directions. It was a follow-up of Dagstuhl seminar 13192 "Tree Transducers and Formal Methods", which now included researchers on string transducers. String transducers have indeed received a lot of attention in the recent years and new important results have been obtained. The aim of this seminar was also to gather researchers from the string and tree transducer communities. There were 31 participants from 10 countries (Sweden, France, Poland, Germany, US, Belgium, UK, Japan, Portugal, and India). These participants were invited by the organizers Emmanuel Filiot, Sebastian Maneth, and Helmut Seidl to give survey talks and shorter talks on their current research.

The seminar started with a survey talk by Emmanuel Filiot on recent results on string transformations, intended to motivate the need to address the theory of string transformations in this seminar, and to introduce some of its main recent breakthroughs. Mikolaj Bojanczyk then gave a survey talk on the notion of *origins* in transformations. Origins are inherent to all known transducer models and making them explicit in the semantics of transducers gave rise to new decidability and definability results. Mikolaj's talk was followed by a session on string transducers. Luc Dartois presented a new and expressive logic to define string transformations. The logic offers good decidability properties with respect to satisfiability and equivalence (under the origin semantics). String transformations, under the origin semantics, can be seen as sets of graphs, called origin graphs. Bruno Guillon presented a characterization of the class of origin graphs

generated by known string transducer models. Finally, Jacques Sakarovitch closed this session by presenting a contribution, and an open problem, on rational base numeration systems and their analysis by means of automata and transducers.

The afternoon session started with a survey talk by Christof Löding on the automatic synthesis of deterministic transducers of strings and trees, from specifications given by non-deterministic transducers. This problem, called the uniformization problem, can be seen as a variant of the classical Church synthesis problem. Two short talks given by Nathan Lhote and Michaël Cadilhac presented their recent contributions on definability problems for string transducers. Nathan Lhote presented decidability results for checking whether a given rational function is first-order definable, and, more generally, whether it is definable by some string transducer whose (input) transition monoid (which disregards the outputs) belongs to some given class, thus lifting to transformations well-known results from the theory of regular languages. Michaël Cadilhac then introduced the notion of C-continuity for string transformations and showed effectiveness of this notion for particular classes C. This notion gives an alternative and machine-independent way, which also takes outputs into account, of defining meaningful subclasses of transformations.

Day 2 (morning) was a session devoted to applications. It started with a survey talk on symbolic transducers, by Margus Veanes. Symbolic transducers have been introduced to address practical issues when dealing with very large alphabets (even infinite). Transition labels are replaced by predicates defining sets of allowed labels. Fundamental decidability results and practical applications were presented. Based on symbolic string transducers, Loris d'Antoni then presented a fully automatic method to invert a practical class of functional (and injective) transformations, implemented in the tool GENIC. Adrien Boiret then addressed the problem of checking equivalence for symbolic top-down tree transducers. Keisuke Nakano presented some fundamental results on B-terms, which model function composition. Finally, the morning session ended with a talk by Olivier Carton about the notion of infinite word compression by transducers and its relation to normality.

Andreas Maletti started the afternoon session with a survey talk on tree transducers and their applications in linguistics. He presented transducer models well-suited to some linguistic applications such as automatic natural language translation, and showed a comparison with deep learning approaches. Frank Drewes then discussed the notion of graph transformations, and especially of DAGs, and proposed in his talk a definition of DAG transducers. Motivated by natural-language interface applications, Johanna Björklund presented a study of the expressivity of checking stack transducers, which extend with outputs the well-known model of checking stack automata. Day 2 ended with a talk by Helmut Seidl on a result obtained together with Sebastian Maneth and Gregor Kemper on the decidability of equivalence for deterministic top-down tree-to-string transducers, which had been a long standing open problem.

The last day was dedicated to results on the class of string transformations defined by two-way transducers, mostly definability problems. There were two talks, one by Olivier Gauwin and one by Jean-Marc Talbot, presenting two different techniques, for deciding whether a transformation in this class can be defined by a one-way transducers, i.e., is rational. It was open whether this problem was solvable in elementary time, and the two presented approaches answer this question positively. Another definability problem about the minimization of the number of registers in subclasses of SSTs (a model introduced by Alur and Cerny, with the same expressivity as two-way transducers), was discussed by Pierre-Alain Reynier. Ismaël Jecker presented a class of transducers with strong constraints on their structures, but still as expressive as too-way transducers, and at the same time enjoying many good algorithmic properties, in particular improving by one exponential a famous result by Hopcroft and Ullman. Finally, Krishna S. concluded the seminar by sketching an alternative proof of a result by Alur et. al. on regular expressions for string transformations.

We had long after-lunch breaks (till 4pm) to give the opportunity to the participants to discuss. It was greatly appreciated by the participants, and some of them initiated new collaborations. The discussions inspired new ideas and we hope that joint papers will be published by the participants.

We warmly thank Schloss Dagstuhl for making this fruitful event possible, and for their help in the organization. It is highly appreciated as organizers, and allowed us to focus only on the scientific aspects of the seminar.

6.25 Ambient Notification Environments

Organizers: Lewis Chuang, Sven Gehring, Judy Kay and Albrecht Schmidt

Seminar No. 17161

Date: April 17–20, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.4.38

Participants: Elisabeth André, Susanne Boll, Stephen Brewster, Barry Brown, Andreas Bulling, Andreas Butz, Keith Cheverst, Lewis Chuang, Raimund Dachselt, Anind K. Dey, David Dobbelstein, Christoph Gebhardt, Sven Gehring, Hans Gellersen, Tom Gross, Niels Henze, Judy Kay, Pascal Knierim, Thomas Kosch, Bob Kummerfeld, Markus Löchtefeld, Evangelos Niforatos, Tadashi Okoshi, Joseph A. Paradiso, Veljko Pejovic, Benjamin Poppinga, Aaron Quigley, Harald Reiterer, Michael Rohs, Albrecht Schmidt, Aurélien Tabard, Pauline Vierne, Alexandra Voit, Frederik Wiehr

Reports indicate that many users interact with their smartphone and wearable devices more than 100 times per day. Oftentimes, these interactions result from direct notifications, presented on the screen or via sounds. New communication applications (e.g. from email to WhatsApp) has increased the frequency of notifications, while social media applications are inherently motivated to entice repeat visits and interactions. The end result is that more and more systems as well as applications compete for the user's attention using notifications. Projecting this into the future, it is apparent that current implementation schemes that rely on direct notifications will not scale. A simple extrapolation of the rate in notifications suggests that a near future whereby users will only attend to notifications with no time leftover for productive work. Therefore, a radical restructure of notification delivery is necessary – specifically, one that keeps the user in-the-loop without consuming all of the user's attention. This is a key challenge. If no alternatives to direct notifications can be realized, current visions of ubiquitous computing and smart environments are likely to be unrealistic, given their anticipated undesirability to the end-user. In the Dagstuhl Seminar 17161 "Ambient Notification Environments", we looked at how novel approaches to notification delivery can address the above issue. We brought together researchers and experts that understand the technical, psychological, and social aspects of notification systems. This facilitated a broad discussion that was fuelled by the discussants' joint expertise in mobile and smart home technologies, in ambient sensing and presentation, and psychological models of human attention, to name a few. This discussion brought to the foreground, many of the underlying challenges that deserves further research. On the hand, a technological push in novel communication and smart devices drives an immediacy in user interaction. On the other hand, this will result in a higher demand on human attention. In the current report, we document the diverse approaches that were proposed as an alternative means towards notifications that are more personalized and contextualized, with the express purpose of reducing user effort. The central aim of the seminar was to understand the challenges and questions that we face when designing future interactive systems and to overcome the mounting notification problem. In addition to the ideas raised in the individual presentations and the group sessions, we jointly identified the following research questions and challenges:

1. How can notifications be designed to be simultaneously non-intrusive and yet noticeable?
2. How can artificial intelligence be designed to provide effective context-aware notifications?
3. What is a suitable notification architecture for integrating user devices, smart (shared) environments, and personal data without compromising personal privacy?
4. What is a suitable conceptualization for notification systems and how will a taxonomy for classifying notification delivery that is centered on user preferences and privacy look like?

Overall, the seminar contributed to a common vision of how interactions between human and systems can and ought to progress – one where technological progress does not necessitate an ever growing burden on the user's attention. In this report, we balance an analysis of the underlying drivers and problems, ideas for novel conceptual and technical approach, and most importantly a set of questions and research challenges in this domain that merit the attention of like-minded researchers.

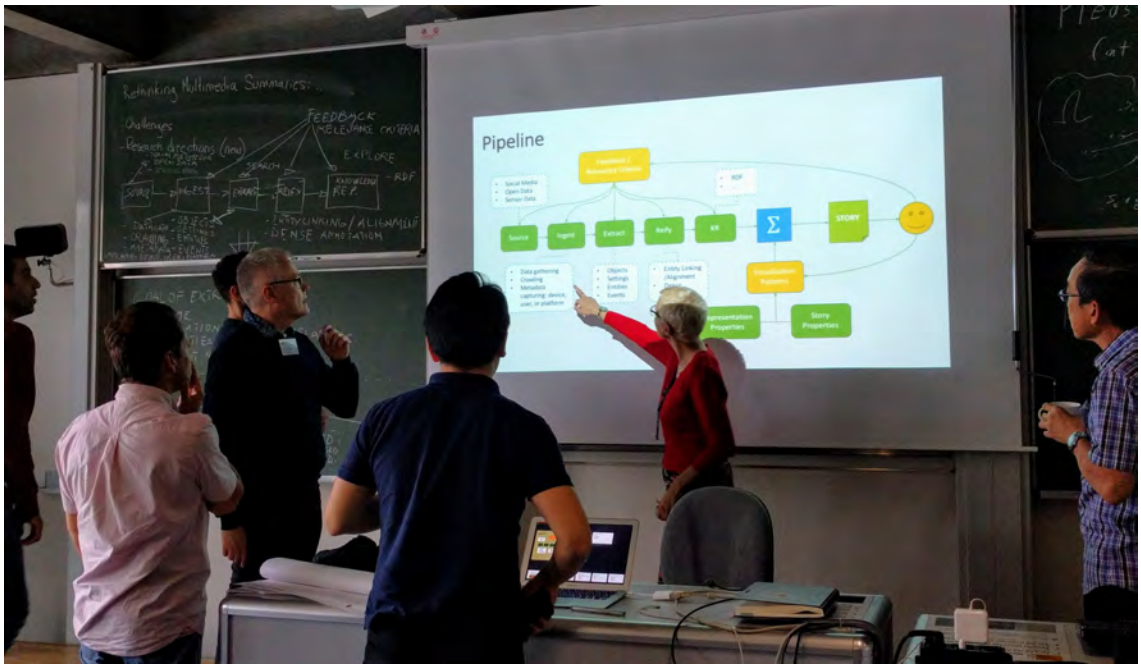

Fig. 6.5
"Rethinking the requirements for #multimedia & #multilingual #summarisation and #storytelling - excellent teamwork at @dagstuhl seminar!" Twitter post by 17301 Dagstuhl Seminar participant Tatjana Gornostaja.
https://twitter.com/TermServ/status/891971467965063168. Photo courtesy of Tatjana Gornostaja.

6.26 Online Privacy and Web Transparency

Organizers: Nataliia Bielova, Nikolaos Laoutaris, Arvind Narayanan, and Nick Nikiforakis

Seminar No. 17162

Date: April 17–20, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.4.83

Participants: Günes Acar, Benoit Baudry, Nataliia Bielova, Joe Calandrino, David Choffnes, Damian Clifford, Angel Cuevas Rumin, Steven Englehardt, Paul Francis, Krishna P. Gummadi, Aniko Hannak, Costas Iordanou, Nicolas Kourtellis, Balachander Krishnamurthy, Nikolaos Laoutaris, Arnaud Legout, Timothy Libert, Arvind Narayanan, Nick Nikiforakis, Lukasz Olejnik, Josep M. Pujol, Konrad Rieck, Dolière Francis Somé, Oleksii Starov, Vincent Toubiana, Michael Carl Tschantz, Peggy Valcke, Robbert van Eijk, Diana Vlad-Calcic

The Dagstuhl Seminar on *Online Privacy and Web Transparency* was the first seminar at Dagstuhl that gathered together researchers working in web applications, online privacy, transparency on the web, privacy enhancing technologies, privacy measurement, and network economics, as well as several representatives of law and policy discipline.

Research context

The web has become an essential part of our society and is currently the main medium of information delivery. Billions of users browse the web on a daily basis, and there are single websites that have reached over one billion user accounts. In this environment, the ability to track users and their online habits can be very lucrative for advertising companies, yet very intrusive for the privacy of users.

Recent research has shown that third-party advertising networks and data brokers use a wide range of techniques in order to track users across the web – these techniques are used to reconstruct browsing sessions and to create profiles of users, inferring, among others, their hobbies, health status, political inclinations, and level of wealth. This information can be used to, not only deliver better targeted advertisements to users, but also to discriminate users, for example by providing customized prices for products based on a user's willingness and ability to pay. To protect users, several solutions have been proposed, ranging from the laws, policies and the W3C Do-Not-Track candidate recommendation, to browser tools developed by companies and volunteers, and other client-side mechanisms proposed by researchers. At the same time, a number of tools have been developed to increase transparency on the web and allow end users to know when they are being tracked and when discrimination happens.

The seminar aimed to address the open questions of how to protect user privacy and how to increase transparency on the web. The key objectives of the seminar are (i) review the state of the art in the field; (ii) identify key technical challenges and brainstorm potential solutions; (iii) understand how computer science research results can influence law and policy; (iv) discuss ethical and legal issues in privacy research.

The seminar brought together scientists from the privacy and transparency communities, as well as policy makers interested in understanding how existing privacy laws and policies can be implemented, and representatives of Internet users organisations. The discussions at this Dagstuhl Seminar were strongly inspired by the following questions and challenges:

Technology

- How can we detect tracking and algorithmic discrimination most effectively? What are the scientific and engineering challenges to overcome? What are the relative merits of automated, semi-automated, and crowdsourced approaches?
- How can we ensure that methodologies, techniques, and tools are shared across different communities working on this topic?
- How can we design the next generation of privacy tools, get users to actively use the tools, and generate data for privacy researchers to scrutinize?
- What are the tracking techniques and data collection practices on mobile devices and how do they compare to those on the web?
- What are the privacy and transparency issues raised by the Internet of Things, and how do we address them?

Law

- Do the current laws and policies cover existing tracking technologies? What is the process for reporting newly discovered tracking techniques to the appropriate Data Protection Authorities?
- Even with the appropriate legislation in place, how can we ensure that companies comply with the law? What can researchers do to help enforce compliance?

Ethical issues

- What is the best way to conduct web privacy research taking ethical issues into account?
- When does a study necessitate ethical review?

Users

- What is the most efficient way to raise user awareness about web tracking and transparency tools?

6.27 Computational Geometry

Organizers: Otfried Cheong, Anne Driemel, and Jeff Erickson

Seminar No. 17171

Date: April 23–28, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.4.107

Participants: Mohammad Ali Abam, Tetsuo Asano, Sang Won Bae, Bahareh Banyassady, Ulrich Bauer, Karl Bringmann, Otfried Cheong, Man-Kwun Chiu, Jinhee Chun, Kenneth L. Clarkson, Michael Gene Dobbins, Anne Driemel, Jeff Erickson, Esther Ezra, Sándor Fekete, Kyle Jordan Fox, Jie Gao, Joachim Gudmundsson, Anna Gundert, Dan Halperin, David G. Kirkpatrick, Michal Kleinbort, Christian Knauer, Matias Korman, Amer Krivosija, Maarten Löffler, Tillmann Miltzow, Joseph S. B. Mitchell, Wolfgang Mulzer, Aleksandar Nikolov, Zuzana Patáková, Jeff M. Phillips, Benjamin Raichel, Abhishek Rathod, Saurabh Ray, Marcel J. M. Roeloffzen, Michael Sagraloff, Alejandro Salinger, Lena Schlipf, Raimund Seidel, Fabian Stehn, Monique Teillaud, André van Renssen, Kevin Verbeek, Bei Wang, David P. Woodruff

Computational Geometry

Computational geometry is concerned with the design, analysis, and implementation of algorithms for geometric and topological problems, which arise naturally in a wide range of areas, including computer graphics, CAD, robotics, computer vision, image processing, spatial databases, GIS, molecular biology, sensor networks, machine learning, data mining, scientific computing, theoretical computer science, and pure mathematics. Computational geometry is a vibrant and mature field of research, with several dedicated international conferences and journals and strong intellectual connections with other computing and mathematics disciplines.

Seminar Topics

The emphasis of the seminar was on presenting recent developments in computational geometry, as well as identifying new challenges, opportunities, and connections to other fields of computing. In addition to the usual broad coverage of new results in the field, the seminar included broad survey talks on monitoring and shape data and on high-dimensional geometric computing, two focus areas that have seen exciting recent progress and that present numerous opportunities for further cross-disciplinary impact.

Computational geometry for monitoring and shape data

The combination of movement and geometry has always been an important topic in computational geometry, initially motivated by robotics and resulting in the study of kinetic data structures. With the advent of widely available location tracking technologies such as GPS sensors, trajectory analysis has become a topic in itself, which has connections to other classical topics in computational geometry such as shape analysis. Still, efficient technologies to perform the most basic operations are lacking. We need data structures supporting similarity queries on trajectory data and geometric clustering algorithms that can handle the infinite-dimensional geometry inherent in the data. A related type of data, namely time series data, has not received much attention in the computational geometry community, despite its universality and its close relation to trajectory data. Shedding light on the interconnections of these topics will promote new results in the field which will address these timely questions.

Computing in high-dimensional and infinite-dimensional spaces

The famous "curse of dimensionality" prevents exact geometric computations in high-dimensional spaces. Most of the data in science and engineering is high-dimensional, rendering classical geometric techniques, such as the sweepline approach, insufficient. One way to address this issue is to use sparsity, but it is not always easy to find a sparse representation of the data. The search of the most efficient representation and how to exploit this representation leads to dimension-reduction techniques, metric embeddings, and approximation algorithms. This line of research has strong ties to machine learning and discrete mathematics as well as computational geometry.

Participants

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

fostering collaboration, and exposing young talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry. The organizers held a lottery for the third time this year; the lottery allows to create space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar has now a more balanced attendance in terms of seniority and gender than in the past. This year, 47 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive. No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar. We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, Schloss Dagstuhl itself is a great strength of the seminar.

6.28 Theory and Applications of Hashing

Organizers: Martin Dietzfelbinger, Michael Mitzenmacher, Rasmus Pagh, David P. Woodruff, and Martin Aumüller

Seminar No. 17181

Date: May 1–5, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.1

Participants: Alexandr Andoni, Martin Aumüller, Hannah Bast, Djamal Belazzougui, Vladimir Braverman, Andrei Z. Broder, Tobias Christiani, Artur Czumaj, Søren Dahlgaard, Martin Dietzfelbinger, Stefan Edelkamp, Tom Friedetzky, Andreas Goerdt, Bernhard Haeupler, Torben Hagerup, Michael Kapralov, Mathias Bæk Tejs Knudsen, Ravi Kumar, Yi Li, Andrew McGregor, Michael Mitzenmacher, Moni Naor, Jelani Nelson, Rasmus Pagh, Konstantinos Panagiotou, Ely Porat, Ilya Razenshteyn, Peter Sanders, Tamás Sarlós, Ludwig Schmidt, Anshumali Shrivastava, Francesco Silvestri, Christian Sohler, Tatiana Starikovskaya, Mikkel Thorup, Stefan Walzer, Udi Wieder, Philipp Woelfel, David P. Woodruff, Qin Zhang

Background

The idea of hashing was proposed in the 1950s as an efficient method for implementing symbol tables in compilers. In succeeding decades it has emerged as an algorithmic tool that goes well beyond its original purpose, providing solutions for a wide range of algorithmic problems. While the theory of hashing may appear mature, in fact, many new advances have been made in recent years. Also, the number of applications has grown to an extent that few people realize how broad the reach of hashing is, or have a comprehensive overview. The aim of this seminar was to bring together researchers with an interest in hashing methods (and more generally random mappings) from various perspectives, spanning theory and a diverse set of application areas. In this way we wanted to identify opportunities for further advancing the field.

Theory The theoretical side includes aspects such as the design of hash functions, the design of algorithm and data structure primitives, and the mathematical analysis of hashing schemes. For hash function design, since Carter and Wegman proposed universal hashing there has been a fertile research agenda identifying sufficient randomness properties of hash functions for various applications, and on the other hand devising resource-efficient hash functions having these properties. While new simple and efficient hash function constructions with strong theoretical properties keep appearing (seminar participants have contributed to this), there are still many applications of hashing for which no efficient hash function construction is known. At the same time it is increasingly clear that traditional measures of hash function performance like collision probability and independence are inadequate to describe the randomness properties needed in many applications.

While hashing is interesting in its own right, it has also become a foundational building block for higher level algorithms and data structures, each of which can in turn often find uses in a variety of application spaces. Various hash-based sketches laid the groundwork for the field of streaming algorithms, by showing how approximate counting could be done effectively. More recently, hashing algorithms have provided frameworks for similarity measures for text and images, data reconciliation, and even fast sparse Fast Fourier transform algorithms. As we construct richer, more elaborate structures on top of the hashing building blocks, we stretch what we require from hashing further.

Mathematical analysis of hashing schemes, an area enriched by early work of Knuth but grown well beyond, has inspired the development of a number of combinatorial methods and results. In the 1990s it was realized that load balancing using the best of two (or more) random choices leads to a much more even distribution, often referred to as the "power of two choices". Another great success in this area, obtained during the last decade, has been the analysis of cuckoo hashing; several seminar participants were involved in this. On the other hand, as questions become answered, new questions arise; in the case of cuckoo hashing, the effectiveness of variants including double hashing and partial-key cuckoo hashing are not understood. Beyond the hashing schemes themselves, data structures and algorithms that make use of lower-level hashing primitives further require and benefit from mathematical analysis. The fundamental connections between hashing, sparse reconstruction algorithms, and various sketching approaches are only now starting to be realized.

Applications Hashing is of course heavily used in information storage and retrieval contexts. For example, the "power of two choices" paradigm(where several seminar participants were among the pioneers) has resulted in extremely scalable and robust architectures for distributed hash tables (also known as key-value stores).

Other applications of hashing are appearing at a tremendous rate, as systems-designers and builders become accustomed to a world where approximate answers (as opposed to exact answers) are not only sufficient, they are necessary for efficiency. Indeed, hashing was one of the key methodologies for handling big data well before "big data" was even a widely used term. Since the seminal paper of Flajolet and Martin that showed how to efficiently compute an approximate count of the number of elements in a data stream, hashing has been a central tool in the design of algorithms for data streams, where the inability to store all the information that passes requires approximations. But in recent years the use of hashing has spread to many other settings where data is stored and accessible, but scalability can be achieved only by resorting to approximation algorithms similar to those developed in the setting of data streams. One early success story in this direction is the method for identifying near-duplicate web pages in Altavista using min-wise hashing. Another is HyperANF, a refined version of the Flajolet-Martin method that was used in 2012 to compute the distance distribution of the Facebook social network, making it by far the largest scale Milgram-like experiment ever performed. Finally, Bloom filters, invented around 1970 to provide a small-memory approximate representation of a set, have become a staple in systems, with countless variations expanding on its initial functionality, such as counts associated with set elements or aging out of set items as new elements are dynamically added.

In the field of machine learning, random mappings of data to lower-dimensional vectors that are easier to handle is of increasing importance for big data applications. This is true in particular since machine learning algorithms often work with kernelized feature vectors whose dimension far exceeds the size of data vectors. Designing randomized mappings that meet the criteria of machine learning applications has been an active research area in recent years. NIPS 2014 awarded one of two best paper awards a paper co-authored by seminar participant Shrivastava that presents a new asymmetric locality-sensitive hash function design for machine learning applications and shows how it leads to significant speedups.

The rich interplay between theory and practice in the field of hashing cannot be overstated. Applications drive the need for new algorithms and data structures based on hashing, as well as the need for more efficient classes of hash function families with provable theoretical guarantees. Specific implementations developed in the field often do not have theoretical guarantees on performance or accuracy, creating new theoretical problems and driving a need to make theory and practice meet.

Industrial relevance. The workshop topic was highly relevant for companies dealing with big data. Three seminar participants are affiliated with Google, one has worked on hashing algorithms at AT&T for a decade, and one is affiliated with VMware. Also, one organizer was affiliated with IBM at the time of the seminar.

Outcome of the seminar

The seminar brought together quite a few leading figures in the area of hashing, mixed with a good fraction of young researchers, some of which are behind many of the most exciting results in the area in recent years. Areas that were particularly well represented were: Analysis of multiple-choice hashing methods, hashing for high-dimensional search problems, applications of hashing in string algorithms, applications of hashing in machine learning, streaming (approximation) algorithms, high-performance hash function construction, and algorithm engineering. Many results in these areas were presented in 18 shorter talks. Four longer talks (by A. Andoni, A. McGregor, U. Wieder, and Q. Zhang) contributed background, overview over new results, and aspects of applications of hashing in industry. Open problems were discussed in an open problem session; four of them are included in this report.

The paper [1] on fillable arrays co-authored by J. Nelson was motivated by a talk given by T. Hagerup (see Section 3.7 of the full report) and can thus be seen as a first direct result of the seminar.

We, the organizers, would like to thank all participants for their contributions in talks, material, and discussions, and in particular the speakers of the longer talks. Many thanks are due to the Dagstuhl staff both in the offices and in the castle for their support in making this seminar a success.

References

1 Jacob Teo Por Loong, Jelani Nelson, Huacheng Yu: Fillable arrays with constant time operations and a single bit of redundancy. CoRR abs/1709.09574 (2017)

6.29 Theory of Randomized Optimization Heuristics

Organizers: Carola Doerr, Christian Igel, Lothar Thiele, and Xin Yao

Seminar No. 17191

Date: May 7–12, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.22

Participants: Youhei Akimoto, Dirk V. Arnold, Anne Auger, Thomas Bäck, Hans-Georg Beyer, Dimo Brockhoff, Maxim Buzdalov, Arina Buzdalova, Francisco Chicano, Benjamin Doerr, Carola Doerr, Anton V. Eremeev, Carlos M. Fonseca, Tobias Friedrich, Christian Gießen, Tobias Glasmachers, Nikolaus Hansen, Matthew W. Hoffman, Christian Igel, Thomas Jansen, Martin S. Krejca, William B. Langdon, Per Kristian Lehre, Johannes Lengler, Frank Neumann, Pietro S. Oliveto, Adam Prugel-Bennett, Jonathan E. Rowe, Günter Rudolph, Jonathan L. Shapiro, Ofer M. Shir, Dirk Sudholt, Andrew M. Sutton, Olivier Teytaud, Lothar Thiele, Heike Trautmann, Carsten Witt, Jing Yang, Xin Yao, Christine Zarges

Randomized search and optimization heuristics such as evolutionary algorithms, ant colony optimization, particle swarm optimization, and simulated annealing, have become established problem solvers. They have successfully been applied to a wide range of real-world applications, and they are applicable to problems that are non-continuous, multi-modal, and/or noisy as well as to multi-objective and dynamic optimization tasks. Theory of randomized optimization heuristics aims at providing mathematically founded insights into the working principles of these general-purpose problem solvers, and at developing new and more powerful heuristic optimization methods in a principled way. The seminar has covered several important streams in this research discipline. Among several other topics, extended discussions have been held on the advantages of population-based heuristics and of non-static parameter choices, optimization problems with constrains, as well as existing and possible connections to research in machine learning.

The seminar continues to be one of the key stimulator for novel ideas, tools, and approaches in the theory of randomized optimization heuristics. Accordingly, the acceptance rate for the invitations has been staying at a very high level.

Topics

The research in theory of randomized optimization heuristics is as broad as the applicability of these methods. The seminar succeeded in covering the various theoretical approaches. There was a focus on important cross-cutting topics, which we briefly outline in the following.

One of the most prominent research areas in the theory of randomized optimization heuristics deals with *runtime* and *convergence analysis*, aiming at proving bounds on the speed of the convergence to an optimal solution. Typical questions concern the advantages of certain algorithmic choices, such as

- the size of the memory (*population*),
- the usage of different sampling strategies (*variation* of previously sampled search points, in particular via *mutation* of one previously evaluated solution candidate and *recombination* of two or more previous search points), and
- the selection strategies (e.g., *elitist selection* which never discards a best-so-far solution vs. the non-elitist Boltzmann strategies found in Simulated Annealing, SSWM, and the Metropolis algorithm).

One of the most relevant objectives in empirical and theoretical works on randomized optimization heuristics is to determine the best parameter settings for the above-described algorithmic components. Given the complex interactions between the parameter values, this *parameter tuning* task is a highly challenging one. It is further complicated by the observation that for most problems the optimal parameter settings change during the optimization process, thus asking for *parameter control* mechanisms that adjust the parameter value to the current state of the optimization. Identifying such reasonable (and possibly provably optimal) ways to update the parameter choices has been one of the intensively discussed topics of the seminar. Significant progress towards a better understanding of different parameter update schemes has been obtained in the last few years, as has been demonstrated by several talks, for example on self-adaptive and self-adjusting parameter updates as well as on estimation of distribution algorithms. Among other results, several connections to related questions in machine learning have been made, motivating the organizers to include machine learning as a focus topic of this seminar.

Randomized search heuristics are currently very popular in general machine learning[46] in the form of *Bayesian optimization*.

[46] One may well argue that randomized search heuristics actually belong to the broader field of machine learning methods.

However, there has been little connection between the research in Bayesian optimization and the established work on randomized search heuristics, and the seminar was a step to change this. The first talk of the seminar was an extended introduction to Bayesian optimization by Matthew W. Hoffman from Google DeepMind, a leading expert in the field. The talk set the stage for informed discussions on similarities and differences between methods—and potential synergies between the research fields. Thompson sampling, an important algorithm in Bayesian optimization, was revisited in the talk by Jonathan Shapiro on dueling bandit problems, which demonstrated randomized search heuristics in a scenario of high commercial relevance. A common application of randomized search heuristics in general machine learning is model selection, for example finding a tailored structure of a neural network. This was addressed in the talk by Olivier Teytaud from Google Brain, who discussed model selection heuristics for large-scale machine learning systems. Randomized search heuristics are also successfully used for reinforcement learning. Arina Buzdalova presented work in which the connection is the other way round: ideas from reinforcement learning are used to improve randomized optimization (by controlling the choice of objectives).

Another intensively discussed topic, highly relevant in both discrete and continuous optimization, was constrained optimization. Here the main research questions concern the different ways to model constrained problems in black-box settings, and suitable algorithmic approaches. In addition to a number of theoretical results on constrained optimization, the need for a well-designed benchmark suite has also been discussed. As a result of one of the breakout sessions of the previous Dagstuhl Seminar 15211 on *Theory of Evolutionary Computation*, Dimo Brockhoff presented the recent extension of the COCO benchmark set (http://coco.gforge.inria.fr/doku.php) to constrained optimization. Dirk Arnold presented some work indicating that this extension of COCO is very timely, and much needed in the randomized search heuristics community. Furthermore, another breakout session has been held this year on the topic of constrained optimization, organized by Frank Neumann, with a focus on the different ways to model soft and hard constraints in discrete black-box optimization.

Organization

The seminar schedule has offered a good flexibility for the participants to propose talks and discussions of different lengths. 29 talks of 10–30 minutes each have been held in total, in the plenary room. These plenary talks were complemented by a introductory tutorial on Bayesian Optimization by Matt Hoffman on Monday morning and by 7 breakout sessions on various topics, including methodology-oriented discussions on the applicability of drift analysis in continuous domains or how to interpret the CMA-ES in the framework of information geometry optimization as well as problem-driven brainstorming on constrained optimization, the role of diversity in heuristic search, preference-based selection, and the method of estimation of distribution algorithms. Another breakout session was devoted to discussing the importance and possible obstacles of bringing theory-and practice-driven research in heuristic optimization closer together. The breakout sessions have been held on Tuesday, Wednesday, and Thursday afternoon, respectively, and have all witnessed high attendance rates. All talks and breakout sessions are summarized in Sections 3 and 4 of the full report.

We would like to express our gratitude to the Dagstuhl staff and all participants for making this Dagstuhl Seminar 17191 on *Theory of Randomized Optimization Heuristics* such a successful event, which has been a pleasure to organize.
Carola Doerr (CNRS and Pierre et Marie Curie University Paris 6, FR)
Christian Igel (University of Copenhagen, DK)
Lothar Thiele (ETH Zürich, CH)
Xin Yao (University of Birmingham, GB and SUSTech Shenzhen, CH)

6.30 Human-Like Neural-Symbolic Computing

Organizers: Tarek R. Besold, Artur d'Avila Garcez, and Luis C. Lamb
Seminar No. 17192

Date: May 7–12, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.56

Participants: Lucas Bechberger, Tarek R. Besold, Jelmer Borst, Artur d'Avila Garcez, James Christopher Davidson, Marc de Kamps, Derek Doran, Ulrich Furbach, Raquel Garrido Alhama, Marco Gori, Pascal Hitzler, Dieuwke Hupkes, Caroline Jay, Kristian Kersting, Kai-Uwe Kühnberger, Oliver Kutz, Luis C. Lamb, Martha Lewis, Isaac Noble, Sarah Schulz, Katja Seeliger, Luciano Serafini, Daniel L. Silver, Michael Spranger, Keith Stenning, Niels A. Taatgen, Serge Thill, Frank Van der Velde, Tillman Weyde

The underlying idea of Human-Like Computing is to incorporate into Computer Science aspects of how humans learn, reason and compute. Recognising the relevance of the scientific trends in big data, data science methods and techniques have achieved industrial relevance in a number of areas, from retail to health, by obtaining insight from large data collections. Notably, neural networks have been successful and efficient at large-scale language modelling, speech recognition, image, video and sensor data analysis [3, 12, 15]. Human beings, on the other hand, are excellent at learning from very few data examples, capable of articulating explanations and resolving inconsistencies through reasoning and communication [7, 9, 16, 17].

Despite the recent impact of deep learning, limited progress has been made towards understanding the principles and mechanisms underlying language and vision understanding. Under this motivation, the seminar brought together not only computer scientists, but also specialists in artificial intelligence (AI), cognitive science, machine learning, knowledge representation and reasoning, computer vision, neural computation and natural language processing. In particular, the methodology of neural-symbolic computation [4, 7, 12], which can offer a principled interface between the relevant fields, especially symbolic AI and neural computation, was adopted in an attempt to offer a new perspective of reconciling large-scale modelling with human-level understanding, thus building a roadmap for principles and applications of Human-Like Neural-Symbolic Computing.

The techniques and methods of neural-symbolic computation have already been applied effectively to a number of areas, leading to developments in deep learning, data science and human-like computing [3]. For instance, neural-symbolic integration methods have been applied to temporal knowledge evolution in dynamic scenarios [6, 10, 14], action learning in video understanding [10], uncertainty learning and reasoning [1], argument learning in multiagent scenarios [7, 8], hardware and software verification and learning [2], ontology learning [13] and distributed temporal deep learning in general, with several applications in computer science [2, 10, 14].

Specifically, in this Dagstuhl Seminar we aimed at: (i) building better bridges between symbolic and sub-symbolic reasoning and learning, and between big data and human-like learning; (ii) comparative analyses and evaluations of the explanatory capacity of language modelling tools and techniques; (iii) design and applications of knowledge extraction methods and techniques towards life-long learning and transfer learning.

The seminar consisted of contributed and invited talks, breakout and joint group discussion sessions, and scientific hackathons. After each presentation or discussion session, open problems were identified and questions were raised. The area is clearly growing in importance, given recent advances in Artificial Intelligence and Machine Learning. In particular, the need for explainability in AI clearly poses relevant questions to learning methodologies, including deep learning. In summary, the main research directions identified by participants are:

- Explainable AI: The recent success of deep learning in vision and language processing, associated with the growing complexity of big data applications has led to the need for *explainable AI models*. In neural-symbolic computing, rule extraction, interpretability, comprehensibility leading to the development of integrated systems, are one of the principled alternatives to lead these efforts [5, 7], as discussed in the Explainability hackathon. Furthermore, the concept of modularity in multimodal learning in deep networks is crucial to the development of the field and can help achieve knowledge extraction (as identified in [5, 6]) which can result in the development of effective knowledge extraction methods towards explainable AI, as discussed in the deep learning with symbols hackathon.

- Hybrid Cognitive Architectures: The development of Cognitive Architectures capable of simulating and explaining aspects of human cognition also remains an important research endeavour. Some cognitive architectures typically consider symbolic representations, whereas others employ neural simulations. The integration of these models remains a challenge and there are benefits on integrating the accomplishments of both paradigms, as identified in the cognitive architectures hackathon.
- Statistical Relational Learning: Logic Tensor Networks (LTNs) [11] provides a model that integrates symbolic knowledge (encoded as first-order logic relations) and subsymbolic knowledge (represented as feature vectors). The LTNs enable the representation of relational knowledge infusion into deep networks, and knowledge completion and distilling through querying the networks. There remains a number of challenges in integrating, explaining and computing symbolic knowledge in deep networks. Both LTNs [11] and Connectionist Modal and Temporal Logics [6, 7, 14] offer effective alternatives towards these research challenges, as explored in the LTN hackathon.

The seminar builds upon previous seminars and workshops on the integration of computational learning and symbolic reasoning, such as the Neural-Symbolic Learning and Reasoning (NeSy) workshop series, and the previous Dagstuhl Seminar 14381: Neural-Symbolic Learning and Reasoning [5].

References

1 Tarek R. Besold, Artur S. d'Avila Garcez, Keith Stenning, Leendert W. N. van der Torre, Michiel van Lambalgen: Reasoning in Non-probabilistic Uncertainty: Logic Programming and Neural-Symbolic Computing as Examples. *Minds and Machines*, 27(1): 37-77, 2017.

2 Rafael V. Borges, Artur S. d'Avila Garcez, Luis C. Lamb. *Learning and Representing Temporal Knowledge in Recurrent Networks*. IEEE Trans. Neural Networks and Learning Systems 22(12):2409-2421, Dec. 2011. IEEE Press.

3 A. d'Avila Garcez, Tarek R. Besold, Luc De Raedt, Peter Földiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kühnberger, Luis C. Lamb, Risto Miikkulainen, Daniel L. Silver. Neural-Symbolic Learning and Reasoning: Contributions and Challenges. Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches. Stanford Univ., March 2015, pp. 18-21, AAAI press, 2015.

4 Artur S. d'Avila Garcez, Dov M. Gabbay, Krysia Broda. *Neural-Symbolic Learning System: Foundations and Applications*. Springer-Verlag, New York, Inc., USA, 2002.

5 A. S. d'Avila Garcez, Marco Gori, Pascal Hitzler, Luís C. Lamb: Neural-Symbolic Learning and Reasoning (Dagstuhl Seminar 14381). Dagstuhl Reports 4(9): 50-84, 2015.

6 Artur S. d'Avila Garcez and Luis C. Lamb. A Connectionist Computational Model for Epistemic and Temporal Reasoning. *Neural Computation*, 18(7):1711-1738, 2006.

7 Artur S. d'Avila Garcez, Luis C. Lamb and Dov M. Gabbay. *Neural-Symbolic Cognitive Reasoning*. Springer Publishing Company, 2009. ISBN: 9783540732457.

8 A. d'Avila Garcez, D.M. Gabbay and Luis C. Lamb: A neural cognitive model of argumentation with application to legal inference and decision making. Journal of Applied Logic, 12(2):109-127, 2014.

9 M. de Kamps and F. van de Velde. From neural dynamics to true combinatorial structures. Behavioral and Brain Sciences, 20, pp. 88-108, 2006.

10 H.L. de Penning, Artur S. d'Avila Garcez, Luis C. Lamb, John-Jules Ch. Meyer: *A Neural-Symbolic Cognitive Agent for Online Learning and Reasoning*. Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence, IJCAI 2011. Pages 1653-1658, 2011.

11 Ivan Donadello, Luciano Serafini, and Artur d'Avila Garcez. *Logic Tensor Networks for Semantic Image Interpretation*. Proceedings of the 26th International Joint Conference on Artificial Intelligence, 2017.

12 Barbara Hammer, Pascal Hitzler (Eds.): Perspectives of Neural-Symbolic Integration. Studies in Computational Intelligence 77, Springer 2007.

13 P. Hitzler, S. Bader, and A. d'Avila Garcez. *Ontology learning as a use case for neural-symbolic integration*. In Proc. Workshop on Neural-Symbolic Learning and Reasoning, NeSy'05 at IJCAI-05.

14 Luis C. Lamb, R.V. Borges and A.S. d'Avila Garcez. *A connectionist cognitive model for temporal synchronisation and learning*. Proc. AAAI Conference on Artificial Intelligence AAAI-07, pages 827-832, AAAI Press, 2007.

15 Y. LeCun, Y. Bengio, G. Hinton. Deep Learning. *Nature*, 521 (7553): 436-444, 2015.

16 G.F. Marcus, S. Vijayan, S.B. Rao, and P.M. Vishton. Rule learning by seven-month-old infants. *Science*, 283(5398):77–80, 1999.

17 K. Stenning and M. van Lambalgen. *Human reasoning and Cognitive Science*. MIT Press, Cambridge, MA, 2008.

18 Alan M. Turing. Computing Machinery and Intelligence, *Mind*, LIX (236): 433–460, 1950.

6.31 Formal Synthesis of Cyber-Physical Systems

Organizers: Calin A. Belta, Rupak Majumdar, Majid Zamani, and Matthias Rungger

Seminar No. 17201

Date: May 14–19, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.84

Participants: Erika Abraham, Murat Arcak, Calin A. Belta, Ivana Cerná, Samarjit Chakraborty, Chih-Hong Cheng, Jyotirmoy Deshmukh, Rüdiger Ehlers, Martin Fabian, Georgios Fainekos, Antoine Girard, Anak Agung Julius, Roozbeh Kianfar, Hadas Kress-Gazit, Jan Kretinsky, Stéphane Lafortune, Jun Liu, Rupak Majumdar, Manuel Mazo, Sanjoy K. Mitter, Sahar Mohajerani, Necmiye Ozay, Ivan Papusha, Nir Piterman, Pavithra Prabhakar, Gunther Reißig, Harald Ruess, Matthias Rungger, Indranil Saha, Anne-Kathrin Schmuck, Sadegh Soudjani, Paulo Tabuada, Wolfgang Thomas, Hazem Torfah, Stavros Tripakis, Ashutosh Trivedi, Jana Tumova, Marcell Vazquez-Chanlatte, Majid Zamani

Cyber-Physical Systems (CPS) are complex systems resulting from intricate interaction of discrete computational devices with the continuous physical plants. Within CPS, embedded control software plays a significant role by monitoring and adjusting several physical variables, e.g. temperature, velocity, pressure, density, and so on, through feedback loops where physical processes interact with computational devices. Recent advances in computation, storage, and networking have made tremendous advances in hardware and system platforms for CPS. With this growing trend in computational devices, embedded control software is becoming more and more ubiquitous in many safety-critical applications including automotive, aerospace, transportation systems, critical infrastructure, energy, robotics, healthcare, and many other domains. Unfortunately, the design of embedded control software nowadays is still based on ad-hoc solutions resulting in brittle and error-prone software, and very high verification and validation costs. In order to detect and eliminate design flaws and inevitable software bugs, a large portion of the design budget is consumed with validation and verification efforts, which are often lengthy. On the other hand, by changing the emphasis from verification to synthesis, it is possible to synthesize correct-by-design embedded control software for CPS while providing formal guarantees of correctness and preventing the need for costly post facto verification.

In recent years, there has been a lot of progress in designing automatic and correct-by-construction techniques for controller synthesis for interacting discrete and continuous systems. These new techniques have combined techniques from continuous control theory as well as from computer science. The focus of this seminar was to provide a state-of-the-art of this nascent but important field, and to describe challenges and opportunities for synthesis techniques to transition to the real world. By the nature of the topic, the participation at the seminar was inter-disciplinary, and consisted of computer scientists and control theorists, both from academia and from industry. Instead of a sequence of presentations of individual research results, the seminar was organized as a sequence of open discussions on topics of common interest to the participants, such as techniques for scalable controller synthesis, identification of application domains and recent success stories, compositionality and system design, end-to-end arguments about systems, as well as education and outreach.

This seminar benefitted the control as well as computer science communities by bridging the gap between many complementary concepts studied in each community. A more detailed survey on the topics of the seminar is in preparation.

Outcomes of the seminar

The seminar focused on the challenges in the application of formal synthesis techniques for automatic, correct-by-construction synthesis of CPS. The seminar had a total of 45 participants with a mix of computer scientists and control theorists.

Sessions

The seminar was organized as a sequence of open discussions led by one or two moderators. Each session had a scribe to note down the discussion. The scribe notes were shared with all participants, who added their comments or filled in more information. The updated notes were used to prepare the session summaries (in the next Section).

The following sessions were organized:

1. Application domains, success stories, and obstacles to adoption
2. Fundamental algorithmic and scalability challenges in formal synthesis
3. Tools and infrastructure
4. Education and outreach

5. Data-driven and search-driven approaches to synthesis
6. Compositionality in synthesis
7. Optimality and completeness
8. Synthesis of distributed protocols from scenarios and requirements.
9. Specification languages
10. Robustness and resiliency
11. Explainability and user interaction
12. End-to-end correctness
13. Formal synthesis challenges in robotics
14. Cyber-Security of CPS

6.32 Challenges and Opportunities of User-Level File Systems for HPC

Organizers: André Brinkmann, Kathryn Mohror, and Weikuan Yu

Seminar No. 17202

Date: May 14–19, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.97

Participants: Stergios V. Anastasiadis, Ned Bass, John Bent, Thomas Bönisch, Luc Bougë, André Brinkmann, Suren Byna, Philip Carns, Lance Evans, Wolfgang Frings, Hermann Härtig, Dean Hildebrand, Nathan Hjelm, Scott Klasky, Michael Kluge, Michael Kuhn, Jay Lofstead, Satoshi Matsuoka, Alberto Miranda, Kathryn Mohror, David Montoya, Federico Padua, Mark Parsons, Maria S. Perez, Franz Josef Pfreundt, Raghunath Raja Chandrasekar, Robert B. Ross, Brad Settlemyer, Osamu Tatebe, Tianqi Xu, Weikuan Yu, Yue Zhu

Seminar Overview

The primary goal of this Dagstuhl Seminar was to bring together experts in I/O performance, file systems, and storage, and collectively explore the space of current and future problems and solutions for I/O on hierarchical storage systems. We had a lively week of learning about each other's approaches as well as unique I/O use cases that can influence the design of a community-driven file and storage system standards. We also engaged in several informal, in-depth discussions on questions surrounding how we should best move forward in the I/O and storage community.

A portion of agenda for this meeting was partitioned into sessions containing short talks. The short talk sessions were grouped into high level topic areas: high performance computing and storage systems today; user needs for I/O; user level file system implementations; object stores and alternatives; and file systems building blocks. The intention behind the short talks was to acquaint the attendees with each other's work and to inspire further discussions. Following each talk topic, we had panel-style discussions with the talk speakers serving as the panel. In these panel-style discussions, the audience had the opportunity to ask questions about the speakers' talks as well as note and discuss commonalities and differences across the presentations.

The remainder of the agenda for the meeting was reserved for open discussions with the entire group. The participants engaged in lengthy discussions on various questions that arose from the talks. Additionally, participants were encouraged to propose and vote for discussion topics on a white board. The proposed topics with the most votes were included in the agenda. The in-depth discussion topics included:

- How are stage-in and stage-out operations actually going to work?
- How can we fairly judge the performance of storage systems – IO 500?
- What is a user-level file system? What do we mean when we say that?
- How can we characterize what users need from storage systems?
- Are we ready to program to a memory hierarchy versus block devices?
- and What should we do about POSIX?

The combination of short talks and open discussions resulted in a fruitful meeting. Since the work of the participants was not necessarily familiar to all, the short talks provided a foundation for getting everyone oriented with each other's efforts. Once that was achieved, we were able to productively dive into the informal topic discussions. Overall, several common themes emerged from the talks and discussions. The participants agreed that these themes were important to address to meet the needs of HPC applications on next-generation storage systems. We include these themes in this report in Section 9 of the full report to serve as suggestions for further investigations.

Report Organization

Here we present an overview of the topics in this report to guide the reader. Our goal in this report is to capture as much information as possible from the seminar so that those who could not attend can benefit from the talks and discussions.

We detail the short talk sessions in Sections 3–7 of the full report. First, we provide a summary of the notes from the session note taker and other comments from the talks and panel discussions. Following this summary we provide a listing of each talk in the session and its abstract.

In Sections 8.1–8.6 of the full report we give summaries of the informal discussion sessions. The summaries in this case are in

outline format in order to capture the conversational and informal nature of the sessions. In many cases, the discussions drew out many interesting questions instead of clear paths forward, so the outline format captured this well.

Following the summaries of the sessions, in Section 9 of the full report we conclude with a discussion of recurring themes, including issues for future discussion and work, that occurred during the meeting. We feel that these themes are the true product of this meeting and can serve as a foundation for future meetings or other community efforts.

6.33 Geometric Modelling, Interoperability and New Challenges

Organizers: Falai Chen, Tor Dokken, Thomas A. Grandine, and Géraldine Morin

Seminar No. 17221

Date: May 28 to June 2, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.140

Participants: Gudrun Albrecht, Marco Attene, Chandrajit Bajaj, Falai Chen, Heidi Dahl, Tor Dokken, Gershon Elber, Daniela Fußeder, Carlotta Giannelli, Christina Gillmann, Ron Goldman, Thomas A. Grandine, Jens Gravesen, Hans Hagen, Stefanie Hahmann, Dianne Hansford, Bert Jüttler, Panagiotis Kaklis, Hongmei Kang, Benjamin Karer, Ligang Liu, Tom Lyche, Carla Manni, Nicolas Mellado, Gëraldine Morin, Georg Muntingh, Jörg Peters, Konrad Polthier, Xiaoping Qian, Ulrich Reif, Malcolm A. Sabin, Maria Lucia Sampoli, Giancarlo Sangalli, Camille Schreck, Hendrik Speleers, Timothy L. Strotman, Georg Umlauf, Tamas Várady, Denis Zorin

This report documents the program and the outcomes of Dagstuhl Seminar 17221 "Geometric Modelling, Interoperability and New Challenges".

The importance of accurate geometric models of shapes, both naturally occurring and man-made, is rapidly growing with the maturing of novel manufacturing technologies and novel analysis technologies. The advent of big data challenges and cloud-based computing serves to confound the distribution and remote access of these geometric models. While previous Dagstuhl seminars on geometric modeling were focused on basic research, this seminar was focused on applications of geometric modeling. We selected four core application areas that stretch the underlying mathematical underpinnings of the discipline to its limits and beyond:

- Big data and Cloud computing
- Multi-material additive manufacturing (3D Printing)
- Isogeometric analysis
- Design optimization

The seminar provided a forum for leading researchers to present new ideas, to exchange scientific insights, and to bring together practical applications and basic research. The previous seminar explored the theory of shape representations, shape transformations, and computational models for each. This seminar explored application of this theory to the four above-mentioned application challenges, strengthening the reliability and performance of applications in engineering, manufacturing, and scientific exploration. The goal of the seminar was to establish a common understanding between the Geometric Modeling research community and the above application fields by addressing the following questions:

- How to handle geometric descriptions that cannot be processed by one computer but have to be distributed among many computers in the cloud?
- How to process huge 3D data sets that may be noisy and incomplete into clean, scalable, and easy to access 3D environments?
- How to turn big, dispersed maybe noisy and incomplete geometric data into clean, scalable, and easy to access 3D information that can be used for change detection and decision making?
- How to represent, control, and process complex, anisotropic, internal material enabled by additive manufacturing, and how to design for additive manufacturing?
- How to perform topology optimization of the internal structure of objects to enable additive manufacturing to reach its full potential?
- How to introduce analysis-based design in CAD systems with isogeometric analysis in mind?

To answer these questions we brought together participants from industry urgently in need of better solutions, researchers in the application areas represented by the four topics, and researchers in the geometric modeling community whose interests align with the four topic areas. The scientific presentations lasted 15 to 20 minutes. Senior researchers gave 4 overview talks on the 4 themes of the seminar. Perspective working groups were organized for each of the four topics.

Fig. 6.6

Dagstuhl Keywords. Picture on flickr by 17392 Dagstuhl Seminar participant m. c. schraefel. https://flic.kr/p/BRLkF9. Photo courtesy of m. c. schraefel.

6.34 Robust Performance in Database Query Processing

Organizers: Renata Borovica-Gajic, Goetz Graefe, and Allison Lee

Seminar No. 17222

Date: May 28 to June 2, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.5.169

Participants: Angelos-Christos Anadiotis, Tahir Azim, Peter A. Boncz, Renata Borovica-Gajic, Surajit Chaudhuri, John Cieslewicz, Thanh Do, Campbell Fraser, Johann-Christoph Freytag, Goetz Graefe, Fabian Hüske, Alfons Kemper, Allison Lee, Thomas Neumann, Anisoara Nica, Glenn Paulley, Ilia Petrov, Bart Samwel, Kai-Uwe Sattler, Knut Stolze, Immanuel Trummer, Srinivas Karthik Venkatesh, Jiaqi Yan, Marcin Zukowski

The Dagstuhl Seminar 17222 on "Robust performance in database query processing" assembled researchers from industry and academia for the third time to discuss robustness issues in database query performance. The seminar gathered 24 researchers around the world working on plan generation and plan execution in database query processing and in cloud-based massively parallel systems with the purpose to address the open research challenges with respect to the robustness of database management systems.

Delivering robust query performance is well known to be a difficult problem for database management systems. All experienced DBAs and database users are familiar with sudden disruptions in data centers due to poor performance of queries that have performed perfectly well in the past. The goal of the seminar is to discuss the current state-of-the-art, to identify specific research opportunities in order to improve the state-of-affairs in query processing, and to develop new approaches or even solutions for these opportunities.

Unlike the previous seminars, the organizers (Renata Borovica-Gajic, Goetz Graefe and Allison Lee) this time attempted to have a focused subset of topics that the participants discussed and analyzed in more depth. From the proposed topics on algorithm choices, join sequences, updates, database utilities, parallelism and skew, column stores, physical database design, and explainability of non-robust query performance, the participants chose four topics and formed four work groups: i) one discussing updates and database utilities, ii) one discussing parallelism and skew, iii) one discussing join sequences, and iv) one focusing on the explanations of the sources of non-robust performance.

Upon choosing the topics of interest, the organizers then guided the participants to approach the topic through a set of steps: by first considering related work in the area; then introducing metrics and tests that will be used for testing the validity and robustness of the solution; after metrics, the focus was on proposing specific mechanisms for the proposed approaches; and finally the last step focused on the implementation policies.

The seminar thus spent its first day on reviewing prior related work, with a special emphasis on the pieces of work that appeared following the previous instances of the seminar: benchmarks (Dagstuhl 12321 [4, 6, 7]), Smooth Scan [2], and Generalized join [3]. Tuesday was spent on defining metrics and tests. On Wednesday, the participants discussed possible alternative approaches and hiked together in the woods. Thursday was focused on driving one chosen approach to specific mechanisms. Finally, we spent Friday on discussing the policies and presented the overall progress.

At the end of the week, each group was hoping to continue their work towards a research publication. The group on parallelism and skew was hoping to publish first a survey on forms of skew and existing remedies for skew. The work group on dynamic join sequences even had a working prototype by the end of the seminar.

References

1 Mert Akdere, Ugur Çetintemel, Matteo Riondato, Eli Upfal, and Stanley B. Zdonik. Learning-based query performance modeling and prediction. pages 390–401, 2012.

2 Renata Borovica-Gajic, Stratos Idreos, Anastasia Ailamaki, Marcin Zukowski, and Campbell Fraser. Smooth Scan: Statistics-oblivious access paths. In *ICDE*, 2015.

3 Goetz Graefe. New algorithms for join and grouping operations. *Computer Science - R&D*, 27(1):3–27, 2012.

4 Goetz Graefe, Wey Guy, Harumi A. Kuno, and Glenn N. Paulley. Robust query processing (dagstuhl seminar 12321). *Dagstuhl Reports*, 2(8):1–15, 2012.

5 Hakan Hacigumus, Yun Chi, Wentao Wu, Shenghuo Zhu, Junichi Tatemura, and Jeffrey F. Naughton. Predicting query execution time: Are optimizer cost models really unusable? pages 1081–1092, 2013.

6 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, DBTest '11, pages 7:1–7:6, New York, NY, USA, 2011. ACM.

7 Martin L. Kersten, Alfons Kemper, Volker Markl, Anisoara Nica, Meikel Poess, and Kai-Uwe Sattler. Tractor pulling on data warehouses. In *Proceedings of the Fourth International Workshop on Testing Database Systems*, DBTest '11, pages 7:1–7:6, New York, NY, USA, 2011. ACM.

8 Tanu Malik, Randal C. Burns, and Nitesh V. Chawla. A black-box approach to query cardinality estimation. pages 56–67, 2007.

9 Microsoft. Microsoft adaptive query processing and diagnostics. https://www.youtube.com/watch?v=szTmo6rTUjM.

10 Wolf Rödiger, Sam Idicula, Alfons Kemper, and Thomas Neumann. Flow-join: Adaptive skew handling for distributed joins over high-speed networks. In *ICDE*, pages 1194–1205, 2016.

11 TPC. Tpc-h benchmark. http://www.tpc.org/tpch/.

12 Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. Lightweight graphical models for selectivity estimation without independence assumptions. *PVLDB*, page 2011, 2011.

13 Kostas Tzoumas, Amol Deshpande, and Christian S. Jensen. Efficiently adapting graphical models for selectivity estimation. *The VLDB Journal*, 22(1):3–27, 2013.

14 Wentao Wu, Yun Chi, Hakan Hacígümüş, and Jeffrey F. Naughton. Towards predicting query execution time for concurrent and dynamic database workloads. *PVLDB*, 6(10):925–936, 2013.

15 Wentao Wu, Xi Wu, Hakan Hacigümüş, and Jeffrey F. Naughton. Uncertainty aware query execution time prediction. *PVLDB*, 7(14):1857–1868, 2014.

6.35 Epistemic Planning

Organizers: Chitta Baral, Thomas Bolander, Sheila McIlraith, and Hans van Ditmarsch

Seminar No. 17231

Date: June 5–9, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.1

Participants: Guillaume Aucher, Chitta Baral, Vaishak Belle, Thomas Bolander, Tran Cao Son, Tristan Charrier, Jens Claßen, Thorsten Engesser, Esra Erdem, Tim French, Malvin Gattinger, Nina Gierasimczuk, Malte Helmert, Andreas Herzig, Mathias Justesen, Ioannis Kokkinis, Filippos Kominis, Barteld Kooi, Louwe B. Kuijer, Jérôme Lang, Yves Lesperance, Kai Li, Yongmei Liu, Robert Mattmüller, Bastien Maubert, Sheila McIlraith, Timothy Miller, Bernhard Nebel, Andrés Occhipinti Liberman, Ron Petrick, Sophie Pinchinat, Ramaswamy Ramanujam, Torsten Schaub, Richard Scherl, Francois Schwarzentruber, Maayan Shvo, Sunil Easaw Simon, Michael Thielscher, Hans Van Ditmarsch, Jan van Eijck, Ivan José Varzinczak, Yanjing Wang, Bruno Zanuttini

This seminar brought together three largely independent research communities: Dynamic Epistemic Logic (DEL), Knowledge Representation and Reasoning (KR&R, subsequently KR) and Automated Planning. All three communities have a tradition of investigating the interaction between dynamical systems and epistemic states, but with a different focus and by different means. In the context of this seminar, despite occasional overlap, DEL has mainly investigated the formal semantics of communication and communicative actions, KR&R has mainly focussed on the theories of action and change, and AP has mainly focussed on computational techniques to automatically generate plans. This seminar aimed to encourage and nurture increasing synergies between these three strong and largely independent research communities leading to frameworks for Epistemic Planning: planning with epistemic states, actions, and goals. The seminar succeeded in strengthening and broadening the cross-disciplinary research community that emerged after a prior Dagstuhl Seminar on the subject, entitled "Planning with Epistemic Goals", that was held in January 2014 (seminar number 14032). This follow-up seminar led to a better understanding and articulation of commonalities, synergies, and deficits between the DEL, KR and Planning communities.

The main components of the seminar were tutorial talks and group work. There were four tutorials, by Andreas Herzig, Bernhard Nebel, Tran Cao Son and Ramaswamy Ramanujam. The four tutorial talks presented epistemic planning from the perspective of four different research communities. Andreas Herzig presented epistemic planning from a knowledge representation perspective, Bernhard Nebel from a classical planning perspective, Tran Cao Son from a theories of action and change perspective, and Ramaswamy Ramanujam from a distributed systems and temporal logic perspective. Abstracts for the four tutorials are included in Section 3 of the full report.

The group work was performed in five separate working groups, each addressing a distinct important objective in epistemic planning of shared interest to the DEL, KR and Planning communities. Specifically, the five groups worked on (1) *Developing Benchmarks for Epistemic Planning*, (2) *Exploring Action Types and their Representations*, (3) *Exploring the Relations Between Knowledge and Belief in Multi-Agent Epistemic Planning*, (4) *Practical Tools, Resources and Computational Techniques*, and (5) *Correspondence Between Planning Problems and Games*. These group themes are briefly described below, and the outcome of the work in each group is documented in Section 4 of the full report.

1. **Developing Benchmarks for Epistemic Planning.** In the planning community benchmarks are common, but in the epistemic planning community they are not. The overall goal of the group working on this theme was to formulate a list of ten benchmark epistemic planning problems, in view of setting goals that can evolve into competitions. In particular, the focus was on targeting planning problems that are truly epistemic, meaning problems in which the epistemic dimension – knowledge and ignorance – cannot easily be disregarded.

   Guiding questions for the work in this group were: What problems help define and circumscribe what we are studying? What are specific tasks that motivate this area of study, e.g. epistemic planning, protocol synthesis, automated diagnosis, verification, and communication? What problems can help drive future research in developing formalisms and implementing systems of epistemic planning? What are the features of the relevant planning problems in terms of knowledge vs. belief, single- vs. multi-agent, communicative vs. sensing vs. ontic actions, deterministic vs. non-deterministic actions, etc. How do we evaluate "hardness" of problems, e.g. in terms of level of nesting of belief/knowledge, types of actions, size of

problem, scalability, and quality of solution. Do benchmarks for some of these problems already exist?

2. **Exploring Action Types and their Representations.** It is important to identify and broaden the list of action types relevant to epistemic planning. In terms of communicative actions we can for instance at least distinguish between announcements, questions, requests and instructions. How these actions are best represented is also an important issue. It should be explored whether formalizations in dynamic epistemic logic are appropriate for planning or whether a more simplified way of representing multi-agent actions is needed. The overall goal of the group working on this theme was to identify, classify and possibly broaden the list of action types relevant to epistemic planning, as well as to explore formalisms for representing these action types.

   Guiding questions for the work in this group were: What are the relevant distinctions between types of actions, e.g. epistemic vs. ontic, deterministic vs. non-deterministic vs. probabilistic, instantaneous vs. durative, sensing vs. announcements, degree of observability (public, private, semi-private), etc. What formalisms can support expressing and distinguishing between these action types (e.g. DEL, Situation Calculus, Knowledge-based Programs)? How are multi-agent actions represented, in particular how are conflicts between concurrently occurring actions specified, and how is observability of concurrent actions specified in terms of the observability of the constituting actions?

3. **Exploring the Relations Between Knowledge and Belief in Multi-Agent Epistemic Planning.** In multi-agent planning an important problem is that knowledge may turn into false belief for some agents after a partially observable action has taken place. This is a problem for several formalisms for epistemic planning, e.g. dynamic epistemic logic, since agents might not be able to recover from false beliefs. It relates to the general issues of devising appropriate formalisms for doxastic planning (treating beliefs instead of knowledge) and how to deal with belief revision in such settings. The overall goal of the work in this group was to identify the theoretical and computational challenges in planning with knowledge vs. planning with belief; when one or the other is appropriate, or both are needed.

   Guiding questions for the work in this group were: How are knowledge and beliefs represented and distinguished from representations of the actual world? How do we formally handle that knowledge may turn into false belief after a partially observable action has occurred? What are the relevant formalisms for planning with knowledge and/or belief and what are their theoretical and computational properties? How do we deal with belief revision in planning? Are there specific types of interesting goals for epistemic planning? For example, in planning with beliefs, goals can be about making some agents have false beliefs. This necessitates formalizing false-belief tasks, lying and deception.

4. **Practical Tools, Resources and Computational Techniques.** The goal of the group working on this theme was to identify practical tools and resources that facilitate the development and experimental evaluation of automated techniques for epistemic planning.

   Guiding questions for the work in this group were: What tools and computational techniques already exist in epistemic planning? What are the models and formulas used? What are the shortcomings and challenges of these tools and computational techniques? Are there tools or computational techniques from other communities that we are not availing ourselves of to the fullest extent (e.g., for the planning people model checking)? What are the trade-offs between different tools? What are the computational complexities of the different approaches and under different assumptions.

5. **Correspondence Between Planning Problems and Games.** This working group was introduced as an additional discussion topic during the seminar, since several participants found it very relevant and important to epistemic planning. The issue is that (epistemic) game theory studies many of the same problems as (epistemic) planning, but mainly by separate research communities using separate vocabularies. The goal of the group working on this theme was to establish formal connections between the area of automated planning and the area of game theory.

   Guiding questions for the work in this group were: What are the formalisms, tools and results from game theory relevant to automated planning? Symmetrically, what are the formalisms, tools and results from automated planning relevant to game theory? Can problems formulated in one of the settings easily be translated into the other? What is gained and lost in such translations?

6.36 Computational Interactivity

Organizers: Xiaojun Bi, Otmar Hilliges, Takeo Igarashi, and Antti Oulasvirta

Seminar No. 17232

Date: June 5–8, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.48

Participants: Emre Aksan, Leif Azzopardi, Nikola Banovic, Anna Maria Feit, Krzysztof Gajos, Otmar Hilliges, Andrew Howes, Takeo Igarashi, Alec Jacobson, Andreas Karrenbauer, Jun Kato, Yuki Koyama, Per Ola Kristensson, Ranjitha Kumar, Wanyu Liu, Roderick Murray-Smith, Antti Oulasvirta, Daniele Panozzo, Jürgen Steimle, Jerry Talton, Harold Thimbleby, John Williamson

The field of Human Computer Interaction (HCI) as a whole has been tremendously successful in the past both in terms of growth and impact of the premier academic conferences and in reshaping the IT industry. However, as we enter the post-PC era new technologies emerge and bring new challenges along that the traditional user-centered-design approach is not well equipped to meet. For example, artificial intelligence, wearable computing, augmented and virtual reality and custom interactive devices enabled by emerging digital fabrication technologies pose increasingly wicked challenges for interaction design, where designers must consider the entire stack from low-level hardware, through software all the way to the human factors, implying that it is no longer feasible to abstract away technology and hence design spaces that explode in their complexity.

In June 2016 we assembled a group of 25 researchers to discuss aspects relating to a computational view of interactions. Through a series of talks, breakout discussions and panel discussions we established a broad consensus that HCI can and should be approached via a computational lense. We also identified several areas in which computational models are already being used to answer HCI research questions and to move the field forward. However, it became clear that the area is in its infancy and that much work is necessary to turn computational HCI into one of the mainstream approaches in the larger research community. Primarily, researchers and students need to begin thinking in computational terms (abstraction, modelling, automation) and need to learn how to incorporate such thinking into the typically more design driven thinking prevalent in current research. Furthermore, it was also discussed at length how state-of-the-art methods in numerical optimization and machine learning can advance HCI research and likewise how HCI research can identify and refine research requirements in these adjacent research communities.

In terms of concrete outcomes, many of the present researchers agreed to contribute to a forthcoming book on "computational interaction" and to write a joint overview article further refining the discussions and outcomes of the Dagstuhl seminar. In summary, a very fruitful and productive seminar led to interesting and in-depth discussions and provided starting points for much collaborative and community-driven future work. We also identified the need for further community building work including establishing of recurring workshops, symposia, and similar outlets, outreach via summer-schools, tutorials and other educational efforts as well as establishment of a sub-committee at ACM SIGCHI the premier venue for HCI research.

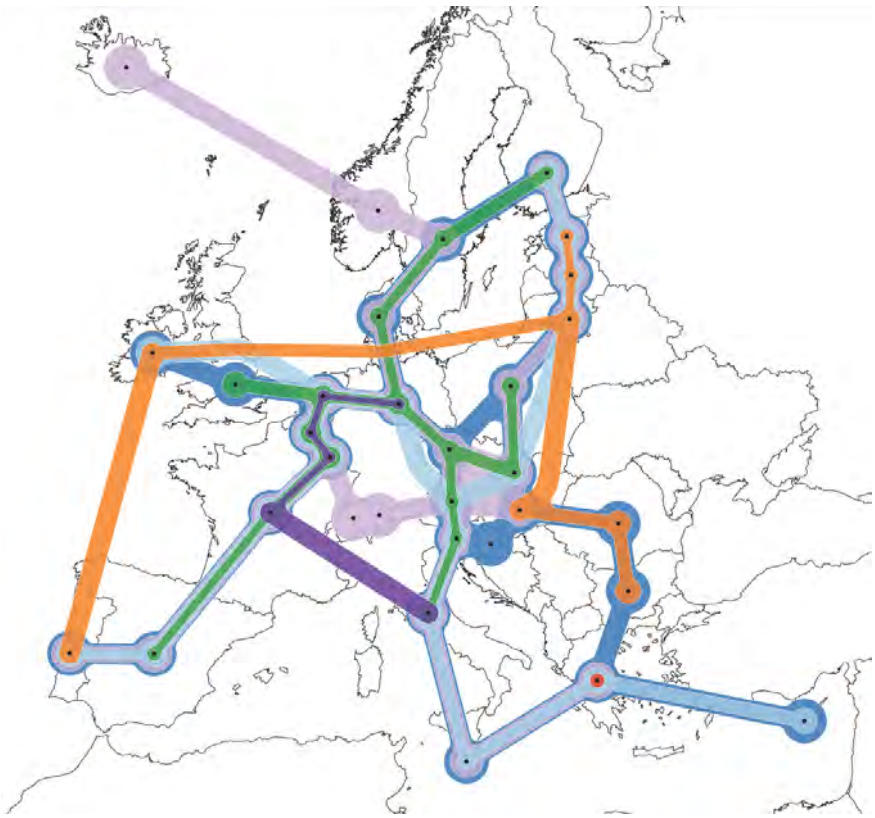

Fig. 6.7

"European countries grouped based on their (past) relations, drawn using an existing technique called KelpFusion. For example, the dark blue shape encloses exactly the EU member states, whereas the light purple one represents the Schengen area".

Press release about Dagstuhl Seminar 17332 – "Scalable Set Visualizations".

https://www.dagstuhl.de/no_cache/en/about-dagstuhl/press/pressemitteilungen/detail/meldung/664/. Photo courtesy of Wouter Meulemans.

6.37 Game Theory Meets Computational Learning Theory

Organizers: Paul W. Goldberg, Yishay Mansour, and Paul Dütting

Seminar No. 17251

Date: June 18–23, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.68

Participants: Yakov Babichenko, Amir Ban, Siddharth Barman, Peter L. Bartlett, Gianluca Brero, Rachel Cummings, Argyris Deligkas, Paul Dütting, John Fearnley, Felix Fischer, Hu Fu, Claudio Gentile, Paul W. Goldberg, Jason D. Hartline, Martin Hoefer, Aleck Johnsen, Thomas Kesselheim, Max Klimm, Sébastien Lahaie, Yishay Mansour, Jamie Morgenstern, Andrés Muñoz Medina, Paresh Nakhe, Denis Nekipelov, Ariel D. Procaccia, Rahul Savani, Sven Seuken, Yaron Singer, Alexander Skopalik, Gregory Valiant, Ellen Vitercik, Bernhard von Stengel, Jens Witkowski, Yair Zick

Algorithmic Game Theory (AGT) has been an identifiable research field for about 20 years by now. It emerged as an important research community in the 1990s, with the ACM-EC conference starting in 1999, and the conferences WINE and SAGT also support this community; in addition, the field is also represented in the main CS theory and AI conferences. Among former Dagstuhl seminars on topics in AGT, there have been a sequence of Dagstuhl seminars on Equilibrium Computation, and another sequence of seminars on Computational Social Choice, and also on Electronic Markets and Auctions.

Machine learning has of course become very pervasive, with the vast accessibility of "big data" and has the motivation to develop new methodologies for harvesting the vast amounts of data, improve our ability to automatically carry out many tasks, from classifying documents and pictures, to identifying normal trends and anomalies.

It is perhaps not surprising that it is timely to investigate the connections between economics and big data, more specifically the interface between game theory and machine learning. Much of econometrics is about handing data and deriving understanding from data.

In the AGT context, this would seem to apply most readily to data emanating from "economic" sources, and of course there are plenty of examples. The most notable of these is learning user preferences from examples.

There have been workshops at the AGT/Machine Learning interface in the ACM-EC conference (the 2017 EC held the 3rd workshop on Algorithmic Game Theory and Data Science) but there is clearly space for more meetings on this topic, and this is the first one to take place at Dagstuhl. As such, it contributed to the development of the European community in this research area (noting that ACM-EC is usually held in the USA).

It was pleasing that the seminar attracted quite a high proportion of participants who were visiting Dagstuhl for this first time, alongside others who have made multiple visits. There was a good balance amongst representatives of Algorithmic Game Theory, Machine Learning, and Economics.

One can classify the AGT/Machine Learning topics as follows.

- Usage of ML ideas (reinforcement learning, multi-arm bandits, etc.) into decision making under uncertainty (and the search for game-theoretic solution notions such as equilibria)
- usage of game-theoretic tools into machine learning approaches (as in Generative Adversarial Nets).
- A basic test case is learning user valuations from historical data. For example, given the outcome of previous auctions to learn the distribution of the users' valuations and the goal is to define near optimal mechanisms. (This is also an aspect in learning "revealed preferences".)
- Query complexity of solution concepts of games, aspects of which are applicable to learning adversary preferences in the context of security/patrolling games.

The seminar was structured around longer invited/tutorial talks (typically lasting 1.5-2 hours), one or two such talks taking place each day. These were followed by shorter contributed talks.

We thank Argy Deligkas for serving as collector of the abstracts.

Keynote/tutorial talks

Monday Sven Seuken, University of Zurich Design of Machine Learning-Based Mechanisms; Yaron Singer, Harvard University: Learning, Optimization, and Noise

Tuesday Claudio Gentile, Universita dell'Insubria: No Regret and Sequential Prediction

Wednesday Denis Nekipelov, University of Virginia: Robust Inference for Non-Robust Models

Thursday Jamie Morgenstern, University of Pennsylvania: The Sample Complexity of Single-Parameter Auction Design

Friday Yakov Babichenko, Technion: Informational Bounds on Equilibria, and its Relation to Learning

Related topics not covered There is ongoing work at the intersection of macroeconomics and machine-learning techniques, that was out of scope of this meeting but may be of interest later. For example, ongoing work on Vector Autoregressive models in the context of multivariate time series modelling, which may later lead to interesting problem in computational learning theory. Another topic that was only touched-on is agent-based models of macroeconomics.

6.38 Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction

Organizers: Rolf Backofen, Yael Mandel-Gutfreund, Uwe Ohler, and Gabriele Varani

Seminar No. 17252

Date: June 18–21, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.86

Participants: Frédéric Allain, Rolf Backofen, Benedikt Beckman, Janusz Bujnicki, Tomaž Curk, Philipp Drewe, André Gerber, Mahsa Ghanbari, Tim Hughes, Eckhard Jankowsky, Hilal Kazan, Grzegorz Kudla, Markus Landthaler, Martin Lewinski, Yael Mandel-Gutfreund, Hannah Margalit, Annalisa Marsico, Daniel Maticzka, Irmtraud Meyer, Quaid Morris, Uwe Ohler, Teresa Przytycka, Andres Ramos, Guido Sanguinetti, Michael Sattler, Peter F. Stadler, Gabriele Varani, Gene Yeo

All living organism must be able to differentially regulate the expression of genes encoded in their genome. Genes are first transcribed into RNA, which are either translated to proteins or functionally active as non-coding RNAs. Beside the direct regulation of the transcription of DNA into RNA, an important additional layer is the direct regulation of RNAs by RNA binding proteins (RBPs). This layer of regulation controls cellular decisions as part of gene expression networks composed of both proteins and RNAs. While being a dark matter of the cell for a long time, recent years have shown the development of sophisticated high throughput experimental technologies that greatly increased our understanding of protein-RNA recognition and regulation. Nevertheless, the quantitative molecular understanding of the transcriptome-level processes remains very limited. Especially complexity **(both in the form of data as in the required computational approaches)** limits the exploitation of these advances towards a quantitative understanding of post-transcriptional regulation. The objective of the seminar to discuss urgently needed computational approaches allowing to exploit the wealth of new data. More specifically, the seminar focused on

- addressing major computational challenges in this field
- mining the extensive genomic information on RNA and associated proteins
- investigation of RNA-protein interactions on an atomic level
- quantitative prediction of cellular regulatory networks and their dynamics.

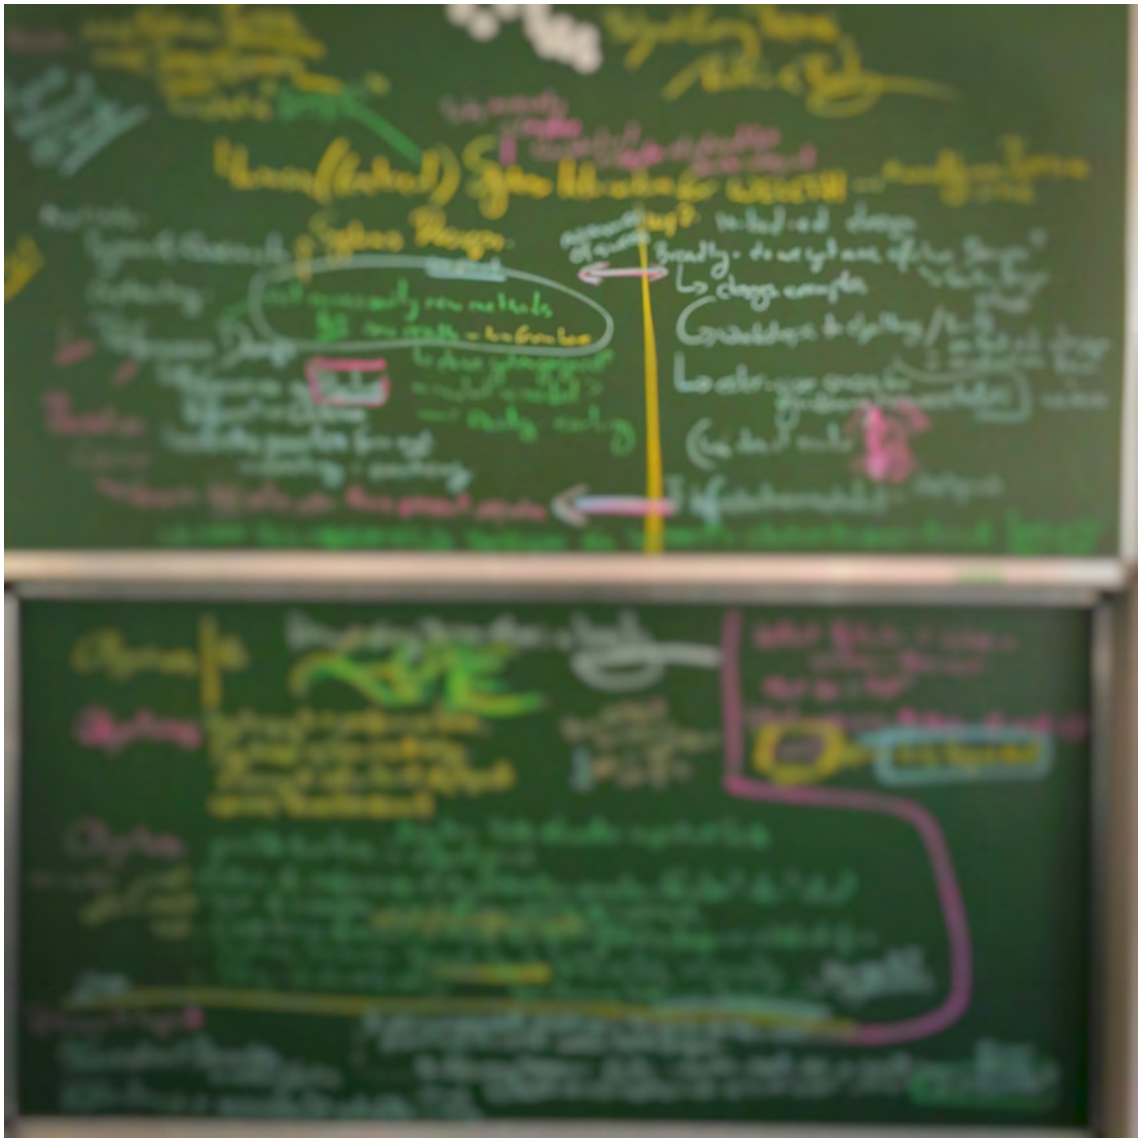

Fig. 6.8
“a two board problem - i thought it might b 3 but it worked out to 2”. Picture on flickr by 17392 Dagstuhl Seminar participant m. c. schraefel. https://flic.kr/p/YYBWsp. Photo courtesy of m. c. schraefel.

6.39 Voting: Beyond Simple Majorities and Single-Winner Elections

Organizers: Dorothea Baumeister, Piotr Faliszewski, Annick Laruelle and Toby Walsh

Seminar No. 17261

Date: June 25–30, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.109

Participants: Haris Aziz, William Bailey, Dorothea Baumeister, Sirin Botan, Sylvain Bouveret, Steven J. Brams, Felix Brandt, Robert Bredereck, Markus Brill, Jiehua Chen, Marcin Dziubinski, Edith Elkind, Ulle Endriss, Piotr Faliszewski, Judy Goldsmith, Bernard Grofman, Marc Kilgour, Werner Kirsch, Christian Klamler, Martin Lackner, Jérôme Lang, Annick Laruelle, Jean-Francois Laslier, Andrea Loreggia, Nicholas Mattei, Vincent Merlin, Stefan Napel, Svetlana Obraztsova, Dominik Peters, Clemens Puppe, Lisa Rey, Jörg Rothe, M. Remzi Sanver, Ann-Kathrin Selker, Sujoy Sikdar, Piotr Skowron, Arkadii Slinko, Krzysztof Sornat, Nimrod Talmon, Kristen Brent Venable, Gerhard J. Woeginger, William S. Zwicker

Computational social choice is an interdisciplinary field of research, focused on the issue of aggregating preferences of agents–perhaps self-interested and strategic – and providing them with joint decisions. Computational social choice combines the tools and approaches of social choice theory, computer science (with particular focus on artificial intelligence and theoretical computer science), economics, political science, and operations research. The distinctive feature of computational social choice–as opposed to the classic social choice theory – is that computational considerations (e.g., efficiency of computing outcomes of the preference aggregation processes) are given significant attention. Further, researchers working on computational social choice often study virtual elections where either people vote through electronic means (such as in Doodle polls for scheduling of meetings) or the elections are used as a tool and the votes are derived automatically in some way (e.g., voting can be used as a selection procedure in a genetic algorithm). Nonetheless, the two research areas are deeply connected and there is significant interaction between them.

One of the most classic problems studied within (computational) social choice regards conducting a single-winner election. For example, consider the situation where members of some society wish to choose their president. They collect the set of candidates (who usually have to register well ahead of time, typically by gathering necessary popular support), then the candidates run their campaigns, argue who would be the best president, air commercials, etc. Eventually, the voters form their preferences (either they simply decide who would be the best president, or they form rankings of the candidates, or they decide on a set of acceptable presidents, depending on the voting rule used) and, on the election day, they cast their votes. In the end, an electoral commission gathers the votes and applies an agreed upon voting rule to decide who would be the next president.

Due to the fantastic progress in social choice (since the middle of the twentieth century) and in computational social choice (over the last fifteen years or so), essentially all the stages of the above-described process are quite well understood. However, In the modern world – especially in the era of ubiquitous use of social media – it appears that there is a great range of preference aggregation settings where the classic approach falls short. For example, consider a situation where a company wants to hire a team of specialists. There may be quite a large number of possible candidates to employ (as opposed to the few candidates in a typical presidential election), each of the candidates may have quite different skills and abilities, various employees may either complement each other or be counter-productive if teamed up (as opposed to our presidential election, where we pick a single person for the whole task). Finally, the recruiting committee typically consists of just a few people (few voters, as opposed to the millions of people voting for presidents), but their preferences might have a very involved structure (for example, if we hire a specialist in X then we also need a specialist in Y, but otherwise a specialist in Z would suffice; on top of that, each member of the recruitment committee may judge candidates' abilities differently).

The goal of the seminar was to bring together researchers who work on various aspects of aggregation problems that go beyond the classic high-stake, rarely conducted, single-winner elections, and to discuss the following issues:

1. Scenarios with multiple winners and/or settings where each aggregation outcome may consist of separate entities (e.g., multi-winner elections and single-winner elections in combinatorial domains).
2. Various non-typical ways of expressing preferences, going from (variants of) non-binary preferences to settings where

the agents can express complex statements, including conditional ones (e.g., CP-nets).

The seminar also dealt with some real-world applications, such the question of drawing constituency boundaries in the United States of America or the choice of the voting rule in EU Council of Ministers.

The seminar brought together 42 researchers from 14 countries, working in artificial intelligence, theoretical computer science, mathematics, economics, social choice, and political science. Discussions regarding the new challenges in the area of preference aggregation should fertilize research in all these areas.

The technical program of the seminar was structured over five working groups. On the first day the participants were invited to give 5-minute presentations of research topics that they found interesting and, based on those presentations, the organizers suggested five working groups:

- Working Group 1: Voting Experiments.
- Working Group 2: Understanding Diversity in Multiwinner Elections.
- Working Group 3: Aggregation Procedures with Nonstandard Input and Output Types.
- Working Group 4: Voting in Larger Contexts.
- Working Group 5: Proportionality in Multiwinner Elections.

The seminar attendees accepted these groups and each chose one to participate in the discussions. During the seminar each group met twice for extended discussions. Also, one afternoon was free for unstructured discussions (of which many were used for in-depth discussions of topics initiated during the working group discussions). Finally, on the last day of the seminar representatives of each group presented the results of their discussions (which ranged from making actual technical contributions to presenting research agendas for future work[47]). The working groups were supported by 21 regular scientific presentations, and extended presentation of the real-life experiments regarding French presidential elections, and 6 survey talks.

The seminar acknowledged that there are new, exciting research topics regarding computational social choice. In particular, multiwinner elections were represented very prominently and it became clear that work on them has only started. We hope and believe that one of the effects of the seminar was convincing many more people (also outside of the core computational social choice community) that studying more general forms of elections (such as multiwinner elections, or elections with nonstandard input formats) is important and promising.

Given the personal feedback we received, we believe that the participants were very happy with the setting of working groups (from the process of forming them based on surveyed interest, through actual meetings, to presentation of results). We also receive very strong, positive feedback regarding leaving one afternoon unstructured, for the participants to use as they preferred. We have seen many ad-hoc discussions, meetings of coauthors, and problem-solving sessions. Indeed, it seems that for seminars with well-established communities, such an unstructured afternoon is far more useful than the traditional excursion (which, on the other hand, seems to be very effective for not as well-established communities).

We are very grateful to all the participants for their contributions, ideas, and discussions, which made this seminar truly enjoyable. We would also like to thank the Schloss Dagstuhl team for their support and excellent organization and patience.

[47] The organizers are already aware of three independent research projects that follow the agendas presented at the seminar.

6.40 Federated Semantic Data Management

Organizers: Olaf Hartig, Maria-Esther Vidal, and Johann-Christoph Freytag

Seminar No. 17262

Date: June 25–30, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.6.135

Participants: Maribel Acosta, Bernd Amann, Sören Auer, Abraham Bernstein, Piero Andrea Bonatti, Carlos Buil-Aranda, Emanuele Della Valle, Michel Dumontier, Johann-Christoph Freytag, Claudio Gutierrez, Peter Haase, Olaf Hartig, Jana Hentschke, Katja Hose, Sabrina Kirrane, Stasinos Konstantopoulos, Jorge Lobo, Pascal Molli, Gabriela Montoya, Themis Palpanas, Ana Maria Roxin, Sherif Sakr, Juan F. Sequeda, Hala Skaf-Molli, Rudi Studer, Joachim Van Herwegen, Maria-Esther Vidal

The Semantic Web is an extension of the World Wide Web in which *structured data and its meaning* is represented in a form that can be readily accessed and exploited by machines. The foundation of this representation is a graph-based data model defined by the Resource Description Framework (RDF). This framework allows for data management approaches that focus on manipulating and using data in terms of its meaning. We refer to this type of data management as *semantic data management*.

In addition to centralized access to RDF datasets, Web-based protocols such as the SPARQL protocol enable software clients to access or to query RDF datasets made available by remote servers. By integrating such remote data sources as members of a *federated system*, software clients may answer cross-dataset queries without having to retrieve various datasets into a single repository. Given such a federation, the complexity of problems of query processing and semantic data management increases due to additional parameters such as variable data transfer delays, a changing availability of federation members, the size of the federation, and distribution criteria followed to place and semantically link data in different datasets of the federation. Moreover, whenever data is replicated across federations, synchronization is required to ensure that all changes are propagated and the semantics of data is preserved. Despite a large number of technologies developed by the Semantic Web and Database communities to address problems of semantic data management, we still observe a significant lack of efficient and effective solutions to the problems of federated semantic data management (FSDM), which prevents the development of real-world applications on top of Semantic Web technologies. Additionally, existing proposals to evaluate such solutions do not sufficiently cover the large number of parameters that affect FSDM and the complexity of tradeoffs. More specifically, variables and configurations that considerably affect the federated semantic data management problems are not sufficiently defined or even considered in state-of-the-art testbeds (e.g., network latency, data fragmentation and replication, query properties, or frequency of updates).

The aim of the Dagstuhl seminar was to gather experts from the Semantic Web and Database communities, together with experts from application areas, to discuss in-depth open issues that have impeded FSDM approaches to be used on a large scale.

The following crucial questions were posed as a basis for the discussions during the seminar:

Q1 *Can traditional techniques developed for federations of relational databases be enriched with RDF semantics, and thus provide effective and efficient solutions to problems of FSDM?*

Q2 *What problems of FSDM present new research challenges that require the definition of novel techniques?*

Q3 *What is the role of RDF semantics in the definition of the problems of FSDM?*

To discuss these questions the participants of the seminar were grouped according to their areas of expertise and interests. In particular, the seminar focused on four main topic areas (see below). The results of the group discussions were presented in plenary sessions and will be compiled into manuscripts with which the seminar outcomes will be disseminated. As a basis of the group work, and to establish a common understanding of key concepts and terminology, the seminar included a few short, survey-style talks on a number of related topics. In particular, these talks covered:

- "Graph data models and graph databases" (given by Olaf Hartig),
- "RDF and semantics" (by Claudio Gutierrez),
- "Policies and access control" (by Sabrina Kirrane and Piero Andrea Bonatti),
- "Database privacy" (by Johann-Christoph Freytag),
- "Distributed database systems" (by Katja Hose), and
- "Federated query processing" (by Maria-Esther Vidal).

In addition to these survey talks, every participant was given the chance to briefly highlight their research as relevant for the seminar. Moreover, in a demo session, some of the participants showcased their FSDM-related systems and tools, which gave interested attendees of this session an opportunity to play with and better understand these systems and tools. The systems and tools demonstrated in this session were the following:

- *Triple Pattern Fragments client* that runs in a browser and executes queries over a federation of Triple Pattern Fragment (TPF) interfaces (demonstrated by Joachim Van Herwegen),
- *Network of Linked Data Eddies* (*nLDE*), an efficient client-side SPARQL query engine for querying server-side data that can be accessed via a TPF interface (demonstrated by Maribel Acosta),
- *Ladda*, a framework for delegating TPF-based query executions among multiple browsers (demonstrated by Hala Skaf-Molli),
- *Quartz*, a system for querying replicated Triple Pattern Fragments (demonstrated by Hala Skaf-Molli),
- *Ontario*, a federated SPARQL query engine for heterogeneous sources represented in different raw formats (demonstrated by Maria-Esther Vidal),
- *UltraWrap*, a framework for integrating relational databases using SPARQL federation (demonstrated by Juan Sequeda),
- *Ephedra*, a SPARQL federation engine that combines SPARQL services with other services (demonstrated by Peter Haase),
- *JedAI*, an entity resolution toolkit (demonstrated by Themis Palpanas), and
- *Exemplar Queries*, a framework for query answering using knowledge graphs (demonstrated by Themis Palpanas).

As mentioned before, besides the short survey talks, the demos, and the participants' presentations, the major focus of the seminar was on discussions in four working groups, where each of these groups addressed a different topic area. The remainder of this section provides a brief overview of the four topic areas covered by the groups and the respective results. More detailed summaries provided by each of the four groups can be found in a separate section of this report.

Graph Data Models. Graph data models such as the RDF data model allow for a representation of both data and metadata using graphs of nodes that represent entities, and edges that model connections between entities. Graph data management encompasses techniques for managing, querying, and analyzing graph data by utilizing graph-oriented operations. SQL-like query languages have been defined for evaluating declarative queries over graph data; additionally, well-known algorithms are utilized for computing graph invariants (e.g., triangle counting or degree centrality) and for solving typical graph problems (e.g., finding shortest paths, traversals, or dense subgraphs). Furthermore, several real-world applications have been built on top of existing graph-based tools (e.g., community detection, centrality analysis, and link prediction). Graphs naturally represent a wide variety of domains (e.g., social networks, biological networks) in which data, interconnectivity, and data topology all are first-class citizens, with RDF data being one example of graph data.

During the Dagstuhl seminar, a working group was formed to discuss whether tools for graph data management are sufficient to model and to manage the semantics in RDF data, taking into account that characteristics of the RDF data model (e.g., blank nodes and SPARQL operators) may affect tractability of the graph-based tasks in a federation of RDF graphs. As a first result of this discussion, the working group made the following observation. In contrast to other graph data models and query languages, the RDF data model is a "universal" data model in the sense that it is designed for sharing data and knowledge in an unbounded space such as the Web. To continue the discussion, the group introduced a definition of the notion of FSDM and identified five principles that characterize FSDM: universality, unboundedness, dynamicity, network protocols, and semantics. Based on further discussion that took into account these principles, the group made two conjectures that they plan to elaborate on in a future publication and that can be summarized as follows. First, it is impossible to build a FSDM system that fully achieves universality, unboundedness, and dynamicity, all at the same time. Second, the concepts of federation and semantics are interdependent and must be tackled together to develop effective and efficient solutions for building FSDM systems.

Federated Query Processing. A vast number of approaches have been developed to provide a unified interface for querying federations of data sources. In the context of federations of RDF datasets, existing approaches focus on two problems: the problem of selecting the RDF datasets required to execute a federated query, and the problem of executing the resulting sub-queries efficiently against the selected data sources. Although federated query processing has been studied extensively, a number of important problems are still open, and more challenges are likely to come up as the complexity of federations increases (e.g., by increasing numbers of federation members, by replication and fragmentation of RDF data, and by federation members that update their RDF data autonomously).

During the Dagstuhl seminar, a working group was formed to discuss the problem of federated query processing over RDF data sources. Challenges imposed by the semi-structured nature of RDF, unpredictable behavior and dynamicity of Web-accessible RDF sources, and the role of the entailment regimes guided the group discussions and allowed for enumerating the main differences with the problem of federated query processing against relational databases. The group focused on the formal definition of the problem, as well as on the formalization of the subproblems of source selection, query decomposition, and query execution. As a first result, the group identified that the entailment regimes to be performed over a federation of RDF sources, as well as data replication and dynamicity, access control policies, and SPARQL query capabilities, play a crucial role in source selection, query decomposition, and query execution. State-of-the-art techniques implemented by existing approaches (e.g., FedX, ANAPSID, or Linked Data Fragments) were discussed and compared based on this formalization; the group concluded that none of existing approaches takes into account all these characteristics of RDF data sources, being required further analysis and work to empower them to solve the formalized problems. Finally, the impact of these characteristics on the performance of SPARQL operators (e.g., join, union, or optional) was discussed. The group concluded that although physical operators implemented by existing approaches are capable of adjusting query execution schedulers to RDF source availability, they are unable to adapt their execution to other RDF source characteristics, e.g., supported entailment regimes or data evolution. These issues remain open as well, and require further study from the semantic data management community.

Access Control and Privacy. Solutions to the problem of modeling access control policies for Web resources have been benefited from Semantic Web technologies. Existing rule-based logic languages rely on ontology-based reasoning tasks to represent reactive policies for access control, and to enforce and to propagate trusted and policy-compliant interactions across resources

in RDF datasets. For instance, the Open Digital Rights Language (ODRL) is a rule-based approach that allows for a description of policies to access and to exchange Web resources. Nevertheless, as per the Linked Data publishing principles, RDF properties associated with any resource can be accessed by de-referencing their corresponding URL. In applications of domains of FSDM such as personalized medicine or finances, only authorized and privacy-respecting access is allowed. Thus, novel approaches are required to bridge the gap between access-control models and unrestricted access to RDF resources.

A working group with a focus on access control and privacy discussed the following open issues: a) formalisms to specify access-control and privacy policies of federation resources and to reason over the meaning of these resources; and b) techniques that enable systems to enforce privacy-aware and security-aware policies whenever a resource is accessed. After concluding that there are too many open challenges to be all solved immediately, the group agreed to focus on access control. Next, the group discussed conceptual access control models and achieved a better understanding of requirements of a conceptual framework to analyze policy-aware federated Semantic Web architectures. Finally, the group defined such a framework and made plans for a publication about it.

Use Cases and Applications. In addition to the first three working groups that focused on various more technical aspects of building FSDM systems, a fourth working group looked into applications of FSDM and use cases in which adopting FSDM would be beneficial. Specifying such use cases, as well as documenting the usage of FSDM systems in existing applications, is important to better understand the requirements and the challenges of FSDM and to derive realistic testbeds for approaches to build FSDM systems.

A key observation of the work group was that approaches to apply FSDM can be categorized into two classes depending on whether they focus i) on explorative, open-domain querying or ii) on controlled, close-domain querying. Then, the working group identified a broad set of general use cases of FSDM. Thereafter, the group defined a framework for developing specific use cases. This framework introduces a set of requirements for the specification of a use case. Finally, the group applied their framework to develop a number of example use cases.

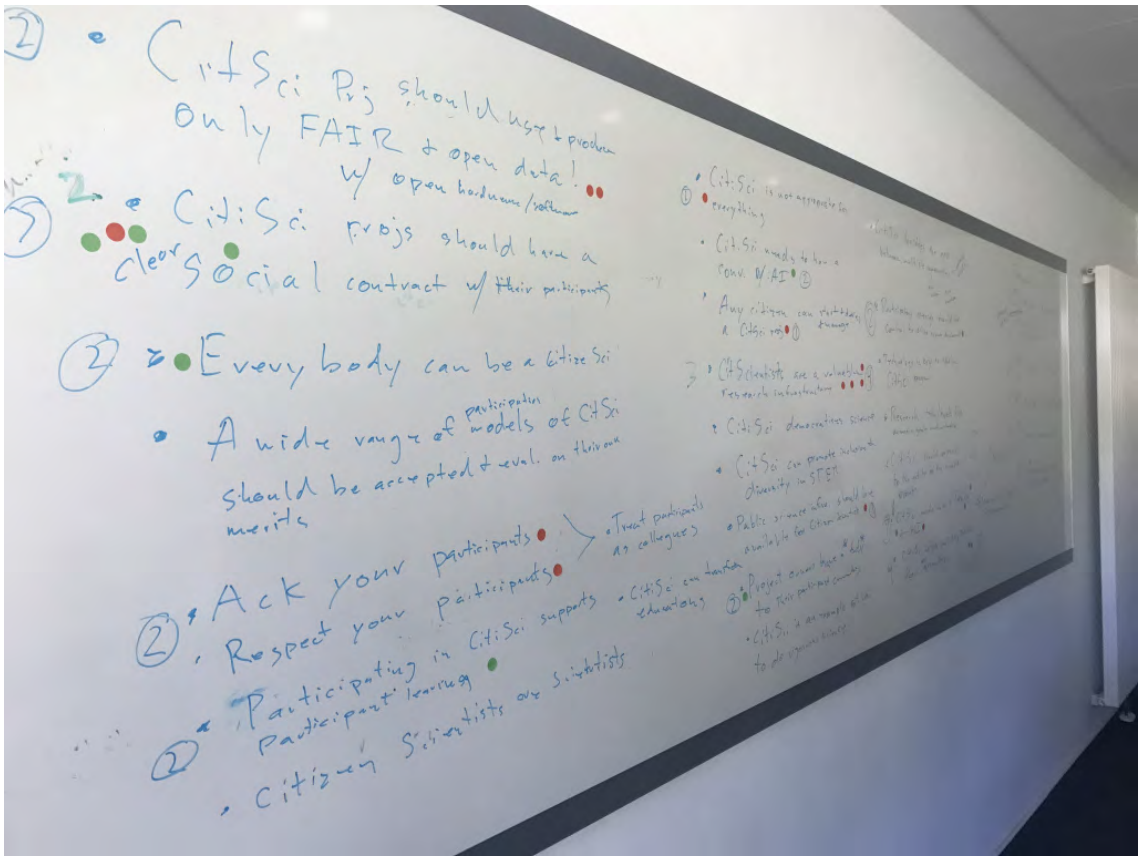

Fig. 6.9
"Trip Report: Dagstuhl Seminar on Citizen Science". Blog post by 17272 Dagstuhl Seminar participant Paul Groth. https://thinklinks.wordpress.com/2017/08/02/trip-report-dagstuhl-seminar-on-citizen-science/. Photo courtesy of Paul Groth.

6.41 Foundations of Wireless Networking

Organizers: Christina Fragouli, Magnús M. Halldórsson, Kyle Jamieson, and Bhaskar Krishnamachari

Seminar No. 17271

Date: July 2–7, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.1

Participants: Giuseppe Caire, Izzat Darwazeh, Michelle Effros, Anja Feldmann, Christina Fragouli, Seth Gilbert, Deniz Gündüz, Martin Haenggi, Magnús M. Halldórsson, Longbo Huang, Kyle Jamieson, Tomasz Jurdzinski, Holger Karl, Christina Kolb, Bhaskar Krishnamachari, Fabian Daniel Kuhn, Muriel Médard, Konstantinos Nikitopoulos, Christian Scheideler, Christian Schindelhauer, Emina Soljanin, Vijay Subramanian, Patrick Thiran, Tigran Tonoyan, Daniela Tuninetti, Elif Uysal-Biyikoglu, Zhiying Wang, Roger Wattenhofer, Jörg Widmer, Dongxiao Yu, Michele Zorzi, Marco Antonio Zúñiga Zamalloa

Wireless communication has grown by leaps and bounds in recent decades, with huge social and societal impact. This is nowhere near saturation, especially with the coming Internet-of-Things on the horizon.

Underlying this technology are fundamental questions: how to efficiently configure and adapt communication links, organize access to the medium, overcome interference, and disseminate information. Unfortunately, the wireless medium is tricky, and the modeling of signal propagation and interference has proved to be highly involved, with additional challenges introduced by issues such as mobility, energy limitations, and device heterogeneity. New technologies such as cooperative MIMO, directional antennas, interference alignment, network coding, energy harvesting and motion control add another layer of complexity, and are yet to be well-understood. Deriving good algorithms and protocols is therefore a non-trivial task.

Communities

Different schools of thought have arisen to tackle these fundamental questions. These come from different backgrounds, involving different types of mathematical tools and different approaches and outlooks. It is not just a theory vs. practice split, but also splits within the theory and the practice camps. In addition to the more established information theory, there has been quite some work on network control theory, and also a budding algorithmic theory; on top of physical-layer hardware experimentation, we see also networking systems research and simulation studies.

We identified the following communities, which are not all exclusive:

Information Theory Characterized by an interest in fundamental information-theoretic capacity bounds; novel communication paradigms such as MIMO, network coding, interference cancellation, interference alignment; estimation and detection under known stochastic models of noise.

Prime publication venues: IEEE Trans. on Information Theory; IEEE International Symposium on Information Theory (ISIT).

Algorithm Theory Characterized by a focus on algorithms, their complexity and effectiveness, with emphasis on rigorous proofs and typically worst-case analysis.

Prime publication venues: PODC, DISC, STOC, SODA, ICALP (Track C)

Experimental Mobile and Wireless Systems Characterized by the design, implementation, and evaluation of practical wireless systems in real testbeds and real-world applications that both both evaluate the efficacy of previously-known techniques and their combinations "in the wild" and add design insight by developing novel heuristic algorithms and architectures that are shown to perform well in practice.

Prime publication venues: MobiCom, MobiHoc, SIGCOMM, NSDI, SenSys, IPSN, BuildSys, MobiSys.

Wireless Network Control and Optimization Theory Characterized by formulation of various problems in wireless networks (typically focusing on medium access, network routing and flow rate control) as continuous control and optimization problems from both deterministic and stochastic perspectives; network utility optimization via primal-dual decomposition, game theory, stochastic decision theory, stochastic multi-armed bandits.

Prime publication venues: Infocom, SIGMETRICS, WiOpt, CDC, IEEE Trans. on Networking, IEEE Transactions on Automatic Control.

Physical Layer and Hardware Design Characterized by the design, implementation, and evaluation of new hardware, signal processing techniques. The theoretical members of

this community have a lot of overlap with information theory, while the more experimental members of this community have a lot of overlap with the experimental wireless and mobile systems community in terms of the problems they consider and their solution approaches.

Prime publication venues: IEEE Trans. on Wireless, IEEE Trans. on Comms., IEEE Globecom, IEEE Vehicular Technology Conf.

Goals of the Seminar

The goal of this Dagstuhl seminar is to bring together top researchers from the different wireless research communities to review and discuss models and methods in order to obtain a better understanding of the capabilities and limitations of modern wireless networks, and to come up with more realistic models and new algorithm and protocol design approaches for future wireless networks that may then be investigated in joint research projects.

An important part of the workshop is to actively promote a dialog between different communities. As a result, we seek researchers that are by nature open to different perspectives and have enough self-confidence to welcome research of a different nature. The objective was for each participant to consciously reflect on the implicit values, identity, shared understandings and skill set that people in his/her community expect, and to articulate these issues to others in order to identify and appreciate commonalities and differences, and the potential gains from forming new bridges.

Seminar Operation

The seminar had varied forms of activities during its operation. As people came from different communities, a major objective was to get to know each other.

"Speak-and-spark" presentations All the participants gave a brief, 5-10 minute pitch talk on a problem that they have been (or would to be) working on. Most of these were given on the first day, which provided a way of introducing one another, as well as a way to spark discussions that could be continued in private, in small groups, or in plenum.

Survey presentations Seven senior researchers were asked to give a one-hour survey talk on a topic of current interest. These were spread over the days excluding the first.

Breakout sessions Several topical issues were identified as particularly suitable for group discussions. The participants voted on the topics of their interest, after which three were selected. The groups were chosen so as to feature representatives from the different communities. Two such rounds of breakout sessions were organized on Tuesday. The discussions were summarized by the group leaders (often with the help of the scribes) for the whole audience on Wednesday morning, followed by discussions.

"Important paper" pitches The participants were encouraged to identify research paper(s) that open "new" research areas and/or pose questions in their community. This was also a means to articulating what researchers in that subfield found essential or influential. These were the presented in 5-10 minutes, followed by open questions.

Plenary discussions Part of the last day's morning was allocated to general discussions around the themes posed during the seminar, with the aim of identifying future problem directions and research areas.

6.42 Citizen Science: Design and Engagement

Organizers: Irene Celino, Oscar Corcho, Franz Hölker, and Elena Simperl

Seminar No. 17272

Date: July 2–5, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.22

Participants: Lora Aroyo, Alessandro Bozzon, Maria Antonia Brovelli, Irene Celino, Oscar Corcho, Dominic di Franzo, Claudia Göbel, Esteban González Guardia, Paul Groth, Lynda Hardman, Franz Hölker, Tomi Kauppinen, Christopher Kyba, Dave Murray-Rust, Maurizio Napolitano, Jasminko Novak, Christopher Phethean, Marisa Ponti, Lisa Posch, Gloria Re Calegari, Neal Reeves, Sven Schade, Sibylle Schroer, Elena Simperl, Alice Verioli, Christopher A. Welty, Andrea Wiggins, Amrapali Zaveri

Citizen science is an approach to science that is enlisting the help of millions of volunteers across a range of academic disciplines to complete tasks that would have otherwise been unfeasible to tackle using expert time or computational methods [2]. While it is a popular and effective way to solve various problems, with many examples of incredible success [1, 3], there remains a number of ongoing challenges that must be addressed in order to ensure the validity of citizen science as a widespread approach to research. The aim of this workshop – organised in partnership with the SOCIAM[48] and Stars4All[49] projects – was to discuss and explore aspects of the future of citizen science, focusing on design factors and engagement strategies, although this naturally required a holistic assessment of citizen science projects, platforms and applications as a whole.

References

1 Jeehung Lee, Wipapat Kladwang, Minjae Lee, Daniel Cantu, Martin Azizyan, Hanjoo Kim, Alex Limpaecher, Snehal Gaikwad, Sungroh Yoon, Adrien Treuille et al. *RNA Design Rules from a Massive Open Laboratory*. Proceedings of the National Academy of Sciences, National Academy of Sciences, 111; 6, 2122–2127, 2014.

2 Chris J Lintott, Kevin Schawinski, Anže Slosar, Kate Land, Bamford Steven, Daniel Thomas, Jordan M Raddick, Robert C Nichol, Alex Szalay and Dan Andreescu et al. *Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey*. Monthly Notices of the Royal Astronomical Society, Blackwell Publishing Ltd Oxford, UK, 389; 3, 1179–1189, 2008.

3 Chris J Lintott, Kevin Schawinski, William Keel, Hanny Van Arkel, Nicola Bennert, Edward Edmondson, Daniel Thomas, Daniel JB Smith, Peter D Herbert, Matt J Jarvis et al. *Galaxy Zoo:'Hanny's Voorwerp', a quasar light echo?*. Monthly Notices of the Royal Astronomical Society, Blackwell Publishing Ltd Oxford, UK, 399; 1, 129–140,2009.

[48] https://sociam.org/about
[49] http://stars4all.eu/

6.43 Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition

Organizers: Sarah Zennou, Saumya K. Debray, Thomas Dullien, and Arun Lakhotia

Seminar No. 17281

Date: July 9–14, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.44

Participants: Radoniaina Andriatsimandefitra, Sebastian Banescu, Thomas Barabosch, Sébastien Bardin, Konstantin Berlin, Paul Black, Cory Cohen, Christian Collberg, Sophia D'Antoine, Mila Dalla Preda, Robin David, Bjorn De Sutter, Saumya K. Debray, Thomas Dullien, Roberto Giacobazzi, Yuan Xiang Gu, Tim Kornau-von Bock und Polach, Arun Lakhotia, Colas Le Guernic, Jean-Yves Marion, Marion Marschalek, J. Todd McDonald, Xavier Mehrenberger, Michael Meier, Bogdan Mihaila, Craig Miles, Asuka Nakajima, Daniel Plohmann, Mario Polino, Pablo Rauzy, Raphael Rigo, Radwan Shushane, Natalia Stakhanova, Ryan Stortz, Dinghao Wu, Yves Younan, Stefano Zanero, Sarah Zennou, Ed Zulkoski

As a follow-up on the previous Dagstuhl Seminar 14241 on the analysis of binaries, the interest in attending this new seminar was very high. The attendance was very diverse, almost half academics and half practitioners.

Talks were arranged by topics and each day ended with an open discussion on one of the three topics: machine learning, obfuscation and practitioners' needs.

Considering the given talks, it appears that the challenges in the realm of general binary analysis have not changed considerably since the last gathering. However, the balance between the topics shows that the academic interest is now more focused on machine learning than on obfuscation. On the contrary practitioners exhibited examples showing that the sophistication level of obfuscations has tremendously increased during this last years.

The open discussions were the most fruitful part of the seminar. The discussions enabled the academics to ask practitioners about the hypotheses that are relevant to build models for their analyses and the problems they face in their daily work. The practitioners gained awareness of the automated tools and techniques that they can expect to see emerge from research labs.

These informal exchanges will be gathered into a separate document and spread to the academic community.

Finally please note that not all people who presented have submitted their abstracts due to the sensitive nature of the content and/or the organization that the participants work for.

6.44 From Observations to Prediction of Movement

Organizers: Mark Birkin, Somayeh Dodge, Brittany Terese Fasy, and Richard Philip Mann

Seminar No. 17282

Date: July 9–14, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.54

Participants: Sean Ahearn, Mark Birkin, Kevin Buchin, Maike Buchin, Urska Demšar, Somayeh Dodge, Damien Farine, Brittany Terese Fasy, Robert Holbrook, Maarten Löffler, Jed Long, Robin Lovelace, Richard Philip Mann, Samuel A. Micka, Harvey J. Miller, Jennifer Miller, David Millman, Nicholas Ouellette, Kristine Pelatt, Andrea Perna, Kamran Safi, Jack Snoeyink, Frank Staals, Kathleen Stewart, Daniel Strömbom, Alexander Szorkovszky, Johan van de Koppel, Carola Wenk, Zena Wood

Dagstuhl Seminar 17282 took place at Schloss Dagstuhl from 9 to 14 July 2017. We had 29 participants and nine invited talks. The main theme of this seminar was the analysis and prediction of movement trajectories. In particular, we focused on the study of movement patterns of individuals, and the interactions of moving agents with each other and with the environment.

Themes Movement analysis is key to understanding the underlying mechanisms of dynamic processes. Movement occurs in *space* and *time* across *multiple scales* and through an embedding *context* that influence how entities move. The importance of spatiotemporal aspect of movement has attracted a wide range of studies. Analysis of movement trajectories is a core element of Movement Ecology in Biology, as well as being important across disciplines as diverse as Geographic Information Sciences (GIS), Transportation, Criminology, Epidemiology, Computer Gaming, and Phylogenetics. Development of efficient algorithms for analyzing and predicting will be vital to realizing the hopes for new generation smart transport systems and smart cities. Furthermore, naturally generated trajectories provide a fascinating context for mathematical and computational study of Geometry and Stochastic Processes.

A trajectory is a time-stamped sequence of locations, representing the movement of entities in space and time. Trajectories are often created by sampling GPS locations, but they can also originate from RFID tags, video, or radar analysis. Time-series of locations can also be associated with other co-temporal data, such as pressure recordings for avian or aquatic animals, activity sensors and accelerometers to measure energy expenditure, or the myriad time-stamped data recorded by modern smartphones alongside GPS locations.

The study of movement involves development of concepts and methods to transform movement observations (trajectory data) to knowledge of the behavior of moving phenomena under known conditions. This knowledge is then used to calibrate simulation models to predict movement and behavioral responses in varying environmental conditions. Figure 6.10 illustrates a continuum encapsulating fundamental areas of movement research for (1) understanding movement processes through trajectory representation and computational movement analysis (the right side of Figure 6.10); and (2) modeling behavior of moving phenomena and prediction of their responses to environmental changes though modeling and simulation approaches (the left side of Figure 6.10). These two processes are tightly connected and feed into each other, often through a validation procedure on the basis of real trajectory observations.

During recent years computational movement analysis tools for trajectory data have been developed within the areas of GIScience and algorithms. Analysis objectives include clustering, similarity analysis, trajectory segmentation into characteristic subtrajectories, finding movement patterns like flocking, and relating patterns to context, and several others. Since these computations are mostly spatial, algorithmic solutions have been developed in the areas of computational geometry and GIScience. The basic analysis tasks for trajectory data are by now comparatively well understood and efficient algorithms have been developed to perform computational movement analysis. However, to be truly effective and to have real-world impact, trajectory analysis has to move beyond 'understanding movement' and tackle substantially more involved questions in 'modeling, simulation, and prediction' of movement responses to a changing environment or as results of (social) interactions.

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 basic analysis tasks for trajectory data are by now comparatively well understood and efficient algorithms have been developed to perform them. However, to be truly effective and have real-world impact, trajectory analysis has to move beyond these basic tasks and tackle substantially more involved questions, prime examples being (social) interaction and collective motion.

Research Approach and Questions Addressed

Trajectories are mathematical objects with geostatistical properties. Movement is a process that occurs as a response to the state of a moving entity across multiple spatial and temporal scales. The state and resulting behavior of moving entities determine the characteristics and capacities of movement (e.g., speeds, directions, accelerations, path sinuosity), which are highly influenced by interaction with environment, geographic context, and other moving entities. Internal properties of the moving agent such as its propensity to explore, or its power and size, typically distinguish the trajectory from that of other agents. As such no element of a trajectory can truly be independent of its other parts. Therefore, we take a view of trajectory analysis that emphasizes the treatment of the whole trajectory as a unit, rather than a series of moment-by-moment steps.

Trajectories are generated by some underlying process, which is typically assumed to integrate both stochastic elements (such as Brownian motion) and more deterministic interactions between the moving agent and the external world. Many research questions can be posed about such processes, but in this seminar we will focused primarily on identifying the forms of interaction, both with other moving agents and with environmental stimuli, and on predicting the characteristics of future trajectories.

In the seminar, we explored the following questions:

- To what extent movement observations convey information on the underlying *behavior* of moving phenomena?
- How susceptible are behaviors of moving agents to environmental changes?
- To what extent changes in the behavior of moving phenomena indicate changes in the environment?
- What does it mean to *predict* a trajectory? Should we focus on predicting the spatial locations or the geometric properties?
- How can we assess a predictive model?
- How can computational geometry help movement prediction?
- What characteristics of motion are indicative of specific trajectory generating processes, and how can we compute these efficiently?
- What is the role of time in trajectory analysis? Where can we analyze the shapes of paths independently of the time stamps and where are these vital to understanding the underlying mechanisms?
- Can we build a classification of trajectory generating mechanisms and associated trajectory properties, such as navigation by waypoints, explorative foraging
- What is the *home range*? Can we have a concrete definition for home range or activity space?
- What is a *collective*?
- Can we make algorithms that work across scales?
- To what extent *goal oriented movement* can be inferred from *local movement patterns*?

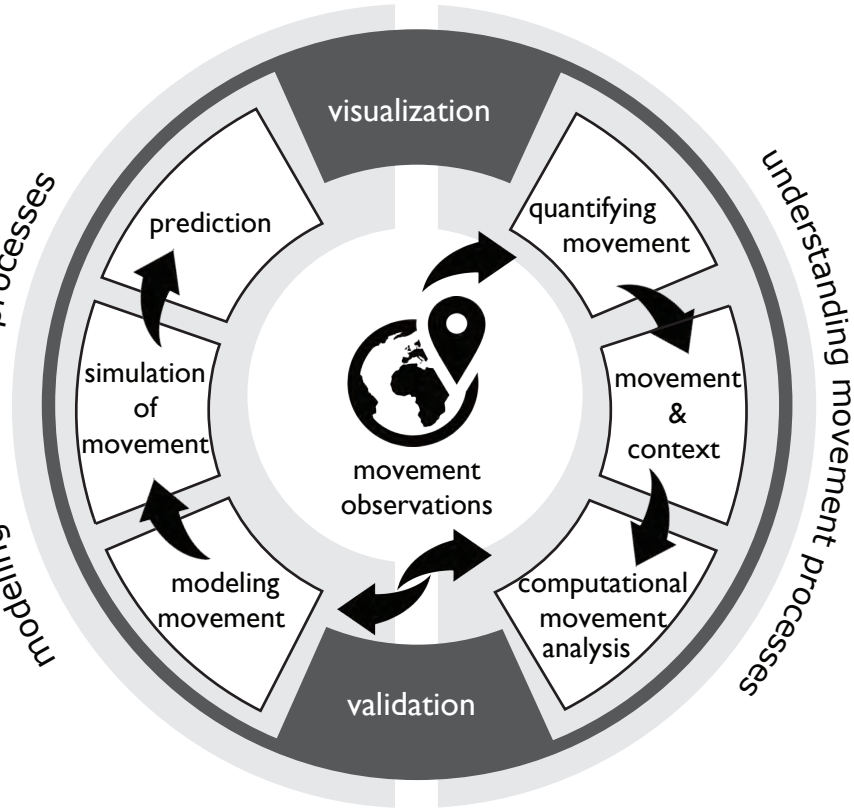

Fig. 6.10
Movement research continuum.

6

6.45 Resource Bound Analysis

Organizers: Marco Gaboardi, Jan Hoffmann, Reinhard Wilhelm, and Florian Zuleger

Seminar No. 17291

Date: July 16–21, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.72

Participants: Sebastian Altmeyer, Robert Atkey, Martin Avanzini, Gilles Barthe, Marc Brockschmidt, Hugues Cassé, Stephen Chong, Ezgi Cicek, Ugo Dal Lago, Norman Danner, Ankush Das, Marco Gaboardi, Deepak Garg, Samir Genaim, Dan R. Ghica, Jürgen Giesl, Reiner Hähnle, Sebastian Hahn, Kevin Hammond, Manuel Hermenegildo, Jan Hoffmann, Martin Hofmann, Steffen Jost, Zachary Kincaid, Jens Knoop, Björn Lisper, Yu David Liu, Alexey Loginov, Hans-Wolfgang Loidl, Antonio Flores Montoya, Georg Moser, Van Chan Ngo, Jan Reineke, Thomas W. Reps, Christine Rochange, Claudio Sacerdoti Coen, Marko van Eekelen, Pedro B. Vasconcelos, Marcus Völp, Peter Wägemann, Reinhard Wilhelm, Lukasz Ziarek, Florian Zuleger

This seminar is dedicated to our friend and colleague Martin Hofmann (1965-2018). Martin's vision and ideas have shaped our community and the way resource analysis is performed and thought about. We are grateful for the time we spent with him and we will sorely miss his ingenuity, kindness, and enthusiasm.

There are great research opportunities in combining the three aforementioned approaches to resource bound analysis. The goal of the Dagstuhl seminar was to bring together leading researchers with different backgrounds in these three areas to address challenging open problems and to facilitate communication across research areas.

To this end, the program included seven tutorials on state-of-the-art techniques in the different communities, and short talks on concrete topics with potential for cross–fertilization. This included combining WCET analysis with higher-level bound analysis techniques, hardware-specific refinement of high-level cost models, and interaction of resource analysis with compilation. Additionally, the seminar included two tools sessions: the first was a presentation of the aiT tool of AbsInt by Simon Wegener; the second was a session with presentations of different tools from different participants. Finally, the seminar included a discussion on open problems in the different areas as well as open problems for cross-fertilization.

The tutorials, the talks solicited from the participants, and the tool and discussion sessions allowed us to identify topics which are of common interest to the three different communities. Some of these topics are

- invariant and flow analysis,
- constraint solving and
- formalisms and logics for resource bounds.

Supporting information about program invariants and the possible control flow are often required by a resource analysis, e.g., the maximal value of a loop counter, or the infeasibility of a program path. The actual resource analysis is often reduced to solving a constraint system, e.g., using techniques from linear programming or recurrence equations. Verification logics for resource bounds as well as programming language formalisms are of common interest as they allow to specify or to guarantee that a program satisfies a required worst case resource bound.

We believe that the further study of these topics promises to increase the connections and to leverage the synergies between the different communities.

6.46 Topology, Computation and Data Analysis

Organizers: Hamish Carr, Michael Kerber, and Bei Wang

Seminar No. 17292

Date: July 16–21, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.88

Participants: Ulrich Bauer, Magnus Botnan, Peer-Timo Bremer, Roxana Bujack, Hamish Carr, Leila De Floriani, Pawel Dlotko, Harish Doraiswamy, Brittany Terese Fasy, Christoph Garth, Ellen Gasparovic, Hans Hagen, Ingrid Hotz, Michael Kerber, Claudia Landi, Heike Leitte, Michael Lesnick, Elizabeth Munch, Vijay Natarajan, Steve Y. Oudot, Valerio Pascucci, Jan Reininghaus, Osamu Saeki, Gerik Scheuermann, Martina Scolamiero, Donald Sheehy, Primoz Skraba, Julien Tierny, Katharine Turner, Bei Wang, Yusu Wang

The Dagstuhl Seminar titled *Topology, Computation and Data Analysis* has brought together researchers with mathematical and computational backgrounds in addressing emerging directions within computational topology for data analysis in practice. The seminar has contributed to the convergence between mathematical and computational thinking, in the development of mathematically rigorous theories and data-driven scalable algorithms.

Context

In the last two decades, considerable effort has been made in a number of research communities into computational applications of topology. Inherently, topology abstracts functions and graphs into simpler forms, and this has an obvious attraction for data analysis. This attraction is redoubled in the era of extreme data, in which humans increasingly rely on tools that extract mathematically well-founded abstractions that the human can examine and reason about. In effect, topology is applied as a form of data compression or reduction: topology is one of the most powerful forms of mathematical compression that we know how to apply to data.

Efforts to apply topology computationally to data, however, have largely been fragmented so far, with work progressing in a number of communities, principally computational topology, topological data analysis, and topological visualization. Of these, computational topology expands from computational geometry and algebraic topology to seek algorithmic approaches to topological problems, while topological data analysis and topological visualization seeks to apply topology to data analysis, of graphs and networks in the first case and of (usually) simulated volumetric data in the second. The research in these communities can roughly be clustered into theory (what are the underlying mathematical concepts), applications (how are they used for data analysis), and computation (how to compute abstractions for real datasets). It is crucial to advances in this area that these three branches go hand-in-hand, and communication between theoretical, applied, and computational researchers are therefore indispensable. On the other hand, there has been surprisingly little communication between the computational topology and topological visualization communities, mostly caused by the fact that each community has its own set of regular venues. As a consequence, the linkages in the two communities have been independent of each other, and results can take years to migrate from one community to the other.

Vision

Our goal was therefore to soften the aforementioned rather strict separation between computational topology and topological visualization by establishing new inter-community ties. The seminar aimed to bring together cross sections of both communities, including researchers with theoretical, applied, and computational backgrounds. By reducing redundancy and accelerating cross-communication, we expected a significant boost to both areas, perhaps even leading to a singular more dynamic community. As a side effect, we also wanted to provide a communication platform within each community between theory and application.

Topics

We identified specific research topics reflecting emerging trends in both communities. These topics were chosen to span the spectrum from the theoretical (category theory), to applicable theory (multidimensional persistent homology), and from applied theory (singularity theory and fiber topology) to the computational (scalable topological computation, applications) aspect.

Category theory: theory and applications. Category theory

has recently gained momentum in computational topology, in particular through sheaves and cosheaves, which are extremely useful as an alternative foundation for level set persistence. Recent work has shown that the data of a Reeb graph can be stored in a category-theoretic object called a cosheaf, and this opens the way to define a metric for Reeb graphs known as the interleaving distance. Sheaves can also be used in deriving theoretical understandings between the Reeb space and its discrete approximations. Research into sheaves and their relationship with computation is, however, in its infancy, and would benefit from pooling the resources of experts in category theory and topological data analysis, to address questions such as how to simplify theories in computational topology, how to reinterpret persistent homology, or how to compare topological structures.

Multidimensional persistent homology. The second area of active research, both mathematically and computationally, is the extension of unidimensional persistence to multidimensional persistence. Mathematically, the lack of a complete discrete invariant for the multidimensional case raises the theoretical question of identifying meaningful topological invariants to compute. Some earlier proposals have been complemented by recent approaches and raise the immediate question of computability and applicability. Besides the invariants themselves, other questions such as the comparison of multidimensional data, or the efficient generation of cell complexes suitable for the multidimensional case are crucial, but hardly studied questions in this context. Computationally, existing algorithms for topological constructs rely on filtrations to encapsulate a sweep order through the data, thus serializing the problem for algorithmic implementation. For multidimensional data, this serialization is hard to achieve, and progress in this area is, therefore, crucial for computational advances in the topological analysis of data.

Singularity theory and fiber topology in multivariate data analysis. Singularity theory and fiber topology both seek to extend Morse theory from scalar fields to multivariate data described as functions mapping $f : \mathbb{X} \to \mathbb{R}^d$. Since multivariate datasets are near-ubiquitous in scientific applications such as oceanography, astrophysics chemistry, meteorology, nuclear engineering and molecular dynamics, advances here are also crucial for topological data analysis and visualization. Methods from computational topology have been developed to support the analysis of scalar field data with widespread applicability. However, very few tools exist for studying multivariate data topologically: the most notable examples of these tools are the Jacobi set, the Reeb space, and its recent computational approximation, the Joint Contour Net. Here, we aim to bring together researchers in singularity theory, fiber topology and topological data analysis to develop new theory and algorithms driving a new generation of analytic tools.

Scalable computation. At the opposite pole from theory is the practical question: how do we apply topological analysis to ever-larger data sets? This question spans questions of algorithmic performance to the accuracy of representation: using the metaphor of compression, do we want lossy or lossless compression, how fast can we perform it, and what do we lose in the process? Moreover, the largest data sets are necessarily computed and stored on clusters, and scalability of topological computation therefore also depends on building distributed and parallel algorithms. For example, the standard algorithm for computing persistent homology is cubic in the number of simplices, but can be speeded up in theory and practice, and further improved by parallel computation. However, many challenges remain, including efficient generation, storage and management of simplicial complexes, streaming computation, I/O efficient computation, approximate computation, and non-simplicial complexes. Some of these approaches have already been applied in topological visualization, and cross-fertilization between the two communities is therefore of great interest.

Participants, Schedule, and Organization

The invitees were chosen according to the topics, bringing together enough expertise for each topic and resulting in a representative subset of both communities. Out of the 37 invited researchers in the first round, 28 accepted our invitation, pointing out the general interest for the seminar topic in both communities.

We decided for a mixed setup with introductory talks, contributed research talks and breakout sessions.

For the first day, we scheduled two overview talks per listed topic, which were delivered by Steve Oudot and Elizabeth Munch (Category theory), Michael Lesnick and Claudia Landi (Multidimensional persistent homology), Osamu Saeki and Julien Tierny (Singularity theory), and Yusu Wang and Valerio Pascucci (Scalable Computation). Further contributed talks by participants took place from Tuesday to Friday morning, resulting in a total of 19 contributed talks.

The afternoons of Tuesday and Thursday were used for breakout sessions. The format was different on the two days. Based on the discussions on Monday, we identified the topics "multivariate topology" and "scalable computation" as topics of general interest. We decided to let every participant discuss both topics, so we organized 4 discussion groups on multivariate topology in the early afternoon, and 3 discussion groups on scalable computation in the later afternoon (plus an alternative group with a different topic). We composed these groups mostly randomly, making sure that members of both communities are roughly balanced in each group. On Thursday afternoon, we let participants propose their topics of interest. 5 groups were formed discussing various aspects raised in contributed talks. On Wednesday and Friday morning, the outcomes of every discussion group were summarized and discussed in a plenary session.

Moreover, the majority of the participants joined an organized excursion to Trier on Wednesday afternoon.

Results and Reflection

The participants gave the unanimous feedback that the breakout sessions were a full success (and several proposed more time for such discussions in possible upcoming seminars). We first let people from a mixed background to discuss rather vague topics on Tuesday, and asked for specific topics on Thursday. Such an organizational plan led to a stimulating working environment, and helped to avoid idle breakout sessions.

We believe that we have fully achieved the goal of softening the separation between the two communities involved in this seminar. We expect visible evidence of newly formed inter-community ties fostered by the seminar, for instance through joint research projects and/or survey articles summarizing major open problems on the interface of both communities. To the best of our knowledge, 3 working groups are being formed and at least 1 position paper is underway that will combine expertise from both communities to tackle key research questions raised during the seminar.

6.47 User-Generated Content in Social Media

Organizers: Tat-Seng Chua, Norbert Fuhr, Gregory Grefenstette, Kalervo Järvelin, and Jaakko Peltonen

Seminar No. 17301

Date: July 23–28, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.7.110

Participants: Tat-Seng Chua, Nicolas Diaz-Ferreyra, Gerald Friedland, Norbert Fuhr, Anastasia Giachanou, Tatjana Gornostaja, Gregory Grefenstette, Iryna Gurevych, Andreas Hanselowski, Xiangnan He, Benoit Huet, Kalervo Järvelin, Rosie Jones, Rianne Kaptein, Krister Lindén, Yiqun Liu, Marie-Francine Moens, Josiane Mothe, Wolfgang Nejdl, Jaakko Peltonen, Isabella Peters, Miriam Redi, Stevan Rudinac, Markus Schedl, David Ayman Shamma, Alan Smeaton, Benno Stein, Lexing Xie

Social media play a central role in many people's lives, and they also have a profound impact on businesses and society. Users post vast amounts of content (text, photos, audio, video) every minute. This user generated content (UGC) has become increasingly multimedia in nature. It documents users' lives, revealing in real time their interests and concerns and activities in society. The analysis of UGC can offer insights to individual and societal concerns and could be beneficial to a wide range of applications, for example, tracking mobility in cities, identifying citizen's issues, opinion mining, and much more. In contrast to classical media, social media thrive by allowing anyone to publish content with few constraints and no oversight. Social media posts thus show great variation in length, content, quality, language, speech and other aspects. This heterogeneity poses new challenges for standard content access and analysis methods. On the other hand, UGC is often related to other public information (e.g. product reviews or discussion of news articles), and there often is rich contextual information linking, which allows for new types of analyses.

In this seminar, we aimed at discussing the specific properties of UGC, the general research tasks currently operating on this type of content, identifying their limitations and lacunae, and imagining new types of applications made possible by the availability of vast amounts of UGC. This type of content has specific properties such as presentation quality and style, bias and subjectivity of content, credibility of sources, contradictory statements, and heterogeneity of language and media. Current applications exploiting UGC include sentiment analysis, noise removal, indexing and retrieving UGC, recommendation and selection methods, summarization methods, credibility and reliability estimation, topic detection and tracking, topic development analysis and prediction, community detection, modeling of content and user interest trends, collaborative content creation, cross media and cross lingual analysis, multi-source and multi-task analysis, social media sites, live and real-time analysis of streaming data, and machine learning for big data analytics of UGC. These applications and methods involve contributions from several data analysis and machine learning research directions.

This seminar brought together researchers from different subfields of computer science, such as information retrieval, multimedia, natural language processing, machine learning and social media analytics. After participants gave presentations of their current research orientations concerning UGC, we decided to split into two Working Groups: (WG-1) Fake News and Credibility, and (WG-2) Summarizing and Storytelling from UGC.

WG-1: Fake News and Credibility

WG-1 began discussing the concept of Fake News, and we arrived at the conclusion that it was a topic with much nuance, and that a hard and fast definition of what was fake and what was real news would be hard to define. We then concentrated on deciding what elements of Fake (or Real) News could be calculated or quantified by computer. This led us to construct a list of text quality measures that have or are being studied in the Natural Language Processing community: Factuality, Reading Level, Virality, Emotion, Opinion, Controversy, Authority, Technicality, and Topicality. During this discussion, WG-1 invented and mocked up what we called an Information Nutrition Label, modeled after nutritional labels found on most food products nowadays. We feel that it would be possible to produce some indication of the "objective" value of a text using the above nine measures. The user could use these measures to judge for themselves whether a given text was "fake" or "real". For example, a text highly charged in Emotion, Opinion, Controversy, and Topicality might be Fake News for a given reader. Just like with a food nutritional label,

a reader might use the Information Nutritional Label to judge whether a given news story was "healthy" or not.

WG-1 split into further subgroups to explore whether current status of research in the nine areas: Factuality, Reading Level, etc. For each topic, the subgroups sketched out the NLP task involved, found current packages, testbeds and datasets for the task, and provided recent bibliography for the topic. Re-uniting in one larger group, each subgroup reported on their findings, and we discussed next steps, envisaging the following options: a patent covering the idea, creating a startup that would implement all nine measures and produce a time-sensitive Information Nutritional Label for any text submitted to it, a hackathon that would ask programmers to create packages for any or all of the measurements, a further workshop around the Information Nutrition label, integration of the INL into teaching of Journalists, producing a joint article describing the idea. We opted for the final idea, and we produced a submission (also attached to this report) for the Winter issue of the SIGIR (Special Interest Group on Information Retrieval) Forum[50].

WG-2: Summarizing and Storytelling from UGC

WG-2 set out to re-examine the topic of summarization. Although this is an old topic, but in the era of user-generated content with accelerated rates of information creation and dissemination, there is a strong need to re-examine this topic from the new perspectives of timeliness, huge volume, multiple sources and multimodality. The temporal nature of this problem also brings it to the realm of storytelling, which is done separately from that of summarization. We thus need to move away from the traditional single source document-based summarization, by integrating summarization and storytelling, and refocusing the problem space to meet the new challenges.

We first split the group into two sub-groups, to discuss separately: (a) the motivations and scopes, and (b) the framework of summarization. The first sub-group discussed the sources of information for summarization including, the user-generated content, various authoritative information sources such as the news and Wikipedia, the sensor data, open data and proprietary data. The data is multilingual and multimodal, and often in real time. The group then discussed storytelling as a form of dynamic summarization. The second group examined the framework for summarization. It identified the key pipeline processes comprising of: data ingestion, extraction, reification, knowledge representation, followed by story generation. In particular, the group discussed the roles of time and location in data, knowledge and story representation.

Finally, the group identified key challenges and applications of the summarization framework. The key challenges include multi-source data fusion, multilinguality and multimodality, the handling of time/ temporality/ history, data quality assessment and explainability, knowledge update and renewal, as well as focused story/ summary generation. The applications that can be used to focus the research includes event detection, business intelligence, entertainments and wellness. The discussions have been summarized into a paper entitled "Rethinking Summarization and Storytelling for Modern Social Multimedia". The paper is attached along with this report. It has been submitted to a conference for publication.

[50] http://sigir.org/forum/

6.48 Scalable Set Visualizations

Organizers: Yifan Hu, Luana Micallef, Martin Nöllenburg, and Peter Rodgers

Seminar No. 17332

Date: August 13–18, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.8.1

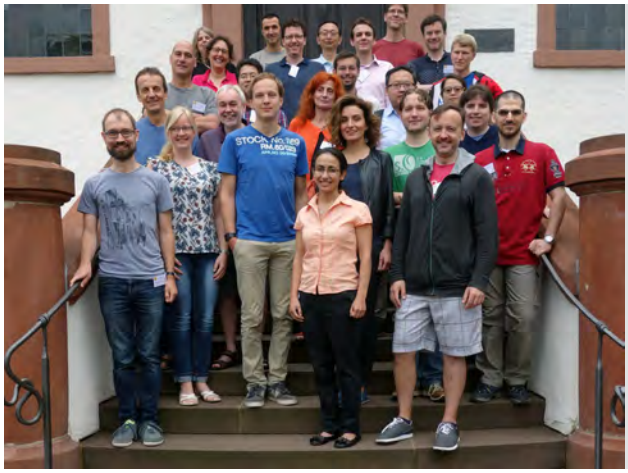

Participants: Daniel Archambault, Robert Baker, Kerstin Bunte, Nan Cao, Thom Castermans, Fadi Dib, Sara Irina Fabrikant, John Howse, Yifan Hu, Radu Jianu, Michael Kaufmann, Andreas Kerren, Stephen G. Kobourov, Martin Krzywinski, Tamara Mchedlidze, Wouter Meulemans, Luana Micallef, Silvia Miksch, Martin Nöllenburg, Sergey Pupyrev, Peter Rodgers, Mereke van Garderen, Alexander Wolff, Hsiang-Yun Wu, Michael Wybrow, Xiaoru Yuan

Sets are a fundamental way of organizing information. Visualizing set-based data is crucial in gaining understanding of it as the human perceptual system is an analytic system of enormous power. The number of different set visualization methods has increased rapidly in recent years, and they vary widely in the visual metaphors used and set-related tasks that they support. Large volume set-based data can now be found in diverse application areas such as social networks, biosciences and security analysis. At present, set visualization methods lack the facility to provide analysts with the visual tools needed to successfully interpret large-scale set data as the scalability of existing metaphors and methods is limited. This seminar provided a forum for set visualization researchers and application users to discuss how this challenge could be addressed.

Existing set visualizations can be grouped into several families of techniques, including traditional Euler and Venn diagrams, but also node-link diagrams, map- and overlay-based representations, or matrix-based visualizations. Inevitably, the approaches taken to drawing these visualizations are diverse, for example node-link diagrams require graph drawing methods, whereas overlay techniques use algorithms from computational geometry. However, they are similar in a number of aspects. One aspect is the underlying set theory. For instance, theoretical results into the drawability of many of these set visualization techniques for different data characteristics is possible (as already done in example Venn and Euler diagram research). Another common aspect is that the visualizations are typically focused on an end-user, so perceptual, cognitive and evaluation considerations are an important concern.

A particularly pressing issue in set visualization is that of scaling representations. The number of data items can be large and many methods aggregate individual items. Yet, even using aggregation, the limit of the most scalable of these methods is considered to be in the region of 100 sets [1]. Typical application areas that make use of sets include, e.g., social networks, biosciences and security analysis. In these applications, there may be many millions of data items in thousands of sets. Other applications have high-dimensional data, where each item is associated with a large number of variables, which poses different scalability challenges for set visualizations.

A distinct feature of set visualization is that visualizations must support set-related, element-related, and attribute-related analysis tasks [1] that involve, e.g., visually evaluating containment relations, cardinalities, unions, intersections, or set differences. For example, bioscience microarray experiments classify large numbers of genes and multiple visualization tools have been developed to visualize this data. However, current efforts can only visualize small sections of the information at once [2]. Similar scalability challenges for set visualizations appear in many other applications as well. Hence, developing effective visualization methods for large set-based data would greatly facilitate analysis of such data in a number of important application areas.

Seminar Goals

The goal of this seminar was to bring together researchers with different backgrounds but a shared interest in set visualization. It involved computer scientists with expertise, e.g., in visualization, algorithms, and human-computer interaction, but also users of set visualizations from domains outside computer science. Despite the large number of set visualization techniques, for which there is often a considerable practical and theoretical understanding of their capabilities, there has only been limited success in scaling these methods. Thus the intended focus of this seminar was to discuss and study specific research challenges for scalable set visualizations concerning fundamental theory, algorithms, evaluation, applications, and users. We started with a

few overview talks on the state of the art in set visualization, but then focused on small hands-on working groups during most of the seminar week. We aimed to accelerate the efforts to improve scalability of set visualizations by addressing open questions proposed by the seminar attendees, in order to produce concrete research outcomes, including new set visualization software and peer-reviewed research publications.

Seminar Program

1. On the first two days of the seminar we enjoyed five invited overview lectures on different aspects of set visualizations. The topics and speakers were chosen as to create a joint understanding of the state of the art of set visualization techniques, evaluations and applications. Silvia Miksch gave a systematic overview of set visualization techniques, grouped by types of visual representations and tasks, with a special focus on set visual analytics. Martin Krzywinski reported about his experiences on using visual analogies for showing set-based data in the area of genomics. Sara Fabrikant took a cartographer's view on visualizing sets and explained how successful cartographic maps work as information displays by taking not only the design but also the context and the user into account. Stephen Kobourov explained how large graph-based set data can be represented using a familiar map metaphor by showing several interesting data sets and their map representations. Finally, John Howse presented how set visualizations can be used as diagrammatic reasoning systems in logic.
2. In the open problem session on the first day of the seminar we collected a list of 13 open research problems that were contributed by the seminar participants. In a preference voting we determined the five topics that raised the most interest among the participants and formed small working groups around them. During the following days the groups worked by themselves, except for a few plenary reporting sessions, formalizing and solving their respective theoretical and practical challenges. Below is a list of the working group topics; more detailed group reports are found in Section 4 of the full report.
   a. **Mapifying the genome:** Can the axis of the entire genome be mapped on a 2-dimensional space based on gene function rather than a 1-dimensional line based on gene position?
   b. **Area-proportional Euler diagrams with ellipses:** Can the use of ellipses extend the size of data that can be drawn with area-proportional Euler diagrams?
   c. **Spatially informative set visualizations:** Can we improve spatial overlay-based set visualizations when allowing some limited displacement of the given set positions?
   d. **Set visualization using the metro map metaphor:** How and under which conditions can the metro map metaphor be used to visualize set systems?
   e. **Visual analytics of sets/set-typed data and time: challenges and opportunities:** What are the main research challenges and opportunities in the context of set visualizations that change over time and how can these be structured?
3. We had a flexible working schedule with a short plenary session every morning to accommodate group reports and impromptu presentations by participants. In two of those Wouter Meulemans and Nan Cao shared recent results of theirs related to set visualization.
4. During the week we encouraged participants to come up with suggestions for further strengthening this growing community of set visualization researchers. In a plenary session on Friday we collected and structured these ideas and made started planning future events related to set visualizations, see Section 6.48.

Future Plans

During the entire seminar, participants actively discussed ways how to disseminate, proliferate and promote scalable set visualization research in diverse specific areas, such as: set theory and diagrammatic reasoning; algorithms and graph theory; information visualization and visual analytics; evaluation, users and application areas. This led to the concretization of the following future milestones, each of which is being coordinated by volunteered seminar participants:

- **Diagrams Workshop in 2018** on Set Visualization and Reasoning (SetVR)
  The workshop aims at promoting set visualization to the Diagrams community, of which well-renowned mathematicians and logicians are members, thus proliferating relevant set theory and diagrammatic reasoning research;
- **IEEE VIS Workshop in 2019** on Set Visualization and Analytics (SetVA)
  The workshop aims at promoting set visualization to the Information Visuaization and Visual Analytics communities, at the premier forum for advances in information and scientific visualization, with the aim to generate new visualization and analytic techniques to handle large set-typed data;
- **Dagstuhl seminar in 2019** on Set Visualization and Analytics (SetVA) over Time and Space
  This seminar has revealed, for the first time, the need for visualization and analytic techniques for the set-typed data that has an element of time and/or space; thus a follow-up Dagstuhl seminar will be organized to discuss this topic, once again among researchers with diverse set visualization backgrounds;
- **Set Visualization Workshop in 2020** in the Computational Geometry Week or collocated with Graph Drawing
  This workshop aims at disseminating set visualization to a more algorithmic and computational geometry research community, to ensure the production of effective, yet efficient and scalable set visualization algorithms;
- **Set Visualization browser**, like http://setviz.net
  This browser will collect and disseminate available set visualization techniques, making them easily accessible through various categorizations, such as the type of data analysis tasks or application areas they target;
- **Set Visualization book**
  The book would serve as a guide for researchers who are new to set visualization and as a review of the current state-of-the-art of set-typed data in the various related domains.

We decided to have an annual set visualization workshop that each year focuses on one of (i) diagrammatic reasoning and logic, (ii) information visualization and visual analytics, and (iii) computational geometry and graph drawing, at premier venues of the respective research communities, to generate further research interest in all of these three diverse areas that are all important for scalable set visualizations.

Evaluation

According to the Dagstuhl survey conducted after the seminar, as well as informal feedback to the organizers, the seminar was highly appreciated. Particularly the small group size, group composition, and the seminar structure focusing on hands-on working groups was very well received. The seminar's goals to identify and initiate collaboration on new research challenges was very successful (also in comparison to other Dagstuhl seminars) as the participants rated the seminar highly for inspiring new research directions, joint projects and joint publications. We are looking forward to seeing the first scientific outcomes of the seminar in the near future and to continuing the efforts to support the growth of the set visualization community.

Acknowledgments

Schloss Dagstuhl was the perfect place for hosting a seminar like this. The unique scientific atmosphere and the historic building provided not only all the room we needed for our program and the working groups, but also plenty of opportunities for continued discussions and socializing outside the official program. On behalf of all participants the organizers want to express their deep gratitude to the entire Dagstuhl staff for their outstanding support and service accompanying this seminar. We further thank Tamara Mchedlidze for helping us collecting the contributions and preparing this report.

References

1 Bilal Alsallakh, Luana Micallef, Wolfgang Aigner, Helwig Hauser, Silvia Miksch, and Peter Rodgers. The State-of-the-Art of Set Visualization. *Computer Graphics Forum*, 35(1):234–260, 2015.

2 Sebastian Behrens and Hans A Kestler. Using VennMaster to evaluate and analyse shRNA data. *Ulmer Informatik-Berichte*, page 8, 2013.

6.49 Computational Counting

Organizers: Ivona Bezáková, Leslie Ann Goldberg, and Mark R. Jerrum

Seminar No. 17341

Date: August 20–25, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.8.23

Participants: Miriam Backens, Eleni Bakali, Alexander Barvinok, Ivona Bezáková, Markus Bläser, Antonio Blanca, Cornelius Brand, Andrei A. Bulatov, Sarah Cannon, Amin Coja-Oghlan, Radu Curticapean, Holger Dell, Martin Dyer, Charis Efthymiou, Jacob Focke, Andreas Galanis, David Gamarnik, Andreas Göbel, Leslie Ann Goldberg, Catherine Greenhill, Heng Guo, Thomas Hayes, Miki Hermann, Thore Husfeldt, Mark R. Jerrum, John Lapinskas, Jingcheng Liu, Kitty Meeks, Sarah Miracle, Viresh Patel, Will Perkins, Dana Randall, Guus Regts, David Richerby, Marc Roth, Piyush Srivastava, Alexandre Stauffer, Daniel Stefankovic, Mingji Xia, Kuan Yang, Yitong Yin, Chihao Zhang, Stanislav Zivny

Computational counting problems arise in practical applications in many fields such as statistical physics, information theory and machine learning. In such a problem, the goal is to compute or to estimate a weighted sum. Some typical computational counting problems include evaluating a probability, the expectation of a random variable, a partition function, or an integral.

The study of the computational complexity of counting problems requires a coherent set of techniques which are different in flavour from those employed in other algorithmic branches of computer science. Relevant techniques include the analysis of Markov chains, the analysis of correlation decay, parameterised algorithms and complexity, and dichotomy techniques for constructing detailed classifications.

Most computational problems are intractable when considered from the perspective of classical complexity, so it is important to find ways to cope with intractability. These include approximation, randomisation, as well as viewing computational counting through the lens of parameterised complexity, where the goal is to find algorithms that are efficient when some key parameter is "small". Intractability thresholds often relate to "phase transitions" as these key parameters vary. Great progress has been made in recent years towards understanding the complexity of approximate counting, based largely on a connection with these phase transitions.

Specific themes identified for consideration at the meeting included:

- *Exact counting*, including classifications, quasi-polynomial and/or moderately exponential algorithms for intractable problems, and parameterised algorithms; also complexity-theoretic limitations to obtaining exact solutions.
- *Approximate counting* including Markov Chain Monte Carlo (*MCMC*) algorithms, and algorithms based on decay of correlations; also complexity-theoretic limitations to obtaining approximate solutions.
- The interplay between *phase transitions* and computational tractability.
- *Constraint satisfaction* problems and the more general *Holant* framework. The partition functions of many models in statistical physics are included within this setting.

In the event, the talks ranged more widely than this list suggests.

Although the topic of Computational Counting has been explored at various meetings, including at Dagstuhl, for a number of years, it continues to retain its freshness. New approaches are found, new insights are gained, and new connections drawn with other areas both inside and outside computer science. Among the new directions that have emerged since the previous Dagstuhl Seminar in this series are the following.

- Results from quantum information theory applied to the apparently unrelated task of classifying the complexity of Holant problems. (Refer to the talk by Miriam Backens.)
- A new paradigm for designing polynomial-time algorithms for approximating partition functions with complex parameters. This is based on Taylor expansion in a zero-free region of the parameter space combined with an ingenious approach to enumerating small substructures. (Refer to talks by Alexander Barvinok, Jingcheng Liu, Viresh Patel and Guus Regts.)
- Emerging connections between the Lovász Local Lemma – specifically the Shearer condition and the Moser-Tardos algorithmic version – and sampling and approximate counting. (Refer to talks by Andreas Galanis and Heng Guo.)

6.50 SLEBOK: The Software Language Engineering Body of Knowledge

Organizers: Benoît Combemale, Ralf Lämmel, and Eric Van Wyk

Seminar No. 17342

Date: August 20–25, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.8.45

Participants: Mathieu Acher, Anya Helene Bagge, Walter Cazzola, Andrei Chis, Benoit Combemale, Thomas Degueule, Sebastian Erdweg, Johannes Härtel, Görel Hedin, Marcel Heinz, Ralf Lämmel, Manuel Leduc, Tanja Mayerhofer, Peter D. Mosses, Gunter Mussbacher, Oscar M. Nierstrasz, Anthony Sloane, Friedrich Steimann, Eugene Syriani, Tijs van der Storm, Eric Van Wyk, Hans Vangheluwe, Jurgen J. Vinju, Markus Völter, Vadim Zaytsev

Over the last 10 years, the field of **Software Language Engineering** (SLE) has emerged based on a strong motivation to connect and integrate different research disciplines such as compiler construction, reverse engineering, software transformation, model-driven engineering, and ontologies. This seminar strives for directly promoting the further integration of said communities with the clear objective of assembling a **Body of Knowledge** on SLE (SLEBoK). The BoK features artefacts, definitions, methods, techniques, best practices, open challenges, case studies, teaching material, and other components that will afterwards help students, researchers, teachers, and practitioners to learn from, to better leverage, to better contribute to, and to better disseminate the intellectual contributions and practical tools and techniques coming from the SLE field.

The following questions and issues provided the guiding principles for the seminar. The first two categories reflect on the result of the past decade and the last category looks forward to the next decade; these categories has been addressed by the seminar attendees in breakout groups.

- **Conceptual model of the SLE field**: What is a comprehensive and objective (validated) classification of SLE approaches? What are appropriate dimensions of such a classification? How to otherwise ontologically organize software language engineering, e.g., in terms of application areas, relationships to other software engineering areas, and fundamental SLE concepts?
- **Curriculum contributions by the SLE field**: What is the suite of formal notions and engineering methods, that one could want to see introduced in a computer science curriculum so that SLE is properly represented? What is a reference curriculum for SLE? What is an appropriate combination of timeless foundations and relevant (current) applications and technologies? How to contribute to or otherwise support a computer science curriculum?
- **Open SLE challenges**: What are the open challenges in SLE and how to lay out a larger research agenda that the community can refer to in the next 10 years? How to connect to important developments such as AI and IoT? How to measure the relevance of the research priorities proposed?

With the SLE field approximately 10 years old, there is a strong support by the community to analyse the situation and to move to the next level of maturity. This Dagstuhl seminar provided the ideal format for such a critical analysis and further development of SLE's foundation. As a result of the work on the above three pillars, the seminar attendees initiated the SLEBOK: **https://github.com/slebok/slebok**.

Fig. 6.11
"Even bugs at #SLEBoK #Dagstuhl are beautiful!" Twitter post by 17342 Dagstuhl Seminar participant Vadim Zaytsev. https://twitter.com/grammarware/status/901369729440518146. Photo courtesy of Vadim Zaytsev / CC BY 3.0 DE.

6.51 Machine Learning and Formal Methods

Organizers: Sanjit A. Seshia, Xiaojin (Jerry) Zhu, Andreas Krause and Susmit Jha

Seminar No. 17351

Date: August 28–September 1, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.8.55

Participants: S. Akshay, Dalal Alrajeh, Rajeev Alur, Nikolay A. Atanasov, Ammar Ben Khadra, Felix Berkenkamp, Ben Caulfield, Swarat Chaudhuri, Luc De Raedt, Jyotirmoy Deshmukh, Dana Fisman, Pranav Garg, Matthias Hein, Holger Hermanns, Susmit Jha, Kristian Kersting, Andreas Krause, Sasa Misailovic, Mayur Naik, Nagarajan Natarajan, Anna Rafferty, Dorsa Sadigh, Stephan Schulz, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Charles Sutton, Josef Urban, Martin Vechev, Eran Yahav, Yisong Yue, Xiaojin Zhu, Sandra Zilles

The seminar was successful in bringing the following two communities together:

- The community that works on machine learning (ML), both on theoretical topics and on applications to areas such as robotics and cyber-physical systems, and
- The community that works on formal methods (FM), both on computational proof techniques and on applications to formal verification and program/controller synthesis.

Both communities have long and vibrant histories, with associated conferences and journals. However, they have rarely intersected. The machine learning community has traditionally focused on *inductive* learning from data, with the data set considered as partial (potentially noisy) observations of some phenomenon. The formal methods community has traditionally emphasized automated *deduction*, e.g., using theorem proving or model checking, as a core reasoning method, with a heavy emphasis placed on formal models and proofs of correctness using those models. However, recent ideas and methods have appeared that demonstrate new connections between the two disciplines, which suggested that the time is ripe for a meeting to promote cross-fertilization between the areas at a deep technical level. This seminar has been a significant step forward to bring the two communities together.

More concretely, the Seminar and the interaction it facilitated has brought three kinds of benefits. First, formal methods can benefit from a more effective use of machine learning techniques particularly in the context of automated synthesis. Similarly, the increasing use of machine learning in applications that require a high level of assurance points to the need for integration with formal methods. However, the potential synergies between the two areas go beyond a simple application of the techniques in one area to the other area. Importantly, there is new fundamental science to be explored in the intersection of machine learning and formal methods, related to the the confluence of inductive and deductive reasoning, and which can inform a range of new industrially-relevant applications as well.

The seminar had about 40 participants from both the FM and ML communities. The organizers took several steps to foster discussion and cross-pollination of ideas between the two communities, including the following:

- The seminar began with a day of tutorials: a half-day tutorial on Machine Learning for Formal Methods participants, and a half-day tutorial on Formal Methods for a Machine Learning audience. These tutorials helped to establish a common vocabulary to discuss ideas, problems and solutions.
- Sessions were organized based on themes that emerged in discussions before the seminar and during the first day. The list of session topics is as follows:
  1. Probabilistic Programming
  2. Teaching and Oracle-Guided Synthesis
  3. Safe Learning-Based Control
  4. Probabilistic Program Analysis
  5. Adversarial Analysis and Repair for Machine Learning
  6. Inductive Synthesis and Learning
  7. Machine Learning for Theorem Proving and Optimization
  8. Explainable and Interpretable Machine Learning
  9. Deep Learning and Verification/Synthesis

  In organizing these sessions, the organizers tried to

combine speakers from both ML and FM areas to foster discussion and comparison of approaches.
- Seating arrangements at meals were organized so that (a) each table had an approximately equal number of participants from both communities, and (b) the seating was randomly changed from meal to meal.
- A joint session was organized with the concurrent seminar on analysis and synthesis of floating-point programs. This session had a panel discussion on floating-point issues in machine learning programs.

After the seminar, we have heard positive feedback from multiple participants. One told us that he started a new research project as a direct result of the seminar. A group of participants are planning to continue the interaction via joint workshops at major venues of both communities such as CAV, PLDI, ICML, NIPS, etc.

6.52 Analysis and Synthesis of Floating-point Programs

Organizers: Eva Darulova, Alastair F. Donaldson, Zvonimir Rakamarić, and Cindy Rubio-González

Seminar No. 17352

Date: August 27–30, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.8.74

Participants: Erika Abraham, George A. Constantinides, Nasrine Damouche, Eva Darulova, James W. Demmel, Anthony Di Franco, Alastair F. Donaldson, Theo Drane, Sam Elliott, Ganesh L. Gopalakrishnan, Hui Guo, Jeffrey K. Hollingsworth, Miriam Leeser, Daniel Liew, Piotr Luszczek, Victor Magron, Matthieu Martel, Guillaume Melquiond, David Monniaux, Magnus Myreen, Santosh Nagarakatte, Pavel Panchekha, Sylvie Putot, Zvonimir Rakamarić, Nathalie Revol, Cindy Rubio-González, Daniel Schemmel, Oscar Soria Dustmann, Zachary Tatlock, Michela Taufer, Laura Titolo, Thomas Wahl

This report documents the program and the outcomes of Dagstuhl Seminar 17352 "Analysis and Synthesis of Floating-point Programs".

Floating-point numbers provide a finite approximation of real numbers that attempts to strike a fine balance between range, precision, and efficiency of performing computations. Nowadays, performing floating-point computations is supported on a wide range of computing platforms, and are employed in many widely-used and important software systems, such as high-performance computing simulations, banking, stock exchange, self-driving cars, and machine learning.

However, writing correct, and yet high-performance and energy-efficient, floating-point code is challenging. For example, floating-point operations are often non-associative (contrary to their real mathematical equivalents), which creates problems when an ordering of operations is modified by either a compiler or due to nondeterministic interleavings of concurrent executions. Furthermore, the underlying floating-point hardware is often heterogeneous, hence different results may be computed across different platforms or even across components of the same heterogeneous platform. Given the underlying complexity associated with writing floating-point code, it is not surprising that there have been numerous software bugs attributed to incorrectly implemented floating-point computations.

Research related to floating-point computation spans a multitude of areas of computer science, ranging from hardware design and architecture, all the way to high-performance computing, machine learning, and software analysis and verification. The objective of this seminar was thus to bring together researchers from several of these areas, which have either traditionally been considered as non-overlapping, or which have arguably enjoyed insufficient interaction despite a clear overlap of interests. The goal in mind here was to provide opportunities to brainstorm new theoretical advances and practical techniques and tools for making floating-point computations performant and correct, and to help foster long term collaborations.

The seminar involved brief presentations from most participants, interspersed with a lot of informal technical discussion, in addition to four breakout sessions based on common themes that arose during informal discussion. In addition, a joint panel session was held between this seminar and the concurrently running "Machine Learning and Formal Methods" seminar. This report presents the collection of abstracts associated with the participant presentations, notes summarising each discussion session, and a transcript of the panel session. We hope that the report will provide a useful resource for researchers today who are interested in understanding the state-of-the-art and open problems related to analysing and synthesising floating-point programs, and as a historical resource helping to clarify the status of this field in 2017.

Fig. 6.12
"dagstuhl b". Picture on flickr by 17392 Dagstuhl Seminar participant m. c. schraefel.
https://flic.kr/p/BSxF9N. Photo courtesy of m. c. schraefel.

6.53 Finite and Algorithmic Model Theory

Organizers: Anuj Dawar, Erich Grädel, Phokion G. Kolaitis, and Thomas Schwentick

Seminar No. 17361

Date: September 3–8, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.1

Participants: Faried Abu Zaid, Albert Atserias, Christoph Berkholz, Dietmar Berwanger, Mikolaj Bojanczyk, Thomas Colcombet, Claire David, Anuj Dawar, Arnaud Durand, Kousha Etessami, Ronald Fagin, Diego Figueira, Jakub Gajarsky, Erich Grädel, Etienne Grandjean, Martin Grohe, Lauri Hella, Sandra Kiefer, Phokion G. Kolaitis, Manfred Kufleitner, Dietrich Kuske, Carsten Lutz, James F. Lynch, Jerzy Marcinkowski, Joanna Ochremiak, Martin Otto, Wied Pakusa, Svenja Schalthöfer, Thomas Schwentick, Luc Segoufin, Sebastian Siebertz, Dan Suciu, Lidia Tendera, Wolfgang Thomas, Szymon Torunczyk, Jouko Väänänen, Jan Van den Bussche, Victor Vianu, Heribert Vollmer, Nils Vortmeier, Scott Weinstein, Gregory Wilsenach, Thomas Zeume

Topic and Goals of the Seminar

Finite and Algorithmic Model Theory (FAMT) encompasses a number of research themes united around common methods for analysing the expressive power of logical formalisms on structures that are either finite or can be finitely represented. These are precisely the structures that can serve as inputs to computation and, for this reason, FAMT is intimately connected to computer science. Over the past decades the subject has developed through a close interaction between theoretical computer science and closely related areas of mathematics, including logic and combinatorics, and a strong research community has been forged, with a common research agenda which has influenced several important areas of computer science. The last Dagstuhl-like meeting of this research community before this seminar had been Aux Houches in 2012.

The principal goals of the seminar have been the following:

1. To identify fresh research challenges in the area of finite and algorithmic model theory, arising from the main application areas and to make new connections between core research in FAMT and emerging application areas, such as logic and learning.
2. To transfer knowledge from emerging methods and techniques in core FAMT to application areas.
3. To strengthen the research community in FAMT, especially by integrating younger members into it.

Organisation and Activities

The organisers developed a schedule consisting of three invited one-hour survey talks, more focussed regular contributions proposed by the participants, an open problem session, and a final discussion about the state and perspectives on the future of FAMT. The three survey talks were given by

- Wied Pakusa (Oxford) on recent achievements concerning the quest for a logic for polynomial time, focussing on Rank Logic and on Choiceless Polynomial Time,
- Dan Suciu (Washington) on highlights of the connections between FAMT and databases.
- Martin Grohe (Aachen) on new developments in machine learning and connections to FAMT.

In addition, 22 other participants gave regular talks on their recent work on topics of FAMT. A further social highlight was a superb concert on Thursday evening performed by Jan Van den Bussche (violin), Wolfgang Thomas (violin) and Jouko Väänänen (piano).

Outcomes

The seminar exceeded our expectations in achieving our principal goals. The invited talk by Pakusa gave an overview on the status of the ongoing pursuit for a logic for PTIME, and demonstrated the depth and technical sophistication that FAMT has reached. This talk was complemented by several insightful presentations on new and deep work on core topics of finite model theory. A particular highlight was the double presentation by Torunczyk and Siebertz, who brought in new methods from stability theory to the study of the finite model theory of sparse structures.

The invited talks by Grohe and Suciu explored new connections between the core methods of finite model theory and emerging areas of applications. These talks were also complemented by several other talks on new directions in FAMT and on interactions of FAMT with other areas. In particular, Grädel's talk focussed on new work in the area of database provenance, while Atserias'

talk discussed a priori unexpected connections between constraint satisfaction and quantum information theory.

Overall, the presentations at the seminar were highly stimulating, and we know through discussions during and after the seminar that the new work presented has motivated others to take up the explorations of these questions. Based on the feedback received, we believe that this seminar will serve as a catalyst for new research directions in FAMT.

The organizers regard the seminar as very successful. As reflected in the final discussion, there was a consistent sentiment expressed by the participants that the FAMT community is in very healthy state. There are interesting new developments and exciting results in different directions, there is a strengthening of traditional connections to areas, such as databases and verification, but also new connections are emerging with such areas as knowledge representation, learning theory, logics for dependence and independence, and quantum information theory. Finally, and perhaps more importantly, there is an infusion of several outstanding young researchers who have the interest and hold the promise to advance FAMT in the years to come.

The participants clearly expressed the wish to a have a next meeting of the FAMT community, be it in Dagstuhl or elsewhere, within the next two to three years.

The organizers are grateful to the Scientific Directorate of the Center for its support of this workshop and the staff of Schloss Dagstuhl for the perfect organisation of our stay and their hospitality.

6.54 Deduction Beyond First-Order Logic

Organizers: Jasmin Christian Blanchette, Carsten Fuhs, Viorica Sofronie-Stokkermans, and Cesare Tinelli

Seminar No. 17371

Date: September 10–15, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.26

Participants: Franz Baader, Christoph Benzmüller, Nikolaj S. Bjorner, Jasmin Christian Blanchette, James Brotherston, Chad E. Brown, Hans de Nivelle, Pascal Fontaine, Carsten Fuhs, Jürgen Giesl, Georges Gonthier, Reiner Hähnle, Swen Jacobs, Moa Johansson, Cezary Kaliszyk, Deepak Kapur, Chantal Keller, Cynthia Kop, Konstantin Korovin, K. Rustan M. Leino, Tomer Libal, Tobias Nipkow, Naoki Nishida, Andrei Paskevich, Ruzica Piskac, Andrei Popescu, Andrew Joseph Reynolds, Philipp Rümmer, Renate Schmidt, Stephan Schulz, Thomas Sewell, Natarajan Shankar, Mihaela Sighireanu, Viorica Sofronie-Stokkermans, Alexander Steen, Sorin Stratulat, Thomas Ströder, Martin Suda, Laurent Théry, Cesare Tinelli, Josef Urban, Christoph Weidenbach

Much research on automated deduction has traditionally focused on automated reasoning in first-order logic. First-order logic with equality is generally considered a sweet spot on the logic design continuum. Yet, from the point of view of several applications it can be too restrictive as a modeling and reasoning tool. In recent years, there has been a realization that while first-order reasoning is very useful to discharge the bulk of proof obligations, it must be tightly integrated with richer features to be useful in many applications. Practical problems often need a mixture of first-order proof search and some more advanced reasoning, for instance, about non-first-order-axiomatisable theories, higher-order formulas, or simply higher-level reasoning steps.

First-order logic cannot be used to finitely axiomatize many interesting theories, such as those including transitive closure operators, inductive predicates, datatypes, and standard arithmetic on integers or reals. Even provers that provide native support for some of these theories typically fail to prove trivial-looking problems because they lack general support for mathematical induction. Some applications need a richer set of constructs than those provided by first-order logic such as, for instance, the separating conjunction ($*$) and magic wand ($-\!\!*$) connectives of Separation Logic or the disjunctive well-foundedness predicates used in HSF, a popular approach to software model checking based on first-order Horn logic.

There are potential synergies between automatic first-order proving and verification methods developed in the context of richer logics. However, they have not received enough attention by the various deduction sub-communities so far. In general, there is a cultural gap between the various deduction communities that hinders cross-fertilization of ideas and progress.

This Dagstuhl Seminar brought together experts in automated reasoning in first-order logic and researchers working on deduction methods and tools that go beyond first-order logic. The latter included specialists on proof methods for induction, proof planning, and other higher-order or higher-level procedures; and consumers of deduction technology whose specification languages contain non-first-order features. The main goal of the seminar was to exchange ideas and explore ways to facilitate the transition from first-order to more expressive settings.

Research questions that were discussed and answered at the seminar included the following:

- What higher-order features do applications need, and what features can be incorporated smoothly in existing first-order proof calculi and provers?
- How can we best extend first-order reasoning techniques beyond first-order logic?
- Can proof-assistant-style automation and first-order reasoning techniques be combined in a synergetic fashion?
- What are good strategies for automatic induction and coinduction or invariant synthesis?
- Is a higher layer of reasoning, in the spirit of proof planning, necessary to solve more difficult higher-order problems?

Fig. 6.13
"Notes from @Dagstuhl day 3 part 1". Twitter post by 17471 Dagstuhl Seminar participant Mirjam P. Eladhari. https://twitter.com/MirjamPE/status/933300298826870784. Photo courtesy of Mirjam P. Eladhari.

6.55 Cybersafety in Modern Online Social Networks

Organizers: Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, and Thorsten Strufe

Seminar No. 17372

Date: September 10–13, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.47

Participants: Zinaida Benenson, Jeremy Blackburn, Emiliano De Cristofaro, Julien Dreux, Manuel Egele, Julien Freudiger, Oana Goga, Seda F. Gürses, Huy Kang Kim, Christiane Kuhn, Srijan Kumar, Ilias Leontiadis, Filippo Menczer, Prateek Mittal, Alexandra Olteanu, Awais Rashid, Ahmad-Reza Sadeghi, Stefan Schiffner, Michael Sirivianos, Gianluca Stringhini, Thorsten Strufe, Savvas Zannettou

The Dagstuhl Seminar 17372 "Cybersafety in Modern Online Social Networks" was a short two and a half day seminar, which took place September 10th–13th, 2017. Its main goal was to bring together researchers from various research areas related to cyberfraud and cybersafety in online social network, and to inspire them to exchange results, practical requirements, and ethical/legal implications related to user-driven research.

First Day. The seminar started with a short self-introduction of all the participants, then, we had an initial brainstorming session to identify main topics of interest, various aspects involved in them, and a balance in terms of interdisciplinary representation. Specifically, we focused on scams happening in online social network and hate speech, while paying special attention to the protection of minors. The aspects discussed were related to algorithmic, user, understanding/modeling, ethical, and privacy aspects of working in this line of research.

The brainstorming session concluded with the discussion of the following tangible research directions:

1. We should work on detection, prevention, and mitigation of hate speech.
2. All solutions should be in accordance of regulations.
3. We should pay particular attention to false positives, as users can easily lose their trust in the platform.
4. We should take into consideration the role of proxies, which act as biases on the data.
5. We should focus on counter-terrorism research that aims to distinguish vulnerable population in order to recruit them for propaganda purposes.

We then had four long talks throughout the day. The first speaker, Jeremy Blackburn (University of Alabama at Birmingham, US), described his work on cyberbullying and hate speech that includes studying behavior on video games as well as fringe Web communities like 4chan. The second speaker, Filippo Menczer (Indiana University – Bloomington, US) ,presented how misinformation is spread on Twitter. Specifically, he presented how false information as well as the respective fact-checking efforts are diffused on the Twitter network. The third speaker, Gianluca Stringhini (University College London, GB) presented his work on cyberfraud and scams, focusing on deceptive techniques employed by malicious users in order to scam benign users on online dating sites. The last speaker of the first day was Awais Rashid (Lancaster University, GB), who described his work related to child sex offenders and how he coordinated with Police bodies in order to undertake research on this topic. Also, he presented the ethical considerations when doing research with sensitive data, like those used for this study.

Second Day. The morning of the second day focused on giving an overview of work done on a variety of topics related to the main topics of the seminar (through short talks from the participants). More specifically, Zinaida Benenson (Universitat Erlangen – Nurnberg, DE), described her work on spear fishing, where malicious users aim to deceive users by sending email that contain malicious URLs. Then, Michael Sirivianos (Cyprus University of Technology- Lemesos, CY) presented his work on how to combat friend spam by analyzing the underlying network of social rejections. The next talk was by Alexandra Olteanu (IBM TJ Watson Research Center – Yorktown Heights, US), who discussed some preliminary results on work done on hate speech. Srijan Kumar (Stanford University, US) then showed how sockpuppet accounts are used in social networks to change and manipulate the opinions of other users of the platform. Savvas Zannettou (Cyprus University of Technology – Lemesos, CY) presented his research on how news propagates across multiple Web communities, and how to measure their influence. Then, Manuel Egele (Boston University, US), presented COMPA, which is a system that captures the behavioral profile of the user in order to identity possible account compromises. Huy Kang

Kim (Korea University – Seoul, KR) talked about malicious users exploiting video games to make money. Next, Oana Goga (MPI-SWS – Saarbrucken, DE) described how online identities can be strengthen by combining multiple weak identities. The last talk was by Julien Freudiger (Apple Computer Inc. – Cupertino, US), who covered public privacy and safety guidelines used at Apple.

The afternoon was dedicated to two parallel working groups focused on discussions around a particular topic, specifically, one was about future directions on hate speech research, and another about ethical considerations that researchers should keep in mind when working with users or user data.

Third Day. The final day of the seminar had two more parallel working groups, one on research directions related to cyberfraud in online social network, and another on algorithmic biases and possible solutions to avoid it. We then had a discussion summarizing the work and the discussion done in the various working groups and ended up with ideas for future events, collaborations, and follow-ups.

Acknowledgments. The organizers of this workshop acknowledge research funding from the European Union's Horizon 2020 Research and Innovation programme under the Marie Skłodowska-Curie Grant Agreement No 691025. The organizers would like to thank the Schloss Dagstuhl for the professional, productive, and enjoyable atmosphere it provides and for their invaluable support. Finally, we are grateful to and Seda Guerses for taking notes during two of the working groups and to Savvas Zannettou for coordinating the writing of this report and taking notes throughout the seminar.

6.56 Recent Trends in Knowledge Compilation

Organizers: Adnan Darwiche, Pierre Marquis, Dan Suciu, and Stefan Szeider

Seminar No. 17381

Date: September 17–22, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.62

Participants: Antoine Amarilli, Lameck Mbangula Amugongo, Paul Beame, Arpita Biswas, Pierre Bourhis, Simone Bova, Florent Capelli, Ondrej Cepek, Zaineb Chelly Dagdia, Arthur Choi, YooJung Choi, Adnan Darwiche, Ronald de Haan, Hélène Fargier, Robert Ganian, Martin Grohe, Henry A. Kautz, Batya Kenig, Frederic Koriche, Oliver Kullmann, Jean-Marie Lagniez, Neha Lodha, Meena Mahajan, Joao Marques-Silva, Pierre Marquis, Wannes Meert, Stefan Mengel, Shin-ichi Minato, Alexandre Niveau, Jakob Nordström, Dan Olteanu, Igor Razgon, Subhrajit Roy, Lisset Y. Salinas Pinacho, Scott Sanner, Rahul Santhanam, Marco Schaerf, Pierre Senellart, Laurent Simon, Friedrich Slivovsky, Dan Suciu, Stefan Szeider, Guy Van den Broeck

Knowledge compilation (KC) is a research topic which aims to investigate the possibility of circumventing the computational intractability of hard tasks, by preprocessing part of the available information, common to a number of instances. Pioneered almost three decades ago, KC is nowadays a very active research field, transversal to several areas within computer science. Among others, KC intersects knowledge representation, constraint satisfaction, algorithms, complexity theory, machine learning, and databases.

The results obtained so far take various forms, from theory (compilability settings, definition of target languages for KC, complexity results, succinctness results, etc.) to more practical results (development and evaluation of compilers and other preprocessors, applications to diagnosis, planning, automatic configuration, etc.). Recently, KC has been positioned as providing a systematic method for solving problems beyond NP, and also found applications in machine learning.

The goal of this Dagstuhl Seminar was to advance both aspects of KC, and to pave the way for a fruitful cross-fertilization between the topics, from theory to practice.

The program included a mixture of long and short presentations, with discussions. Several long talks with a tutorial flavor introduced the participants to the variety of aspects in knowledge compilation and the diversity of techniques used. System presentations as well as an open problem session were also included in the program.

Fig. 6.14
"maybe don't swing any higher…" a yes on the vr swing. Picture on flickr by 17392 Dagstuhl Seminar participant m. c. schraefel. https://flic.kr/p/YvyhhU. Photo courtesy of m. c. schraefel.

6.57 Approaches and Applications of Inductive Programming

Organizers: Ute Schmid, Stephen H. Muggleton and Rishabh Singh

Seminar No. 17382

Date: September 17–20, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.86

Participants: Aws Albarghouthi, Peter Buhler, Lidia Contreras-Ochando, Andrew Cropper, Luc De Raedt, Richard Evans, Cesar Ferri Ramirez, Elena Glassman, Katsumi Inoue, Frank Jäkel, Susumu Katayama, Martin Möhrmann, Stephen H. Muggleton, David Nieves Cordones, Hila Peleg, Oleksandr Polozov, Fabrizio Riguzzi, Ute Schmid, Sebastian Seufert, Michael Siebers, Rishabh Singh, Armando Solar-Lezama, Claes Strannegård, Janis Voigtländer, Christina Zeller

Inductive programming (IP) addresses the automated or semi-automated generation of computer programs from incomplete information such as input-output examples, constraints, computation traces, demonstrations, or problem-solving experience [5]. The generated – typically declarative – program has the status of a hypothesis which has been generalized by induction. That is, inductive programming can be seen as a special approach to machine learning. In contrast to standard machine learning, only a small number of training examples is necessary. Furthermore, learned hypotheses are represented as logic or functional programs, that is, they are represented on symbol level and therefore are inspectable and comprehensible [8, 17, 18]. On the other hand, inductive programming is a special approach to program synthesis. It complements deductive and transformational approaches [2, 14, 20]. In cases where synthesis of specific algorithm details that are hard to figure out by humans inductive reasoning can be used to generate program candidates from either user-provided data such as test cases or from data automatically derived from a formal specification [16].

Inductive program synthesis is of interest for researchers in artificial intelligence since the late sixties [1]. On the one hand, the complex intellectual cognitive processes involved in producing program code which satisfies some specification are investigated, on the other hand methodologies and techniques for automating parts of the program development process are explored. One of the most relevant areas of application of inductive programming techniques is end-user programming [3, 4, 12]. For example, the Microsoft Excel plug-in Flashfill synthesizes programs from a small set of observations of user behavior [6–8]. Related applications are in process mining and in data wrangling [11]. Inductive programming in general offers powerful approaches to learning from relational data [13, 15] and to learning from observations in the context of autonomous intelligent agents [10, 17]. Furthermore, inductive programming can be applied in the context of teaching programming [19, 21].

Recent Trends and Applications

When the first Dagstuhl Seminar on Approaches and Applications of Inductive Programming took place in 2013, the following trends could be identified:

- Combining different approaches to inductive programming to leverage their complementary strengths.
- New inductive programming approaches based on adapting and using well-developed techniques such as SAT-solving.
- Putting inductive programming to application, for example in the areas of automated string manipulations in spreadsheets or web programming.
- Applying concepts of inductive programming to cognitive models of learning structural concepts.

One of the major outcomes of the first Dagstuhl Seminar was a joint publication in the Communications of the ACM [8] where these trends and first applications and results were described. In the seminar 2015, the following additional trends were identified:

- Evaluation of inductive programming approaches – in relation to general intelligence and in relation to standard machine learning.
- Application of inductive programming to teaching programming.
- Inductive programming as a model of human inductive learning.

The main outcomes of the second seminar were (1) a joint publication in the Artificial Intelligence Journal with respect to the evaluation of computational models solving intelligence test problems – among them inductive programming systems [9], (2) a joint publication addressing comprehensibility as a second criterium to evaluate machine learning approaches besides

accuracy [18], and (3) a NIPS'2016 workshop on Neural Nets and Program Induction[51].

Based on the results of the second seminar, the focus of the third seminar has been on the following aspects:

- Identifying the specific contributions of inductive programming to machine learning research and applications of machine learning, especially identifying problems for which inductive programming approaches more suited than standard machine learning approaches, including deep learning.
- Establishing criteria for evaluating inductive programming approaches in comparison to each other and in comparison to other approaches of machine learning and providing a set of benchmark problems.
- Discussing current applications of inductive programming in enduser programming and programming education and identifying further relevant areas of application.
- Establishing stronger relations between cognitive science research on inductive learning and inductive programming.

In the seminar, we brought together researchers from different areas of computer science – especially from machine learning, AI, declarative programming, and software engineering – and researchers from cognitive psychology interested in inductive learning as well as in teaching and learning computer programming. Furthermore, participants from industry presented current as well as visionary applications for inductive programming.

The seminar was opened with lecture style talks introducing the four major approaches of inductive programming: Inductive functional programming, inductive logic programming, inductive probabilistic logical programming, and programming by example.

Talks covered current developments of IP algorithms, challenging applications –especially in data wrangling and in education –, and relations of IP to cognition.

In addition, several system demons and tutorials were given: Igor and EasyIgor (by Sebastian Seufert and Ute Schmid), MagicHaskeller (by Susumu Katayama), Sketch (by Armando Solar-Lezama), PROSE (by Oleksandr Polozov), Slipcover (by Fabrizio Riguzzi), and TaCLe (by Luc De Raedt).

The following topics were identified and further discussed in working groups during the seminar:

- How to determine relevancy of background knowledge to reduce search?
- Integrating IP with other types of machine learning, especially Deep Learning?
- Data wrangling as exiting area of application.

Additional topics identified as relevant have been anomaly detection, noise, robustnes, as well as non-example based interaction (e.g., natural language).

Concluding remarks and future plans

In the wrapping-up section, we collected answers to the question

"What would constitute progress at Dagstuhl 2019?"

The most prominent answers were

- make available systems, data sets (via IP webpage[52])
- compare systems
- common vocabulary, work on applications attempted by others: drawing problems, string transformation, general ai challenge, benchmarks starexec, learn robot strategy, grammar learning what is inductive programming
- open problems

As the grand IP challenge we came up with: **An IP program should invent an algorithm publishable in a serious journal (e.g., an integer factorization algorithm) or win a programming competition!**

References

1 A. W. Biermann, G. Guiho, and Y. Kodratoff, editors. *Automatic Program Construction Techniques*. Macmillan, New York, 1984.

2 Rastislav Bodik and Emina Torlak. Synthesizing programs with constraint solvers. In *CAV*, page 3, 2012.

3 A. Cypher, editor. *Watch What I Do: Programming by Demonstration*. MIT Press, Cambridge, MA, 1993.

4 Allen Cypher, Mira Dontcheva, Tessa Lau, and Jeffrey Nichols, editors. *No Code Required: Giving Users Tools to Transform the Web*. Elsevier, 2010.

5 P. Flener and U. Schmid. Inductive programming. In C. Sammut and G. Webb, editors, *Encyclopedia of Machine Learning*, pages 537–544. Springer, 2010.

6 Sumit Gulwani. Automating string processing in spreadsheets using input-output examples. In *38th Symposium on Principles of Programming Languages*. ACM, 2011.

7 Sumit Gulwani, William R. Harris, and Rishabh Singh. Spreadsheet data manipulation using examples. *Communications of the ACM*, 55(8):97–105, 2012.

8 Sumit Gulwani, José Hernández-Orallo, Emanuel Kitzelmann, Stephen H. Muggleton, Ute Schmid, and Benjamin G. Zorn. Inductive programming meets the real world. *Commununications of the ACM*, 58(11):90–99, 2015.

9 José Hernández-Orallo, Fernando Martínez-Plumed, Ute Schmid, Michael Siebers, and David L Dowe. Computer models solving intelligence test problems: Progress and implications. *Artificial Intelligence*, 230:74–107, 2016.

10 P. Langley and D. Choi. A unified cognitive architecture for physical agents. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, Boston, MA, 2006. AAAI Press.

11 Vu Le and Sumit Gulwani. Flashextract: A framework for data extraction by examples. *ACM SIGPLAN Notices*, 49(6):542–553, 2014.

12 Henry Lieberman, editor. *Your Wish is My Command: Programming by Example*. Morgan Kaufmann, San Francisco, 2001.

13 D. Lin, E. Dechter, K. Ellis, J.B. Tenenbaum, and S.H. Muggleton. Bias reformulation for one-shot function induction. In *Proceedings of the 23rd European Conference on Artificial Intelligence (ECAI 2014)*, pages 525–530, Amsterdam, 2014. IOS Press.

[51] https://uclmr.github.io/nampi/
[52] www.inductive-programming.org

14 Zohar Manna and Richard Waldinger. A deductive approach to program synthesis. *ACM Transactions on Programming Languages and Systems*, 2(1):90–121, 1980.

15 Stephen H. Muggleton, Dianhuan Lin, and Alireza Tamaddoni-Nezhad. Meta-interpretive learning of higher-order dyadic datalog: predicate invention revisited. *Machine Learning*, 100(1):49–73, 2015.

16 Reudismam Rolim, Gustavo Soares, Loris D'Antoni, Oleksandr Polozov, Sumit Gulwani, Rohit Gheyi, Ryo Suzuki, and Bjoern Hartmann. Learning syntactic program transformations from examples. *arXiv preprint arXiv:1608.09000*, 2016.

17 Ute Schmid and Emanuel Kitzelmann. Inductive rule learning on the knowledge level. *Cognitive Systems Research*, 12(3):237–248, 2011.

18 Ute Schmid, Christina Zeller, Tarek Besold, Alireza Tamaddoni-Nezhad, and Stephen Muggleton. How does predicate invention affect human comprehensibility? In *International Conference on Inductive Logic Programming*, pages 52–67. Springer, 2016.

19 Rishabh Singh, Sumit Gulwani, and Armando Solar-Lezama. Automated feedback generation for introductory programming assignments. *ACM SIGPLAN Notices*, 48(6):15–26, 2013.

20 Douglas R. Smith. The synthesis of LISP programs from examples: A survey. In *Automatic Program Construction Techniques*, pages 307–324. Macmillan, 1984.

21 Christina Zeller and Ute Schmid. Automatic generation of analogous problems to help resolving misconceptions in an intelligent tutor system for written subtraction. In *Proceedings of the Workshop on Computational Analogy at the 24th International Conference on Case Based Reasoning (ICCBR 2016, Atlanta, GA, 31th October to 2nd November 2016)*, 2016.

6

6.58 Deep Learning for Computer Vision

Organizers: Daniel Cremers, Laura Leal-Taixé, and René Vidal

Seminar No. 17391

Date: September 24–29, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.109

Participants: Joan Bruna Estrach, Pratik Chaudhari, Daniel Cremers, Alexey Dosovitskiy, Vittorio Ferrari, Thomas Frerix, Jürgen Gall, Silvano Galliani, Ravi Garg, Raja Giryes, Kristen Grauman, Benjamin Haeffele, Philip Häusser, Caner Hazirbas, Iasonas Kokkinos, Hildegard Kühne, Christoph H. Lampert, Laura Leal-Taixé, Stephane Mallat, Michael Möller, Emanuele Rodolà, Niko Sünderhauf, Rahul Sukthankar, Anton van den Hengel, Jan Van Gemert, Andrea Vedaldi, René Vidal, Christoph Vogel

The paradigm that a machine can learn from examples much like humans learn from experience has fascinated researchers since the advent of computers. It has triggered numerous research developments and gave rise to the concept of artificial neural networks as a computational paradigm designed to mimic aspects of signal and information processing in the human brain.

There have been several key advances in this area including the concept of back- propagation learning (essentially gradient descent and chain rule differentiation on the network weight vectors) by Werbos in 1974, later popularized in the celebrated 1984 paper of Rumelhart, Hinton and Williams. Despite a certain success in pattern recognition challenges like handwritten digit classification, artificial neural networks dropped in popularity in the 1990s with alternative techniques such as support vector machines gaining attention.

With increasing computational power (and in particular highly parallel GPU architectures) and more sophisticated training strategies such as layer-by-layer pretraining, supervised back-propagation and dropout learning, neural networks regained popularity in the 2000s and the 2010s. With deeper network architectures and more training data, their performance has drastically improved. Over the last couple of years, they have outperformed numerous existing algorithms on a variety of computer vision challenges such as object recognition, semantic segmentation and even stereo and optical flow estimation.

The aim of this Dagstuhl Seminar is to bring together leading experts from the area of machine learning and computer vision and discuss the state-of-the-art in deep learning for computer vision. During our seminar, we will address a variety of both experimental and theoretical questions such as:

1. In which types of challenges do deep learning techniques work well?
2. In which types of challenges do they fail? Are there variations of the network architectures that may enable us to tackle these challenges as well?
3. Which type of network architectures exist (convolutional networks, recurrent networks, deep belief networks, long short-term memory networks, deep Turing machines)? What advantages and drawbacks does each network architecture bring about?
4. Which aspects are crucial for the practical performance of deep network approaches?
5. Which theoretical guarantees can be derived for neural network learning?
6. What properties assure the impressive practical performance despite respective cost functions being generally non-convex?

6.59 Body-Centric Computing

Organizers: Steve Benford, Kristina Höök, Joseph Marshall, Florian Mueller, and Dag Svanes

Seminar No. 17392

Date: September 24–29, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.9.126

Participants: Josh Andres, Kathrin Gerling, Stefan Göbel, Nicholas Graham, Marianne Graves Petersen, Perttu Hämäläinen, Martin Jonsson, Elena Márquez Segura, Joseph Marshall, Anna Lisa Martin-Niedecken, Jochen Meyer, Florian Mueller, Corina Sas, Thecla Schiphorst, m.c. schraefel, Anna Stahl, Dag Svanes, Jakob Tholander, Elise van den Hoven, Rainer Wieching, Florian Wolling

The rise of technology that can support the active human body – in contrast to the previously prevalent paradigm of interacting with computers while sitting still – such as wearables, quantified self systems and mobile computing highlights an opportunity for a new era of "body-centric computing". However, most work in this area has taken quite an instrumental perspective, focusing on achieving extrinsic performance objectives. Phenomenology, however, highlights that it is also important to support the experiential perspective of living an active life, that is, technology should also help people focus on their lived experiences to deepen their understanding and engagement with their own bodies. We find that despite the work on embodiment, the use of technology to support the corporeal, pulsating, felt body has been notably absent. We believe the reason for this is due to limited knowledge about how to understand, analyse and correlate the vast amount of data from the various sensors worn by individuals and populations in real-time so that we can present it in a way that it supports people's felt experience. In order to drive such an agenda that supports both instrumental and experiential perspectives of the active human body, this seminar brought together leading experts from industry and academia, including those who are central to the development of products and ideas relating to wearables, mobile computing, quantified self, data analysis and visualization, sports science, exertion games, computer sports science as well as phenomenology. The goal was to address key questions around the use of sensor data to support both instrumental and experiential perspective of the active human body and to jump-start collaborations between people from different backgrounds to pioneer new approaches for a body-centric computing future.

6

6.60 Quantum Cryptanalysis

Organizers: Michele Mosca, Nicolas Sendrier, Rainer Steinwandt, and Krysta Svore

Seminar No. 17401

Date: October 1–6, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.1

Participants: Gorjan Alagic, Shi Bai, Gustavo Banegas, Daniel J. Bernstein, Jean-François Biasse, Alexei Bocharov, Johannes A. Buchmann, Yfke Dulek, Serge Fehr, Tommaso Gagliardoni, Vlad Gheorghiu, Maria Isabel González Vasco, Sean Hallgren, Peter Hoyer, Andreas Hülsing, David Jao, Stacey Jeffery, Elena Kirshanova, Stavros Kousidis, Thijs Laarhoven, Bradley Lackey, Tanja Lange, Yi-Kai Liu, Alexander May, Michele Mosca, Michael Naehrig, Anderson Nascimento, Maria Naya-Plasencia, Phong Q. Nguyen, Ray Perlner, Martin Roetteler, Alexander Russell, John M. Schanck, Claus Peter Schnorr, Nicolas Sendrier, Daniel Smith-Tone, Rainer Steinwandt, Adriana Suárez Corona, Tsuyoshi Takagi, Jean-Pierre Tillich, Dominique Unruh, Frank K. Wilhelm

Motivation and scope

Like its predecessors, this fourth installment of a Dagstuhl seminar on *Quantum Cryptanalysis* was devoted to studying cryptographic solutions that might be suitable for standardization in the post-quantum setting and to studying quantum attacks against currently deployed cryptographic solutions. Two main thrusts were of particular interest:

Algorithmic innovation. Quantum resources can be used in various way for attacking cryptographic solutions, and the seminar included multiple presentations on exploiting quantum resources for cryptanalytic purposes. Both attacks on symmetric and asymmetric primitives were considered, and there were lively discussions on the feasibility of mounting particular types of attacks. Complementing the presentations on quantum attacks, the program included presentations on advanced classical algorithms, raising the question of identifying possibilities to speed up such classical attack venues through quantum "subroutines."

Quantum resource estimation. It goes without saying that asymptotic improvements are of great interest when trying to tackle computational problems underpinning the security of cryptographic constructions. However, when looking at an actually deployed scheme, quantifying the exact resources (such as the number of qubits) needed by an attacker is relevant to judge the practical impact of a proposed attack strategy. The seminar included presentations on the estimation of resources for attacking some prominent cryptographic schemes.

As expected from a seminar with this title, many talks were indeed devoted to cryptanalysis, but the program also included presentations on establishing provable security guarantees in a post-quantum scenario. With the field becoming more mature, we did not schedule much time for survey talks. However, we did include a presentation on the *status of the development of quantum computers* in the program, thereby helping to get a better idea of potential obstacles when trying to implement quantum cryptanalytic attacks.

Organization

This was the fourth Dagstuhl seminar devoted entirely to quantum cryptanalysis, and as in the prior editions the set of participants included both experts in quantum algorithms and experts in classical cryptography. Some of the participants had already participated in earlier editions of this seminar series, but a number of colleagues attended such a seminar — or any Dagstuhl event — for the first time. In total, we had 42 participants from academia, government, and industry. This time we also included an open problem session in the program, which will hopefully help to stimulate further work in this vibrant research area. In the schedule we tried to leave sufficient time for discussions and for collaborative work in smaller groups. In line with the Dagstuhl tradition, no presentations were scheduled for Wednesday afternoon, and the seminar participants could devote the afternoon to a hike, an excursion, or to their research.

Results and next steps

Over the course of the years, communication and collaboration between the classical cryptographic and the quantum algorithmic research communities has intensified, and many colleagues cross traditional discipline boundaries. As evidenced in the seminar, available quantum cryptanalytic results can go well beyond asymptotic statements and include rather fine-grained resource counts. The seminar covered the analysis of both symmetric and asymmetric primitives, and ongoing efforts toward standardizing quantum-safe cryptographic solutions are likely to stimulate more progress, in particular on the quantum cryptanalysis of asymmetric cryptographic primitives.

6.61 Hyperspectral, Multispectral, and Multimodal (HMM) Imaging: Acquisition, Algorithms, and Applications

Organizers: Gonzalo R. Arce, Richard Bamler, Jon Yngve Hardeberg, Andreas Kolb, and Shida Beigpour

Seminar No. 17411

Date: October 8–13, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.14

Participants: Gonzalo R. Arce, Richard Bamler, Gladimir V. G. Baranoski, Shida Beigpour, Ela Claridge, Donald G. Dansereau, Hilda Deborah, Nicolas Dobigeon, Paul D. Gader, Peter Gege, Sony George, Jon Yngve Hardeberg, Peter Haring Bolivar, Andreas Kolb, Dietrich Paulus, Christoph Pomrehn, Noël Richard, Antonio Robles-Kelly, Christoph Markus Schikora, Jean-Baptiste Thomas, Shoji Tominaga, Devis Tuia, Eva M. Valero Benito, Rudolf Verdaasdonk, Anna Katharina Wigger, Masahiro Yamaguchi, Naoto Yokoya, Xiaoxiang Zhu

On the last day of the seminar, the attendees had a very intense discussion about the usefulness of the seminar itself, the grand challenges related to the highly interdisciplinary field of research, and the next steps that should be taken in order to further improve on the cross-fertilizing effects in hyperspectral, multispectral, and multimodal imaging.

Take Home Messages

All attendees agreed on the high quality and open mindedness of the discussions at both, the group level, e.g., in the plenary sessions and the working groups, and also on personal level. All participants assess this Dagstuhl seminar as a great success, especially due to the interdisciplinary discussion and the new insights resulting from this. Despite differences of the individual fields present in the seminar, e.g., remote sensing, color reproduction, and material classification, and despite the wide variety of applications such as medical, environmental monitoring, and arts, a large set of common questions and problems could be identified. All attendees highly appreciated the fact, that unlike in conferences, which usually have a rather narrow perspective on HMM challenges and solutions, as they usually address a single community with a very similar perspective on the field, this seminar brought together people with very different points of view.

This Dagstuhl seminar was a starting point of a number of connections that could be established directly, and several mid-, and maybe even long-term collaborations and joint research actions. There have been several highlights related to the full pipeline from data acquisition, via data processing to applications.

Grand Challenges

On the basis of common and interdisciplinary ground setup in this seminar, several challenges have been identified, which the seminar's attendees see as important to be addressed in further research and engineering work.

Data Acquisition Independent of the specific range of the addressed electromagnetic spectrum, the seminar participants see a severe restriction in the usage of HMM sensors due to their inflexibility, e.g., in selecting spectral bands, bulkiness, high calibration efforts, and acquisition speed (see also the working group report on this topic). Enhancing on these limiting factors has the potential to bring about fundamentally new spins in various application domains. Some approaches presented at the workshop have the potential to push back these limits to some degree. On the other hand, most likely there will be no general purpose HMM acquisition device available in the next decades that covers the majority of application requirements. Still, the seminar attendees agree on the importance of enhancing the applicability of existing and future acquisition devices towards more flexible band selection, fast and efficient, (semi-)automated calibration, and, for some applications, compactness. Ideally, future research provides means for an abstract definition of application specific characteristics from which a specific selection and/ or instantiation of an acquisition device can be deduced.

Data Processing and Validation Regarding data processing and validation, three main topics have been discussed: The usefulness and limitations of machine learning and, especially, deep learning (see working group report), the importance of verified and metric data, and the need for a proper reference and benchmarking data set. Even though there are and have been ongoing activities in spectral normalization and validation, e.g., on the level of CIE or other standardization institutions, or in the field of metrology,

there still is the lack of widely existing and accepted methods and data even if restricted to specific fields of application.

The seminar participants see a huge potential in all three areas. Still, major obstacles have to be overcome in order to leverage these potentials. In machine learning/ deep learning, one main issue is the lack of guarantees that the results obey specific constraints to, e.g., physical limits or relations. The lack of verified, metric data, and proper reference and benchmarking data, on the other hand, can only be overcome if there is a stronger common basis for best practice within and, even more important, between the HMM sub-disciplines.

Information Exchange The existence of common information bases is tightly linked to the prior point regarding data processing and validation. So far, there are only few options and pseudo-standard for sharing data and algorithms. While there are good examples, e.g., Open CV library in computer vision, setting up this kind of "standard" is, and will be, much harder in the diverse and partially fragmented HMM research domain. Apparently, this chicken-egg problem can only be solved from within the involved research domains themselves by the normative power of fact of the actions taken by the researchers themselves.

Next Steps

Participants discussed various options for future activity as a follow-up on this Dagstuhl seminar. As one essential restriction of the discussion, attendees became aware of their own limitations in knowing all relevant work and requirements existing in the HMM research subfields. Therefore, the obvious approach to enhance the fields' convergence by publications, e.g., a special issue or book, and/ or workshops has not the highest priority, even though an introductory workshop or piece of literature for 1st year PhD students would be highly appreciated.

However, participants of Dagstuhl seminar see two main options to proceed in order to keep the initiated process of convergence going and to improve on at least two of the three main challenges identified, i.e., regarding data processing and validation and the exchange of information.

HMM Webpage: As there is a severe lack in common information widely used and recognized, the group of researchers who attended Dagstuhl seminar see the potential of a common, web-based information platform.

In this respect, Masahiro Yamaguchi is open to provide the already established web-link multispectral.org and Andreas Kolb will investigate options for setting up and hosting this kind of platform. In any case, this kind of activity needs to rest on several shoulders, thus the attendees are called to follow through with the activities, on the operative level.

Follow-up Dagstuhl Seminar: As Dagstuhl supports follow-up seminars, the attendees agree on the usefulness of having this kind of seminar in order to evaluate the common, interdisciplinary activities that arose from the first Dagstuhl seminar. In case of a new edition of the workshop, participants agree on having more industrial partners involved.

6.62 Internet of People

Organizers: Elizabeth M. Belding, Jörg Ott, Andrea Passarella, and Peter Reichl

Seminar No. 17412

Date: October 8 – 11 , 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.42

Participants: Panayotis Antoniadis, Chiara Boldrini, Dimitris Chatzopoulos, Nicola Dell, Peter Fatelnig, Markus Fiedler, Huber Flores, Heikki Hämmäinen, Tobias Hoßfeld, Paul Houghton, Pan Hui, Teemu Kärkkäinen, Eemil Lagerspetz, Anders Lindgren, Pietro Lio, Kirsi Louhelainen, Leonardo Maccari, Jörg Ott, Maximilian Ott, Andrea Passarella, Daniele Quercia, Peter Reichl, Jatinder Singh, Thorsten Strufe, Gareth Tyson, Ellen Zegura, Mariya Zheleva, Martina Zitterbart

The key objective of the seminar was to bring together a diverse group of researchers and practitioners to reflect on technological and social issues related to the emerging concept of Internet of People (IoP). The group of attendees was composed of 28 people with diverse expertise on the various areas of Internet, coming from Europe, US, Asia and Australia.

The group worked for two and a half days, and the work was organised on (i) seed talks, (ii) snippet talks on selected research topics related to IoP, and (iii) parallel group work. The group sessions were particularly productive, and attendees worked on many topics. Specifically, they covered the following topics: (i) IoP definition, (ii) IoP use cases, (iii) IoP and people; (iv) Privacy, security and trust; (v) IoP architecture, and (vi) transition towards IoP. Over the last day, the group again split in three sub-groups, to focus on conclusions and follow-up activities. Specifically, the three groups produced (i) guidelines for IoP toolkits, (ii) a possible IoP research agenda, and (iii) an IoP manifesto.

We managed to bring together a quite balanced group of 32 people with expertise in the design and implementation of wireless ad hoc networks of various types, human-computer interaction, community informatics, urban interaction design, ethnography, media studies, arts and design.

The main take-home message from the seminar is that IoP is an emerging research topic with a lot of potential. It spans many aspects, including but not limited to the set of topics identified for the group work. Each of the group works provided concrete guidelines on the selected topics, possibly providing focused research agenda for the future.

Most of all, we are very happy that attendees greatly enjoyed the seminar, including those attending for the first time a Dagstuhl event (about one third). We do believe that the seminar laid the grounds for future fruitful collaborations, and helped a lot in shaping the key ideas of the emerging research topic of IoP.

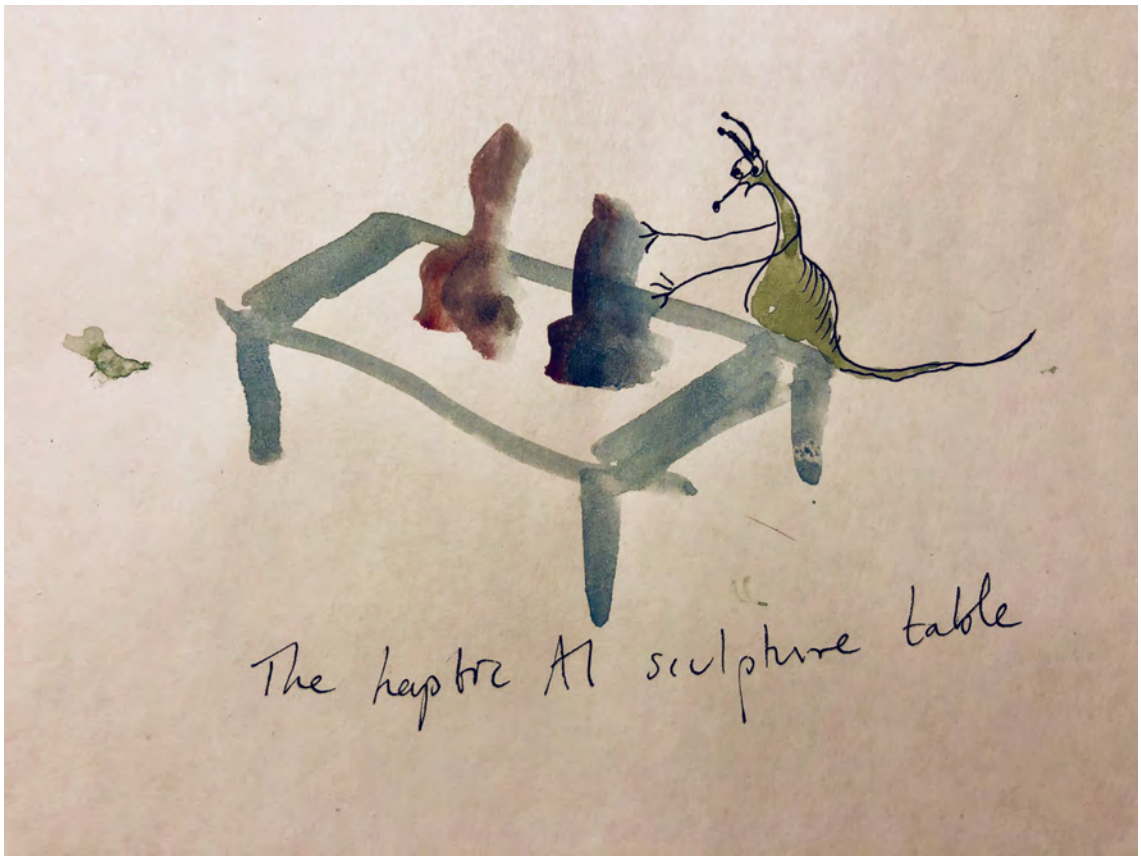

Fig. 6.15
"Notes from @Dagstuhl seminar on AI driven game design". Twitter post by 17471 Dagstuhl Seminar participant Mirjam P. Eladhari. https://twitter.com/MirjamPE/status/933045410880479232. Photo courtesy of Mirjam P. Eladhari.

6.63 Computational Proteomics

Organizers: Bernhard Küster, Kathryn Lilley, and Lennart Martens

Seminar No. 17421

Date: October 15–20, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.69

Participants: Magnus Arntzen, Nuno Bandeira, Harald Barsnes, Sebastian Böcker, Robert Chalkley, John Cottrell, Ileana M. Cristea, Bernard Delanghe, Eric Deutsch, Viktoria Dorfer, Julien Gagneur, Laurent Gatto, Marco Hennrich, Nico Jehmlich, Lukas Käll, Oliver Kohlbacher, Jeroen Krijgsveld, Bernhard Küster, Lydie Lane, Kathryn Lilley, Frédérique Lisacek, Lennart Martens, Gerben Menschaert, Bart Mesuere, Thilo Muth, Magnus Palmblad, Phillip Pope, Hannes Röst, Timo Sachsenberg, Veit Schwämmle, Stephen Tate, Elien Vandermarliere, Hans Vissers, Olga Vitek, Juan Antonio Vizcaino, Pieter-Jan Volders, Mathias Walzer, Ana L. Wang, Mathias Wilhelm, Paul Wilmes, Dennis Wolan, Henrik Zauber

The Dagstuhl Seminar 17421 "Computational Proteomics" discussed in-depth the current challenges facing the field of computational proteomics, while at the same time reaching out across the field's borders to engage with other computational omics fields at the joint interfaces. The issues that were discussed reflect the emergence of novel applications within the field of proteomics, notably proteogenomics (the identification of proteins based on sequence data obtained from prior genomics and/or transcriptomics analyses), and metaproteomics (the study of the combined proteome across an entire community of (micro-)organisms). These two new proteomics approaches share several challenges, which predominantly revolve around the sensitive identification of proteins from large databases while maintaining an acceptably low false discovery rate (FDR). The ramifications of these issues, and possible solutions, were first introduced in short but thought-provoking talks, followed by a plenary discussion to delineate the initial discussion sub-topics. Afterwards, working groups addressed these initial considerations in great detail.

In addition, both proteogenomics and metaproteomics suffer from coverage issues, as neither is currently capable of providing anywhere near a complete view on the true complexity of the (meta-)proteome. This issue is exacerbated by the fact that the true extent of the proteome remains unknown, and is likely to be time-dependent as well. As a result, a separate working group was created to discuss the issues and possible remedies related to proteome coverage.

The field of proteomics has, however, not only extended into novel application areas, but meanwhile also continues to see a strong development of novel technologies. Over the past few years, the most impactful of these is data-independent acquisition (DIA), which comes with its own unique computational challenges. On the one hand, the analysis of DIA data currently relies heavily on spectral libraries, which have so far been a rather niche product in proteomics (as opposed to, for instance, metabolomics, where spectral libraries have a much longer and much more fruitful history), while on the other hand, FDR estimation remains contested in DIA approaches. As a result, two further working groups were established during the seminar, one on the applications for, and methods to create spectral libraries, and the other on the specific challenge of calculating a reliable FDR when performing spectral library searching.

Another key topic of the seminar was the (orthogonal) re-use of public proteomics data, which focused on the provision of metadata for the assembled proteomics data, as this is the key bottleneck facing researchers who wish to perform large-scale re-analysis of public proteomics data, especially when the objective is to obtain biological knowledge. A working group was therefore created to explore the issues with metadata provision, and to explore means to ameliorate the current suboptimal metadata reporting situation.

Throughout the seminar, the topic of visualizing the acquired data and the obtained results cropped up with regularity. A corresponding working group was therefore set up to delineate the state-of-the-art in proteomics data visualization, and to explore the issues with, and opportunities of advanced visualizations in proteomics.

As a last core topic, a short introductory talk and subsequent working group was dedicated to the education of computational proteomics researchers, with special focus on their ability to work at the interfaces with other omics fields (genomics, transcriptomics, and metabolomics). This working group assembled an extensive list of already available materials, along with an overview of the different roles and specializations that can be found across informaticians, bio-informaticians, and biologists,

and how each field should evolve in order to bring these more closely together in the future.

In addition to above mentioned topic introduction talks, and the associated working groups, two talks illustrated specific topics of the seminar. Paul Wilmes showed his recent work in bringing metaproteomics together with advanced metatranscriptomics and metagenomics, showing that the flexible use of sequence assembly graphs at the nucleotide level opens up many highly interesting possibilities at the proteome level through enhanced identification. Nevertheless, it was observed that there is strong enrichment for genes with unknown function at the protein identification level, highlighting quite clearly that we have yet to achieve a more complete biochemical understanding of microbial ecosystems. Finally, Magnus Palmblad delighted the participants with a highly original talk on the exploration of mass spectrometry data (of both peptides as well as small molecules) through the five senses (sight, hearing, touch, smell, and taste).

6.64 Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions

Organizers: Anshu Dubey, Paul H. J. Kelly, Bernd Mohr, and Jeffrey S. Vetter

Seminar No. 17431

Date: October 22–27, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.84

Participants: Sadaf Alam, Michael Bader, Carlo Bertolli, Mauro Bianco, Alexandru Calotoiu, Bradford Chamberlain, Aparna Chandramowlishwaran, Kemal A. Delic, Christophe Dubach, Anshu Dubey, H. Carter Edwards, Jan Eitzinger, Todd Gamblin, Lin Gan, William D. Gropp, Philipp Gschwandtner, Mary W. Hall, Robert J. Harrison, Alexandra Jimborean, Paul H. J. Kelly, Andreas Klöckner, Kathleen Knobe, Seyong Lee, Naoya Maruyama, Chris Maynard, Simon McIntosh-Smith, Richard Membarth, Lawrence Mitchell, Bernd Mohr, Raymond Namyst, Simon Pennycook, Istvan Reguly, P. (Saday) Sadayappan, Sven-Bodo Scholz, Michelle Mills Strout, Nathan Tallent, Christian Terboven, Didem Unat, Ana Lucia Varbanescu, Jeffrey S. Vetter, Mohamed Wahib, Michele Weiland, Robert Wisniewski, Michael Wolfe

This report documents the program and the outcomes of Dagstuhl Seminar 17431 "Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions".

Performance Portability is a critical new challenge in extreme-scale computing. In essence, performance-portable applications can be efficiently executed on a wide variety of HPC architectures without significant manual modifications. For nearly two decades, HPC architectures and programming models remained relatively stable, which allowed growth of complex multidisciplinary applications whose lifecycles span multiple generations of HPC platforms.

Recently, however, platforms are growing much more complex, diverse, and heterogeneous - both within a single system and across systems and generations. Details already known from planned future systems indicate that this trend will continue (at least for the foreseeable future). Current and planned future large-scale HPC systems consist of complex configurations with a massive number of components. Each node has multiple multi-core sockets and often one or more additional accelerator units in the form of many-core nodes or GPGPUs, resulting in a heterogeneous system architecture. Memory hierarchies including caches, memory, and storage are also diversifying in order to meet multiple constraints: power, latency, bandwidth, persistence, reliability, and capacity. These factors are reducing portability, and forcing applications teams to either spend considerable effort porting and optimizing their applications for each specific platform, or risk owning applications that perform well on perhaps only one architecture. The latter option would still require porting and optimizing effort for each new generation of systems.

This Dagstuhl Seminar represented a unique opportunity to bring together international experts from the three research communities essential to tackling this performance portability challenge: developers of large-scale computational science software projects whose lifetime will span multiple generations of systems, researchers developing relevant parallel programming or system software technologies, and specialists in profiling, understanding, and modelling performance. The major research questions for the seminar were:

- To understand challenges, design metrics, and prioritize potential solutions for performance portability: Solutions will need to synthesize existing concepts across multiple fields, including performance and productivity modeling, programming models and compilation, architectures, system software.
- Management of data movement in complex applications: Diverse data movement patterns dictated by different devices form one of the largest impediments to portable performance. Addressing it will require cross-cutting solutions supporting more than one abstraction, and will allow scientists to balance tradeoffs in these factors prior to design, development, or procurement of an architecture, software stack, or application.
- Composability: Many applications require flexibility and composability because they address different physical regimes either within the same simulation, or in different instances of simulations.
- Pathways to impact on the research community: As the community becomes more reliant on both more complex architectures and software stacks, it is especially important that we develop the conceptual tools to facilitate research and practical solutions for performance portability. The impact of ignoring this topic could be potentially devastating to the quality and sustainability of computational science software, and consequently on the science and engineering research they support. Thus a key element of the seminar will be to tackle this challenge in major science community software projects.

The seminar started with a series of flash talks, where participants introduced themselves in a two-minute one-slide presentation summarizing their contribution or interest in the seminar by providing two to three bullet points on (i) Challenge/Opportunity (WHY?) (ii) Timeliness (WHY NOW?) (iii) Approaches (HOW?) and (iv) IMPACT (SO WHAT?). Each day started with a longer keynote presentation by a representative of one of the major stakeholders in the field, followed by short presentations by participants grouped in sessions with a common relevant theme. Each keynote or short talk session ended with an extensive question-and-answer session and open discussion slot in which all the speakers from the session took part.

The overall conclusion shared by all participants was that performance portability in extreme scale computing can be achieved, especially if parallel applications are designed with performance portability in mind from the beginning. Model complexity and performance portability both require that frameworks be designed with composable components incorporating layers of abstraction so that trade-offs can be reasoned about. Making legacy application performance portable still requires enormous efforts and expertise. In many instances it will likely require extensive refactoring. Similar design principles regarding formulation of a flexible and composable framework apply for legacy software refactoring, along with strong emphasis on rigorous verification built into the process. The seminar recognized the challenges faced by the applications in adopting abstractions; converting research prototypes to reliable production-grade product. The adverse structure of incentives for both applications and abstractions, and the complexity of formulating a process or collaboration between the two communities, may be bigger barriers than technical challenges in making performance portability feasible. It is critical that the involved communities and stakeholders are made aware of these challenges while seeking solutions for sustainable computational science projects.

6.65 Big Stream Processing Systems

Organizers: Tilmann Rabl, Sherif Sakr, and Martin Hirzel

Seminar No. 17441

Date: October 29 to November 3, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.10.111

Participants: Pramod Bhatotia, Albert Bifet, Michael H. Böhlen, Angela Bonifati, Jean-Paul Calbimonte, Paris Carbone, Emanuele Della Valle, Javier D. Fernández-García, Ashish Gehani, Manfred Hauswirth, Günter Hesse, Martin Hirzel, Ali Intizar, Asterios Katsifodimos, Nikos Katsipoulakis, Henning Kropp, Danh Le Phuoc, Alessandro Margara, Gianmarco Morales, Christoph Quix, Tilmann Rabl, Sherif Sakr, Kai-Uwe Sattler, Matthias J. Sax, Martin Strohbach, Hong-Linh Truong, Akrivi Vlachou, Thomas Weise, Yongluan Zhou

As the world gets more instrumented and connected, we are witnessing a flood of digital data that is getting generated, in a high velocity, from different hardware (e.g., sensors) or software in the format of streams of data. Examples of this phenomena are crucial for several applications and domains including financial markets, surveillance systems, manufacturing, smart cities and scalable monitoring infrastructure. In these applications and domains, there is a crucial requirement to collect, process, and analyze big streams of data in order to extract valuable information, discover new insights in real-time and to detect emerging patterns and outliers. Recently, several systems (e.g., Apache Apex, Apache Flink, Apache Storm, Heron, Spark Streaming,) have been introduced to tackle the real-time processing of big streaming data. However, there are several challenges and open problems that need to be addressed in order improve the state-of-the-art in this domain and push big stream processing systems to make them widely used by large number of users and enterprises. The aim of this seminar was to bring together active and prominent researchers, developers and practitioners actively working in the domain of big stream processing to discuss very relevant open challenges and research directions. The plan was to work on specific challenges including the trade-offs of the various design decisions of big stream processing systems, the declarative stream querying and processing languages, and the benchmarking challenges of big stream processing systems.

On Monday morning, the workshop officially kicked off with a round of introductions about the participants where adhoc clusters for the interests of the participants have been defined. The clusters have been revolving around the topics of systems, query languages, benchmarking, stream mining and semantic stream processing. The program of the seminar included 4 tutorials, one per day. On Monday, Martin Strohbach from AGT International presented different case studies and scenarios for large scale stream processing in different application domains. On Tuesday, we enjoyed the systems tutorial which has been presented by Paris Carbone from KTH Royal Institute of Technology, Thomas Weise from Data Torrent Inc. and Matthias J. Sax from Confluent Inc. Paris presented an interesting overview of the journey of stream processing systems, Thomas presented the recent updates about the Apache Apex system while Matthias presented an overview about the Apache Kafka and Kafka Streams projects. On Wednesday, Martin Hirzel from IBM TJ Watson Research Center presented a tutorial about the taxonomy and classifications of stream processing languages. On Thursday, Tilmann Rabl from TU Berlin presented a tutorial about the challenges of benchmarking big data systems in general in addition to the specific challenges for benchmarking big stream processing systems. All tutorials have been very informative, interactive and involved very deep technical discussions. On Thursday evening, we had a lively demo session where various participants demonstrated their systems to the audience on parallel round-table interactive discussions. On Wednesday, the participants split into two groups based on common interest in selected subset of the open challengers and problems. The selected 2 topics of the groups were systems and query languages. Thursday schedule was dedicated to working group efforts. Summary about the outcomes of these 2 groups is included in this report. It is expected that work from at least one of the groups to be submitted for publication, and we expect further research publications to result directly from the seminar.

We believe that the most interesting aspect of the seminar was providing the opportunity to freely engage in direct and interactive discussions with solid experts and researchers in various topics of the field with common focused passion and interest. We believe that this is a unique feature for Dagstuhl seminars. We received very positive feedback from the participants and we believe

that most of the participants were excited with the scientific atmosphere at the seminar and reported that the program of the seminar was useful for them. In summary, we consider the organization of this seminar as a success. We are grateful for the Dagstuhl team for providing the opportunity and full support to organize it. The success of this seminar motivated us to plan for future follow-up seminars to continue the discussions on the rapid advancements on the domain and plan for narrower and more focused discussion with concrete outputs for the community.

6.66 Towards Performance Modeling and Performance Prediction across IR/RecSys/NLP

Organizers: Nicola Ferro, Norbert Fuhr, Gregory Grefenstette, Joseph A. Konstan

Seminar No. 17442

Date: October 29 to November 3, 2017 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.7.10.139

Participants: Pablo Castells, Elizabeth M. Daly, Thierry Declerck, Michael D. Ekstrand, Nicola Ferro, Norbert Fuhr, Werner Geyer, Julio Gonzalo, Gregory Grefenstette, Joseph Konstan, Tsvi Kuflik, Krister Lindén, Bernardo Magnini, Jian-Yun Nie, Raffaele Perego, Bracha Shapira, Ian Soboroff, Nava Tintarev, Karin Verspoor, Martijn Willemsen, Justin Zobel

Information systems, which manage, access, extract and process non-structured information, typically deal with vague and implicit information needs, natural language and complex user tasks. Examples of such systems are information retrieval (IR) systems, recommender systems (RecSys), and applications of natural language processing (NLP) such as e.g. machine translation, document classification, sentiment analysis or search engines. The discipline behind these systems differs from other areas of computer science, and other fields of science and engineering in general, due to the lack of models that allow us to predict system performances in a specific operational context and to design systems ahead to achieve a desired level of effectiveness. In the type of information systems we want to look at, we deal with domains characterized by complex algorithms, dependent on many parameters and confronted with uncertainty both in the information to be processed and the needs to be addressed, where the lack of predictive models is somehow bypassed by massive trials of as many combinations as possible.

These approaches relying on massive experimentation, construction of testbeds, and heuristics are neither indefinitely scaled as the complexity of systems and tasks increases nor applicable outside the context of big Internet companies, which still have the resources to cope with them.

The workshop was organized as follows. The first day was devoted to plenary talks focused on providing a general introduction to IR, RecSys, and NLP and on digging into some specific issues in performance modeling and prediction in these three domains. The second day, participants split into three groups – IR, RecSys, and NLP – and explored performance modeling and prediction issues and challenges within each domain; the working groups then reconvened to present the output of their discussion in a plenary session in order to cross-fertilize across disciplines and to identify cross-discipline themes to be further investigated. The third day, participant split into groups which explored these themes – namely measures, performance analysis, documenting and understanding assumptions, application features, and modeling performance – and reported back in plenary sessions to keep all the participants aligned with the ongoing discussions. The fourth and fifth days have been devoted to the drafting of this report and the manifesto originated from the workshop.

This documents reports the overview of the talks given by the participants on the first day. The outcomes of the working groups – both within-discipline themes and cross-discipline themes – as well as the identified research challenges and directions are presented in the Dagstuhl Manifesto corresponding to this Perspectives Workshop [1].

Acknowledgements. We thank Schloss Dagstuhl for hosting us.

References

1 N. Ferro, N. Fuhr, G. Grefenstette, J. A. Konstan, P. Castells, E. M. Daly, T. Declerck, M. D. Ekstrand, W. Geyer, J. Gonzalo, T. Kuflik, K. Lindén, B. Magnini, J.-Y. Nie, R. Perego, B. Shapira, I. Soboroff, N. Tintarev, K. Verspoor, M. C. Willemsen, and J. Zobel. Manifesto from Dagstuhl Perspectives Workshop 17442 – Towards Performance Modeling and Performance Prediction across IR/RecSys/NLP. *Dagstuhl Manifestos, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Germany*, 7(1), 2018.

Fig. 6.16
"Notes from @Dagstuhl seminar on AI driven game design". Twitter post by 17471 Dagstuhl Seminar participant Mirjam P. Eladhari. https://twitter.com/MirjamPE/status/933045410880479232. Photo courtesy of Mirjam P. Eladhari.

6.67 New Challenges in Parallelism

Organizers: Annette Bieniusa, Hans-J. Boehm, Maurice Herlihy, and Erez Petrank

Seminar No. 17451

Date: November 5–10, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.1

Participants: Sarita Adve, Deepthi Devaki Akkoorath, Hagit Attiya, Naama Ben-David, Annette Bieniusa, Hans-J. Boehm, Irina Calciu, Nachshon Cohen, Dave Dice, Sandhya Dwarkadas, Panagiota Fatourou, Pascal Felber, Michal Friedman, Tim Harris, Danny Hendler, Maurice Herlihy, Antony Hosking, Idit Keidar, Alex Kogan, Petr Kuznetsov, William M. Leiserson, Yossi Lev, Victor Luchangco, Virendra Marathe, Maged M. Michael, J. Eliot B. Moss, Madan Musuvathi, Vijay Nagarajan, Erez Petrank, Michael Philippsen, Mirko Rahn, Michael L. Scott, Marc Shapiro, Nir Shavit, Matthew Sinclair, Michael F. Spear, Viktor Vafeiadis, Haris Volos, Heike Wehrheim, Reinhard Wilhelm, Felix Wolf, Sebastian Wolff

Improving the performance and reliability of parallel and concurrent programs is an ongoing topic in multiprocessor software and hardware research. Despite numerous efforts, the semantics of weak memory models remains subtle and fragile. There has not been established a generally accepted way of defining their semantics, and the specifications of programming languages supporting weak memory models with shared accesses are clearly inadequate. In addition, new advances in hardware are adding further complexity. For example, recently, non-volatile memory (NVM) generated a lot of interest in different communities: Hardware designers coming up with instruction sets and layouts of NVM, system designers integrating NVM into the storage stack, programming language designers proposing library interfaces, and theoreticians developing a new theory of persistence under concurrent access and algorithms adapted for persistency.

This Dagstuhl Seminar on "Future Challenges in Parallelism" touched on different aspects on topics in this broad area of research. In this report, we briefly give a summary of the presentations and discussions that took place during the seminar.

Presentations started with an introductory broad overview talk on the non-volatile memory technology. This survey was followed by shorter talks ranging from hardware techniques for efficiently controlling write-back ordering from caches to theoretical foundations and design of specific data structures.

There was agreement that non-volatile memory is likely to become commercially important in the near future, and that it is tempting to exploit it to provide persistence of user data structures. However, there was little agreement on detailed assumptions and direction. The emphasis of the presentations was on manually designed data structures programmed to a near-hardware-level interface. Some participants expressed concerns that this was too low-level and that the community should instead focus on constructs at the level of durable transactions. Transactional semantics are likely to play an important role: When restarting an application from its persisted state, this state must be consistent in order to prevent data corruption and loss. Much of the work presented in the workshop assumed that we will have non-volatile memory combined with visibly volatile caches that require explicit flush operations to persist data. But it was pointed out that the problem could be greatly simplified by either providing sufficient battery backup to ensure that the entire cache is flushed on failure, or providing other hardware support.

Some of the participants discussed the definitions of correctness. On the one hand, the standard definition of durable linearizability is a strong requirement that typically brings a large performance overhead. On the other hand, the weaker buffered linearizability does not compose well. Other participants suggested some hardware modifications that could make the life of the programmer easier. For example, a discussion emerged on whether we could pin a cache line to make sure it is not written back to memory. We also tackled the programmability of systems with non-volatile memories. How difficult should it be to program them? Are application programmers expected to employ it directly or only via dedicated data structures provided in libraries? The experience report of porting the application memcached to non-volatile memory raised a lot of interest with the participants. It turned out that the task was rather difficult due to complex interactions between the different modules in the application, in particular between modules that required persistence and modules that did not. The lack of tools was strongly felt, and the obtained performance was not satisfactory. The conclusion was that applications had better be redesigned from scratch to work with non-volatile memory. The general feeling at the end of the seminar was that we are in the beginning of exciting times for research on non-volatile memories and that the discussions must and will continue.

Memory models formed the second major thread of presentations and discussions, with participants expressing widely

different viewpoints and technical directions. At one extreme, Madan Musuvathi presented evidence that a simple interleaving-based "sequentially consistent" semantics can be provided at reasonable cost, together with an argument that this is a good direction for future programming languages. At the other extreme, Viktor Vafeiadis argued that a weaker "acquire-release" memory model is easier to reason about, an argument that was backed up by model-checking time measurements. Needless to say, this was followed by lively discussion resulting, we believe in at least a more thorough understanding of different perspectives by everyone involved. There were also several brief presentations and extensive discussion on different approaches for addressing the long-standing C++ and Java (among others) out-of-thin-air problem. Current semantics for these languages allow outcomes that are universally accepted as absurd, but which we do not know how to prohibit in any precise way. It is clear that none of the solutions are quite ready to be adopted, but there are encouraging results along several different paths. There is a consensus that this problem makes formal reasoning about programs nearly impossible and that it is a serious obstruction for tool development. There was less consensus about the extent to which it obstructs day-to-day programming efforts.

In conclusion, the seminar inspired discussions and proposed challenging problems to tackle for the research community. As the discussions showed, designing sound and performant parallel systems require the cooperation of researchers on hardware and software level, with both theoretical and practical analyses and evaluations.

6.68 Algorithmic Cheminformatics

Organizers: Jakob L. Andersen, Christoph Flamm, Daniel Merkle, and Peter F. Stadler

Seminar No. 17452

Date: November 5–10, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.28

Participants: Jakob Lykke Andersen, Sarah Berkemer, Mary Ann Blätke, Erick Chastain, Cinzia Di Giusto, Peter Dittrich, Mark Dörr, Rolf Fagerberg, Harold Fellermann, Jérôme Feret, Christoph Flamm, Walter Fontana, David Gilbert, Monika Heiner, Marc Hellmuth, Wim Hordijk, Rojin Kianian, Daniel Merkle, Markus E. Nebel, Nikolai Nøjgaard, Marco Peressotti, Guillermo Restrepo, Heinrich W. Schmidt, D. Eric Smith, Peter F. Stadler, Mirco Tribastone, Klaus-Peter Zauner

Cheminformatics is the application of algorithms, combinatorial approaches, and formal methods from Computer Science to problems in Chemistry. While being formally a very old research field, building the theoretical foundations for Cheminformatics seen from the perspective of state-of-the-art theoretical Computer Science is not at all established research. The second edition of the seminar on "Algorithmic Cheminformatics" brought together researchers working in Chemistry, Cheminformatics, and most importantly the relevant fields in Computer Science related to it. In contrast to the first Dagstuhl meeting in 2014, we specifically focused on the analysis of the behaviour of chemical systems in terms of reaction networks. This includes both networks inferred from experimental data, as well as networks implicitly specified by for example formal grammars. We integrated experts in concurrency theory, in particular using process calculi, Petri nets, and related formal approaches. State-of-the-art results in these fields are hardly used to infer qualitative and/or quantitative properties of chemical reaction systems, which are highly concurrent systems by nature. Most current modeling approaches in chemistry are either very abstract and aimed at formal algebraic properties of reaction networks, or use precise modeling on a very fine grained level such as the quantum mechanical one where computational costs prevent handling of more than a few molecules. In this seminar we therefore seeked to advance discrete modeling approaches for Systems Chemistry. In addition to bringing together the experts in the respective fields from Computer Science, we also invited wet-lab chemists in order to cross-fertilize the fields and generate mutually beneficial activities.

6

6.69 Connecting Visualization and Data Management Research

Organizers: Remco Chang, Jean-Daniel Fekete, Juliana Freire, and Carlos E. Scheidegger

Seminar No. 17461

Date: November 12–17, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.46

Participants: Sihem Amer-Yahia, Leilani Battle, Carsten Binnig, Tiziana Catarci, Remco Chang, Surajit Chaudhuri, Stephan Diehl, Harish Doraiswamy, Steven M. Drucker, Jason Dykes, Jean-Daniel Fekete, Danyel Fisher, Juliana Freire, Michael Gleicher, Hans Hagen, Gerhard Heyer, Heike Hofmann, Daniel A. Keim, Tim Kraska, Heike Leitte, Zhicheng Liu, Volker Markl, Alexandra Meliou, Torsten Möller, Dominik Moritz, Hannes Mühleisen, Arnab Nandi, Behrooz Omidvar-Tehrani, Themis Palpanas, Carlos E. Scheidegger, Gerik Scheuermann, Michael Sedlmair, Thibault Sellam, Juan Soto, Richard Wesley, Wesley J. Willett, Eugene Wu, Yifan Wu

What prevents analysts from acquiring wisdom from data sources? To use data, to better understand the world and act upon it, we need to understand both the computational and the human-centric aspects of data-intensive work. In this Dagstuhl Seminar, we sought to establish the foundations for the next generation of data management and visualization systems by bringing together these two largely independent communities. While exploratory data analysis (EDA) has been a pillar of data science for decades, maintaining interactivity during EDA has become difficult, as the data size and complexity continue to grow. Modern statistical systems often assume that all data need to fit into memory in order to support interactivity. However, when faced with a large amount of data, few techniques can support EDA fluidly. During this process, interactivity is critical: if each operation takes hours or even minutes to finish, analysts lose track of their thought process. Bad analyses cause bad interpretations, bad actions and bad policies.

As data scale and complexity increases, the novel solutions that will ultimately enable interactive, large-scale EDA will have to come from truly interdisciplinary and international work. Today, database systems can store and query massive amounts of data, including methods for distributed, streaming and approximate computation. Data mining techniques provide ways to discover unexpected patterns and to automate and scale well-defined analysis procedures. Recent systems research has looked at how to develop novel database systems architectures to support the iterative, optimization-oriented workloads of data-intensive algorithms. Of course, both the inputs and outputs of these systems are ultimately driven by people, in support of analysis tasks. The life-cycle of data involves an iterative, interactive process of determining which questions to ask, the data to analyze, appropriate features and models, and interpreting results. In order to achieve better analysis outcomes, data processing systems require improved interfaces that account for the strengths and limitations of human perception and cognition. Meanwhile, to keep up with the rising tide of data, interactive visualization tools need to integrate more techniques from databases and machine learning.

This Dagstuhl seminar brought together researchers from the two communities (visualization and databases) to establish a research agenda towards the development of next generation data management and interactive visualization systems. In a short amount of time, the two communities learned from each other, identified the strengths and weaknesses of the latest techniques from both fields, and together developed a "state of the art" report on the open challenges that require the collaboration of the two communities. This report documents the outcome of this collaborative effort by all the participants.

6.70 A Shared Challenge in Behavioural Specification

Organizers: Klaus Havelund, Martin Leucker, Giles Reger, and Volker Stolz

Seminar No. 17462

Date: November 12–15, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.59

Participants: Wolfgang Ahrendt, Cyrille Artho, Domenico Bianculli, Borzoo Bonakdarpour, Stijn de Gouw, Cindy Eisner, Bernd Finkbeiner, Adrian Francalanza, Sylvain Hallé, Martin Leucker, Zhiming Liu, Keiko Nakata, Dejan Nickovic, Gordon Pace, Nicolas Rapin, Giles Reger, Kristin Yvonne Rozier, César Sánchez, Torben Scheffel, Gerardo Schneider, Julien Signoles, Volker Stolz, Hazem Torfah, Dmitriy Traytel, Nobuko Yoshida

This seminar dealt with the issue of behavioural specification from the viewpoint of runtime verification. Runtime verification (RV) as a field is broadly defined as focusing on processing execution traces (output of an observed system) for verification and validation purposes. Of particular interest is the problem of verifying that a sequence of events, a trace, satisfies a temporal property, formulated in a suitable formalism. Examples of such formalisms include state machines, regular expressions, temporal logics, context-free grammars, variations of the mu-calculus, rule systems, stream processing systems, and process algebras. Of special interest is how to specify data-rich systems, where events themselves carry data. Applications cover such domains as security monitoring and safety monitoring.

Such techniques are characterised by highly expressive languages for specifying behaviour, enabled by the concreteness of dealing directly with single runtime traces, which makes the verification problem tractable. However, this permitted expressiveness has also led to a divergence in such languages. The aim of this Dagstuhl Seminar was to shed light on the similarities and differences between these different formalisms, and specifically, suggest directions for future collaboration and research. This effort can potentially lead to an attempt to standardize an RV formalism.

The seminar included a mixture of tool developers, theoreticians, and industry experts and the above aim was addressed by two main activities.

The first activity was that each tool developer was asked to produce a brief summary of their specification language in the form of a set of short examples. These were then presented as talks during the Seminar, alongside other general contributed talks on issues surrounding behavioural specification. The examples were uploaded to a shared repository (which will be available via runtime-verification.org) and eleven participants added their tool descriptions and examples to this repository, producing a lasting resource from the seminar.

The second activity was carried out through eight working groups formed during the Seminar to discuss topics raised by the talks. The results of this working groups are detailed in this report. We take this opportunity to detail the topics (in the form of questions) proposed during the seminar that were not chosen for discussion in working groups:

- *Where should we get specifications from?* This question addressed both the issue of designing specification languages that can be usable by engineers but also the trending topic of inferring specifications from various artifacts and how specification languages can support this.
- *How can we measure specification quality?* What is a good specification, or when is one specification better than another? This might be related to coverage of the system being specified, or might be about interpretability or some other measure of usability.
- *How do we ensure our specification language is not broken?* This question was inspired by the experience of one speaker with developing the industrial-strength PSL language and the issues surrounding getting it right.
- *How can we balance different levels of abstraction (e.g. local and global behaviour) in a specification?* It was noted that specification languages are often closely associated with specifications at a certain level of abstraction. Is this an inherent restriction or a positive feature? Should we build specification languages with a certain level of abstraction in mind?
- *How do we unify the different uses of a specification?* This was inspired by the observation that a specification may be used to explain behaviour, check behaviour, or synthesize behaviour, and different presentations may be preferred in these different contexts.

This seminar was the first time the runtime verification community has reflected on the broad issue of specification and has fed into further developments including new perspectives for the international runtime verification competition, a proposed shared challenge involving the NASA core flight system, and the first informal survey and categorisation of actively developed runtime verification tools.

6.71 Artificial and Computational Intelligence in Games: AI-Driven Game Design

Organizers: Pieter Spronck, Elisabeth André, Michael Cook, and Mike Preuß

Seminar No. 17471

Date: November 19–24, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.86

Participants: Elisabeth André, Dan Ashlock, Sander C.J. Bakkes, Rafael Bidarra, Cameron Browne, Alex J. Champandard, Simon Colton, Michael Cook, Peter I. Cowling, Steve Dahlskog, Mirjam P. Eladhari, Reynald Francois, Matthew J. Guzdial, Amy K. Hoover, Pier Luca Lanzi, Marc Erich Latoschik, Antonios Liapis, Jialin Liu, Simon M. Lucas, Joshua Allen McCoy, Alexander Nareyek, Mark J. Nelson, Ana Paiva, Diego Perez Liebana, Mike Preuß, Sebastian Risi, Spyridon Samothrakis, Jacob Schrum, Emily Short, Adam M. Smith, Gillian Smith, Pieter Spronck, Anne Sullivan, Shoshannah Tekofsky, Tommy Thompson, David Thue, Julian Togelius, Michael Treanor, Vanessa Volz, Georgios N. Yannakakis, R. Michael Young, G. Michael Youngblood, Jichen Zhu

The video game industry has been developing rapidly in the past decade. Whereas ten years ago video games were almost exclusively aimed at entertainment, nowadays they are used in a variety of places in everyday life. All kinds of organizations now use video games for simulation and training. Educational institutes use video games to enrich and replace parts of courses. Governmental and health care agencies use video games to educate people and stimulate them to lead more productive lives. On top of that, the entertainment-focused games industry continues to grow and is a major industry both culturally and financially.

Two parallel developments can be observed in the games industry. On the one hand, the high-profile entertainment games ("triple-A games") see a steady increase of time and financial resources invested in their development, to keep up with technological advances and to be able to compete in a tough market. On the other hand, the number of smaller, special-purpose games in development (including so-called "serious games") increases dramatically too, in particular in research, training, and education. Moreover, as the pervasiveness of video games increases, so does the the number of people involved in creating them. The job of creating games is no longer limited to specialist programmers and artists. Instead, those who need to use the games become heavily involved in their creation.

The serious-games domain poses additional challenges to game development beyond all the challenges already posed by games for entertainment, namely the need for a strong relation with the "real world". Serious games often have a purpose in training, which entails that the game worlds must be a realistic depiction of the actual environment in which user functions, in particular where "behaviors" are concerned. The big data revolution means that huge quantities of data about the real world are becoming available along with the means of processing them, which may offer the possibility to automatically construct games on the basis of such data. This is a particularly enticing notion, given the financial constraints for constructing serious games, which means that professional content producers might not be available.

Furthermore, new computer games are expected to much better incorporate the different needs of a wide variety of customers, to provide more alternative modes, solution paths, incentives, emotional states, and difficulty levels. Game design, and especially balancing, must take this into account. However, this increases the complexity of design and production considerably, such that AI-based tools that can assist the human developer or even partly automatize processes are more desired than ever.

Summarizing, we note the following four trends in modern game development:

1. Technological advances have lead to an increased challenge in developing modern video games, even for expert game developers
2. There is an increased need for non-experts to be able to design and develop games
3. There is an increased need for realism in the virtual world behaviors, in particular in the area of serious games
4. A greater variety of players and a better availability of data about players leads to the need for more variable and better customizable games, which require a more complex development process.

A solution for each of these issues can be found in the application of artificial-intelligence (AI) techniques to drive the design and development of games. From the perspective of AI-driven game design, AI supports or even takes over the role of the human game developer in creating particular elements of a game, and even complete new games.

While the game industry tends to use a small selection of well-known algorithms to generate elements of game worlds (in

particular where graphics and animation are concerned), the use of AI to create new environments, new behaviors of virtual characters, new narratives, new game rules, or even new gameplay mechanisms is at present limited to a very small number of researchers. We see it, however, as the main direction in which innovation in game design and development can be found.

AI-driven game design sees applications in the design of virtual worlds, virtual characters, narratives, and game mechanics. Moreover, it can be used to assist in the human design process, and to adapt games automatically after publication. Finally, it can support the automated analysis of generated game elements. Each of these topics is a research domain by itself, which requires an interdisciplinary approach which borrows from computer science, psychology, cognitive science, and even the creative arts. A common ground is found in artificial intelligence techniques, in particular machine learning.

For this seminar, we brought together computer scientists and creative experts with the common goals of gaining a deeper understanding of various aspects of games, and of further improving games, in particular by using AI-techniques used to generate games and game elements. The goal was to look beyond what is currently possible and in use, and take steps towards the future of AI-driven game development. Besides theoretical discussions, part of the seminar was spent on trying to achieve first practical results.

Reports on the discussions and results achieved are found on the following pages. All in all, the organizers and participants deemed the seminar a great success, and are eager to continue into some of the directions that were focused on during the week at Schloss Dagstuhl.

6.72 Addressing the Computational Challenges of Personalized Medicine

Organizers: Niko Beerenwinkel, Holger Fröhlich, and Susan A. Murphy

Seminar No. 17472

Date: November 19–22, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.130

Participants: Rudi Balling, Niko Beerenwinkel, Holger Fröhlich, Oliver Kohlbacher, Santosh Kumar, Thomas Lengauer, Marloes Maathuis, Yves Moreau, Susan A. Murphy, Teresa Przytycka, Michael Rebhan, Hannes Röst, Andreas Schuppert, Matthias Schwab, Rainer Spang, Daniel Stekhoven, Jimeng Sun, Andreas Weber, Daniel Ziemek, Blaz Zupan

Personalized medicine (PM) is understood as a non-traditional medical approach, in which patients are stratified based on their disease subtype, disease risk, disease prognosis or treatment response using specialized diagnostic tests. High promises for the whole health care sector are associated with PM, and correspondingly the topic has received a lot of attention during the last years. PM is tightly connected to and dependent on computational sciences (computer science, mathematical modeling, computational statistics, bioinformatics). Currently, shortcomings of computational methodology constitute an important bottleneck for PM, which hinders full realization.

The goal of the planned seminar was to bring together an international and interdisciplinary group of experts in different computational science disciplines in order to discuss, how some of the major existing computational challenges could be better addressed in the future, namely:

1. how to enhance prediction performance of computational models for PM
2. how to improve their interpretability
3. how to validate and implement them in practice

The seminar joined together expertise that is usually scattered across different disciplines. The seminar had a strict focus on computational methodology, but few selected non-computational scientists closed the gap to the application field.

Fig. 6.17
"They might not celebrate Thanksgiving in Germany, but it turns ouy we can import it. #dagstuhl". Twitter post by 17471 Dagstuhl Seminar participant Anne Sullivan.
https://twitter.com/anneandkita/status/933459571510784000. Photo courtesy of Anne Sullivan.

6.73 Reliable Computation and Complexity on the Reals

Organizers: Norbert T. Müller, Siegfried M. Rump, Klaus Weihrauch, and Martin Ziegler

Seminar No. 17481

Date: November 26 to December 1, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.11.142

Participants: Andrej Bauer, Henning Behnke, Jens Blanck, Franz Braße, Florian Bünger, Pieter Collins, George F. Corliss, Tibor Csendes, Eva Darulova, Hubert Glesener, Daniel Graça, Stef Graillat, Peter Hertling, Fabian Immler, Luc Jaulin, Akitoshi Kawamura, Sunyoung Kim, Michal Konecny, Margarita Korovina, Vladik Kreinovich, Ivo List, Wolfram Luther, Fritz Mayer-Lindenberg, Atsushi Minamihata, Norbert T. Müller, Mitsuhiro T. Nakao, Eike Neumann, Katsuhisa Ozaki, Sewon Park, Michael Plum, Amaury Pouly, Robert Rettinger, Fabrice Rouiller, Siegfried M. Rump, Michael Sagraloff, Matthias Schröder, Svetlana Selivanova, Florian Steinberg, Akitoshi Takayasu, Holger Thies, Warwick Tucker, Joris van der Hoeven, Klaus Weihrauch, Nobito Yamamoto, Yuuka Yanagisawa, Chee K. Yap, Martin Ziegler, Paul Zimmermann

The seminar was a meeting between two groups of researchers working in the related areas of reliable computing and of computational complexity on real numbers. While the first area originates in numerical analysis, the second area goes back to the roots of computer science and and computability.

Reliable computations aims to produce correct answers to numerical problems with mathematical rigor. This includes to prove that the problem is solvable and to compute mathematically correct error bounds for the solution. Reliable numerical computations solely use floating-point arithmetic to take advantage of the tremendous speed. Naturally that poses limits on the problems which can be solved, in particular the condition number. However, in contrast to purely numerical methods, no false answers are possible: Either a true error bound is computed or, a corresponding error message is given. There is a history of reliable numerical computations. In the early days, interval arithmetic was often used in a rather naive way. Still the computed results were correct, however, often wide or no bounds at all were computed. Meanwhile it is well understood how to derive effective methods for reliable numerical computations, avoiding wide bounds and pushing the set of solved problems to the limit of that of purely numerical algorithms. A number of interesting and hard mathematical problems have been solved using reliable numerical computations. This includes the famous Kepler conjecture, the existence of mutually distinct solutions to certain partial differential equations, and more. Needless to say that solving a mathematical problem requires rigorous solutions of all particular problems.

Computable analysis is a branch of computability theory studying those functions on the real numbers and related structures which can be computed by machines such as digital computers. The increasing demand for reliable software in scientific computation and engineering requires a sound and broad foundation not only of the analytical/ numerical but also of the computational aspects of real number computation. The branch of computable analysis based on the definition by Grzegorczyk and Lacombe of computable real functions (TTE, "Type 2 Theory of Effectivity") has turned out to be particularly useful for investigating computability on uncountable sets. As a central concept computability appears as a specialization of continuity. Meanwhile computability of numerous analytic problems has been investigated (from basic analysis, functional analysis, ordinary and partial differential equations, analytic functions, measure theory, dynamical systems etc.). All these examples demonstrate the usefulness of the concept.

Once a problem has been shown computable, a natural next question asks for the computational efficiency of such a solution. This is where real analysis meets (discrete) complexity theory with notions of runtime and memory/space: asymptotically with respect to $n \to \infty$ for approximating the output up to absolute error 2^{-n}. The famous Bailey-Borwein-Plouffe method for instance permits to compute billions of digits of transcendental within minutes; while Bloch's constant, although proven computable, is still not known up to error 2^{-5}. In fact the distinction between polynomial and exponential time, in the discrete realm gauged for instance by complexity classes P, NP, #P, and PSPACE, re-emerges in the real case: The bit-cost of computing the maximum of an arbitrary fixed smooth (i.e. infinitely often differentiable) polynomial-time computable $f : [0;1] \to [0;1]$ has been shown to correspond to P-vs-NP; that of Riemann integration to #P; and that of solving an ordinary differential equation to PSPACE. On analytic functions on the other hand these operations map polynomial-time computable instances back to polynomial-time computable results.

For practical purposes and in the spirit of "algorithm engineering", the asymptotic results from complexity theory have to be refined by considering the efficiency of actual implementations. Corresponding software libraries are usually called "exact real

arithmetic" (ERA) and implement real numbers in the sense of TTE. ERA implementations exist in many languages, like C, C++ JAVA, Haskell or OCaml. Internally, ERA has to perform operations on infinite data like $\{0,1\}^\omega$. The user interface, however, hides the details and offers operations and functions on "exact" real numbers. In consequence, users do not need to care about aspects like rounding or truncation errors or the specification of precisions. Instead, they can concentrate on the mathematical part of the problem under consideration. As computable real functions have to be continuous, it is impossible to implement some widely used real functions (like testing on equality). In consequence, ERA cannot simply copy the double precision interface one-to-one, but needs to go its own ways. Additionally, for the reason of efficiency the representations used in TTE have to be carefully revised. The resulting speed is comparable to the use of multiple precision floating point numbers, but now without any need for manual precision control.

6.74 Computational Metabolomics: Identification, Interpretation, Imaging

Organizers: Theodore Alexandrov, Sebastian Böcker, Pieter Dorrestein, and Emma Schymanski

Seminar No. 17491

Date: December 3–8, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.12.1

Participants: Hayley Abbiss, Theodore Alexandrov, Manor Askenazi, Eric Bach, Ruth Birner-Grünberger, Sebastian Böcker, Corey Broeckling, Christoph Büschl, Ricardo Da Silva, Pieter Dorrestein, Edward M. Driggers, Kai Dührkop, Madeleine Ernst, P. Lee Ferguson, Nils Hoffmann, Julijana Ivanisevic, Fabien Jourdan, Alexander Kerner, Oliver Kohlbacher, Martin Krauss, Martin Loos, Marcus Ludwig, Marnix Medema, Kris Morreel, Rolf Müller, Steffen Neumann, Louis-Felix Nothias-Scaglia, Andrew Palmer, Alan Race, Stacey Reinke, Simon Rogers, Juho Rousu, Sarah Scharfenberg, Jennifer Schollee, Emma Schymanski, Jan Stanstrup, Michael Andrej Stravs, Raf Van de Plas, Justin van der Hooft, Kirill Veselkov, Ron Wehrens, David Wishart, Michael Anton Witting, Nicola Zamboni

Metabolomics is the study of metabolites (the small molecules involved in metabolism) in living cells, cell populations, organisms or communities. Metabolites are key players in almost all biological processes, play various functional roles providing energy, building blocks, signaling, communication, and defense and serve as clinical biomarkers for detecting medical conditions such as cancer. Small molecule drugs (many of which are derived from metabolites) account for 90 % of prescribed therapeutics. Complete understanding of biological systems requires detecting and interpreting the metabolome in time and space.

Mass spectrometry is the predominant analytical technique for detecting and identifying metabolites and other small molecules in high–throughput experiments. Huge technological advances in mass spectrometry and experimental workflows during the last decade enabled novel investigations of biological systems on the metabolite level. Research into computational workflows, the simulation of tandem mass spectra, compound identification and molecular networking have helped disentangle the vast amount of information that mass spectrometry provides. Spatial metabolomics on different spatial scales from single cells to organs and organisms has posed data analysis challenges, in particular due to an unprecedented data volume generated that grows quadratically with the increase of spatial resolution.

Continued improvements to instruments, resolution, ionization and acquisition techniques mean that metabolomics mass spectrometry experiments can generate massive amounts of data, and the field is evolving into a "big data" science. This is particularly the case for imaging mass spectrometry, where a single dataset can easily be many gigabytes or even terabytes in size. Despite this dramatic increase in data, much of the data analysis in metabolomics is still performed manually and requires expert knowledge as well as the collation of data from a plethora of sources. Novel computational methods are required to exploit spectral and, in the case of imaging, also spatial information from the data, while remaining efficient enough to process tens to hundreds of gigabytes of data.

Dagstuhl Seminar 17491 on Computational Metabolomics: Identification, Interpretation, Imaging built on the success of the first Computational Metabolomics Dagstuhl Seminar (15492) in 2015. A number of topics overlapped with the 2015 seminar, while the focus on imaging introduced new perspectives, participants and topics. In contrast to the first seminar, 17491 was a large seminar, with 45 very active participants and a large portion of young scientists. From the first hours of the seminar, effort was made to integrate these young scientists in the discussions and presentations and this paid off leading to lively discussions involving all participants. Many participants were new to Dagstuhl and the concept of Dagstuhl seminars, which led to a seminar that was a combination of being semi-structured and spontaneous. Very positive feedback was received from all during a comprehensive feedback session before lunch on Friday, including constructive ideas for a new focus for a possible new seminar in 2019.

On the scientific side, the seminar covered numerous topics which were found to be most relevant for the computational analysis of mass spectrometry data, and ranged from the "dark matter in metabolomics" to "integrating spatial and conventional metabolomics"; see the full report for a comprehensive description.

The seminar has fully achieved its key goals: to foster the exchange of ideas between the experimental and computational communities; to expose the novel computational developments and challenges; and, to establish collaborations to address grand and priority challenges by bridging the best available data with the best methods.

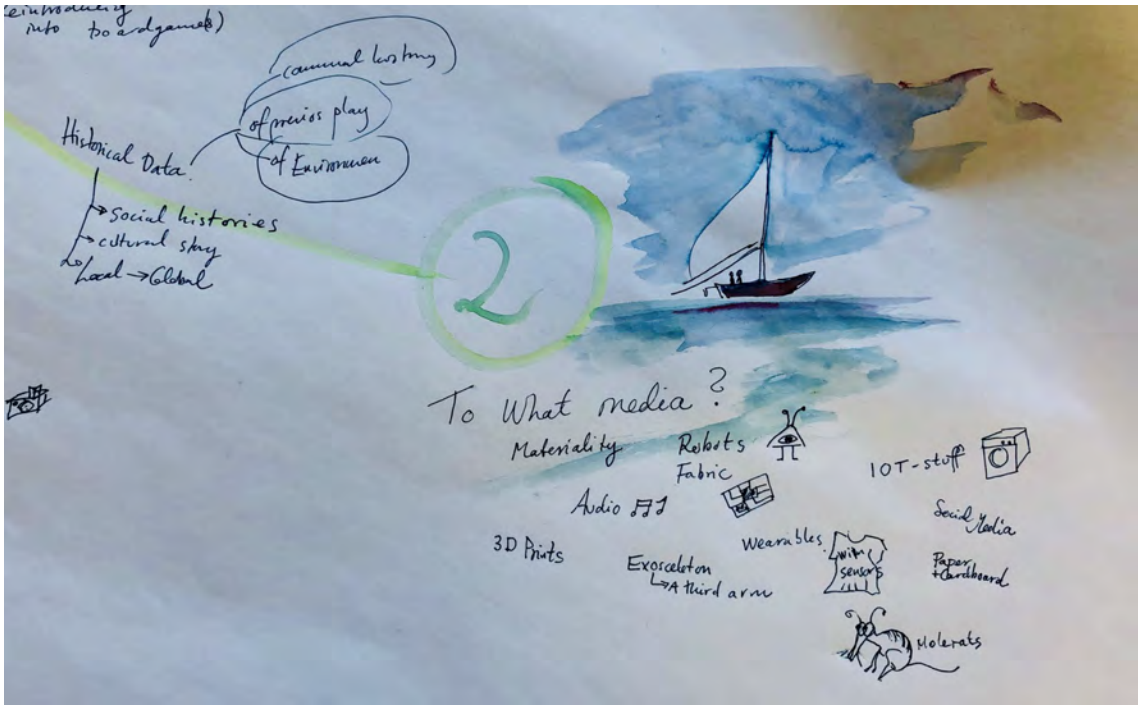

Fig. 6.18
"Notes from @Dagstuhl day 3 part 1 ". Twitter post by 17471 Dagstuhl Seminar participant Mirjam P. Eladhari. https://twitter.com/MirjamPE/status/933300298826870784. Photo courtesy of Mirjam P. Eladhari.

6.75 Multi-Level Modelling

Organizers: João Paulo A. Almeida, Ulrich Frank, and Thomas Kühne

Seminar No. 17492

Date: December 3–8, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.12.18

Participants: João Paulo Almeida, Mira Balaban, Tony Clark, Victorio Albani de Carvalho, Dirk Draheim, Philipp Martin Fischer, Ulrich Frank, Cesar Gonzalez-Perez, Georg Grossmann, Ta'id Holmes, Manfred Jeusfeld, Agnes Koschmider, Anne Koziolek, Thomas Kühne, Vinay Kulkarni, Wendy MacCaull, Bernd Neumayr, Chris Partridge, Iris Reinhartz-Berger, Michael Schrefl, Matt Selway, Maarten Steen, Friedrich Steimann, Manuel Wimmer, Dustin Wüest

Multi-Level Modeling (MLM), i.e., the explicit exploitation of multiple levels of classification when modeling, represents a significant extension to the traditional two-level object-oriented paradigm with the potential to dramatically improve upon the utility, reliability and complexity level of models. It therefore has benefits in many domains of modeling such as software engineering, process modeling and enterprise modeling. Research into multi-level modeling has increased significantly over the last few years, manifesting itself in lively debates in the literature, four international workshops (MULTI 2014–2017), a published journal theme issue (SoSyM), a special issue for the EMISA journal (in preparation), and a growing number of tools and languages. While the enthusiasm around MLM provides momentum to this promising research area, the recent speed of growth and the focus on exploring new language features has raised some challenges, including the following:

growing diversity of approaches On the one hand diversity is welcome in order to spawn a competition of ideas but on the other hand it can slow down the approach's growth and industry adoption unless steps are taken to consolidate ideas and bundle resources.

lack of integration with related disciplines In particular, ontology engineering has overlapping application areas and could form a powerful synergy with MLM due to its complementary strengths and weaknesses. Moreover, areas such as logic, philosophy, and linguistics are also highly relevant for the further advancement of MLM.

neglect of real world applications It is natural to initially focus on core principles when developing a new approach, but at some point it becomes important to make the transition into industrial practice in order to validate claims about the utility and need for MLM, and to promote the uptake of MLM in new domains and industries.

Goals

In order to address the aforementioned challenges the seminar brought together researchers and industry practitioners from the areas of conceptual modeling, ontologies, and formal foundations. In particular, to further the coherence of future research into multi-level modeling, we aimed at

- having some consolidating discussion on terminology and scope;
- strengthening (formal) foundations; and
- identifying objective criteria for comparing competing approaches, e.g., by developing respective benchmarks in cooperation with modelers from industry.

Working Groups

A talk on "What is Multi-Level Modeling?" by Thomas Kühne (cf. talk abstract) set the stage for the terminology discussion and presented results from a survey that the organizers ran prior to the seminar. The survey results were in good agreement with the ideas the talk put forward in terms of what core multi-level modeling concepts are, such as "multiple levels of abstraction", "classification as the core abstraction principle", "modeling the real world" (as opposed to engineering languages). The survey design carefully avoided introducing bias, hence there were no multiple-choice questions on subjects like this one. The lack of answer standardization required a manual allocation into answer categories such as the aforementioned ones, but only clear cut cases were counted. Overall, there was very little controversy over what multi-level modeling constitutes. Furthermore, the survey nicely confirmed that the initial seminar goals were congruent with what most participants found to be interesting and important work in the area of multi-level modeling.

Two further talks were aimed at supporting the formation of working groups:

1. A talk on foundations and ontologies by João Paulo A. Almeida contrasted ontology engineering to language engineering, asked which questions should be addressed by a foundation, and explored some answers.
2. A talk on applications by Ulrich Frank elaborated on challenges for multi-level modeling in practical applications.

Subsequently four working groups were established by identifying the themes that both aligned with the original workshop goals and garnered the highest interest among participants.

A working group on "Foundations" formed and decided to start on investigating which ontological commitments and metaphysical choices may be required or useful for a foundation of multi-level modeling (cf. "Foundations" group report, Section 4.1 of the full report).

A working group on "Applications" set out to identify promising application domains for MLM, find common properties for such application domains, identify anticipated benefits of MLM, and determine evaluation criteria for MLM methods (cf. "Applications" group report, Section 4.2 of the full report).

A further working group on "Dynamic Aspects" focused on a sub-area of enterprise modeling, i.e., modeling process-related and/or dynamic behavior aspects (cf. "Dynamic Aspects" group report, Section 4.3 of the full report).

We strove to both address the originally planned goals of the seminar but also to allow new goals to be formed, based on the final composition of the participants. As a result, the group formation process yielded one more group that focused on transformations in the context of multi-level modeling (cf. "Transformations" group report, Section 4.4 of the full report).

Due to the overlap between foundation work and the area of "integration with ontologies", a group dedicated exclusively to exploring synergies between MLM and ontology engineering did not emerge. We are hopeful that the working group on "Foundations" will explore more of the synergy aspect in future collaborations.

Industry Focus

As closing the gap between academia and industry with respect to multi-level modeling was a primary goal of the seminar, we had a total of six talks by industry representatives. These talks gave the speakers an opportunity to explain actual needs, set challenges, comment on utility, etc., plus allowed the audience to inquire about hurdles for adoption, etc. Please see the industry talk abstracts included in this report for further details.

We, the organizers, are extremely grateful to the staff of Dagstuhl for providing a perfect seminar venue and to the participants who not only made this seminar a success but also provided a wealth of generous positive feedback.

The organizers, Dagstuhl, 2018
João Paulo A. Almeida
Colin Atkinson
Ulrich Frank
Thomas Kühne

6.76 Testing and Verification of Compilers

Organizers: Junjie Chen, Alastair F. Donaldson, Andreas Zeller, and Hongyu Zhang

Seminar No. 17502

Date: December 10–13, 2017 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.7.12.50

Participants: Edward E. Aftandilian, Marcel Beemster, Junjie Chen, Nathan Chong, Eric Eide, Hugues Evrard, Dan Hao, John Hughes, Dan Iorga, Julia Lawall, Daniel Lehmann, Thibaut Lutz, David MacIver, Jessica Paquette, David J. Pearce, Michael Pradel, John Regehr, Raimondas Sasnauskas, Marija Selakovic, Gerson Sunyé, Nikolai Tillmann, Yingfei Xiong, Francesco Zappa Nardelli, Andreas Zeller, Hongyu Zhang

The full report documents the Dagstuhl Seminar 17502 "Testing and Verification of Compilers".

Compilers underpin all software development, but bugs in them can be particularly hard to notice if they result in "silent failure", where a program appears to work but is subtly miscompiled. Thus a compiled program may behave erroneously even when the source form of it appears entirely correct.

Despite the common wisdom that "it is never the compiler's fault", bugs in compilers are in fact relatively common, and finding them is a challenging and active area of research.

This seminar brought together researchers in that area with a broader group of researchers and practitioners in software testing and verification, and in compiler development itself, to share their experiences and discuss the open questions and challenges that the field presents. The goal was to brainstorm new ideas for how to approach these challenges, and to help foster longer-term collaborations between the participants.

The seminar involved a number of talks from participants about their particular areas of work and research, followed by working groups where various specific challenges were discussed. It then concluded with an open panel session on the challenges and concepts of compiler testing and verification.

The full report presents the collection of abstracts associated with the participant presentations, followed by notes summarising each discussion session and the concluding panel, which we provide as a resource for researchers who are interested in understanding the state of the art and open problems in this field.

7 Öffentlichkeitsarbeit
Public Relations and Outreach

7.1 Pressemitteilungen und Medienarbeit / 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[53] 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. 1 400 Abonnenten verbreitet. Zunehmend nutzen aber auch Seminarteilnehmer den Dienst, um ihre Eindrücke vom Seminar mitzuteilen. Darüber hinaus werden über den Twitter-Nutzer @dblp_org Informationen über die Bibliographiedatenbank dblp an ca. 550 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ gepflegt (derzeit über 630 Mitglieder), mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Weiterhin werden dort interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben.

Regular press releases showcase and disseminate information about current computer science topics in a comprehensible manner and clarify the concept behind Schloss Dagstuhl. Press releases and media reports that come to the center's attention are available on the Schloss Dagstuhl website[53].

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 1,400 followers, but is also increasingly used by Dagstuhl Seminar participants to share their impressions. Additionally, information about the dblp computer science bibliography is sent using the Twitter account @dblp_org, having about 550 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with more than 630 members), which supports the networking of participants in Dagstuhl Seminars. Additionally, interesting news about Schloss Dagstuhl are announced there.

7.2 Fortbildung / 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

[53] https://www.dagstuhl.de/about-dagstuhl/press/

zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch können seit 2011 bis zu fünf Lehrer aus anderen Bundesländern teilnehmen. Mehr Informationen zur Lehrerfortbildung 2017 gibt es auf der Webseite der Veranstaltung[54].

participate. Details on the workshop in 2017 are available at the event webpage[54].

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.

Den ersten Vortrag in 2017 hielt der Präsident der Leibniz-Gemeinschaft, Prof. Matthias Kleiner, am 19. Mai 2017. Unter dem Titel „Die Vielfalt der Wissenschaften“ gab er einen breiten Einblick in die vielfältigen Zukunftsthemen, denen sich die Leibniz-Forscherinnen und -Forscher widmen, seien es der Erhalt der Artenvielfalt, das Erkunden der Milchstraße und neuer Planeten, oder die Entwicklung individualisierter Medizin. Kleiner betonte dabei die Wichtigkeit, relevante Themen dabei auch öffentlich und konstruktiv-kontrovers zu debattieren, damit gesellschaftliches und politisches Handeln von wissenschaftlichen Erkenntnissen profitieren kann.

Am 24. November 2017 führte Prof. Michael Backes vom CISPA Helmholtz Center i.G. die Reihe mit einem Vortrag zum Thema „Autonome Systeme der Zukunft – Sicherheit und Datenschutz als Standortvorteil“ fort. An greifbaren Beispielen wie selbstfahrenden Autos, dem smarten Eigenheim oder autonomen Industrierobotern zeigt Prof. Backes auf, wie nachhaltgig sich unsere Gesellschaft in den kommenden Jahren verändern wird — und welche Chancen und Risiken in diesem Wandel liegen.

Die Vorträge fanden regen Anklang und lösten angeregte Diskussionen aus, und werden im kommenden Jahr gewiss fortgesetzt werden.

7.3 *“Dagstuhler Gespräche”*

In order to open its doors a bit further for the general public and the local region, Schloss Dagstuhl, together with the town of Wadern, initiated a new series of events: the *Dagstuhler Gespräche* (“Dagstuhl conversations”). The interested public will be introduced to a broad spectrum of topics from computer science, as well as to practical applications of those topics in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

The first talk of 2017 was given by the President of the Leibniz Association, Prof. Matthias Kleiner, on May 19. Under the title “Die Vielfalt der Wissenschaften” (The Variety of Sciences), the talk gave an insight into the various research topics Leibniz researchers are working on, be it the conservation of species, exploring the Galaxy and new planets, or developing individualized medicine. Kleiner emphasized the importance of discussing relevant topics publicly, constructively, and controversially, so that social and political action can benefit from scientific knowledge and research discoveries.

On November 24, 2017, Prof. Michael Backes of the CISPA Helmholtz Center i.G. continued the series with a talk on the topic of “Autonome Systeme der Zukunft – Sicherheit und Datenschutz als Standortvorteil” (Autonomous Systems of the Future – Security and Data Protection as a Locational Advantage). Backes used tangible examples, such as autonomous driving, smart homes, or autonomous industrial robots, to illustrate the lasting societal changes of the coming years — as well as the opportunities and risks that come with this change.

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

[54] https://www.dagstuhl.de/17503

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.

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 unserem 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 Mittagsbuffet sowie ein Menü am Abend an. Montags und freitags wird aus logistischen Gründen auch am Mittag ein

Whereas the two main lecture halls are equipped with several blackboards, whiteboards are provided in the other rooms.

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 the apprentice 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 hot drinks are served outside the lecture halls. During the longer coffee break in the afternoon, hot drinks

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. Schloss Dagstuhl kommt für die Unterkunft und Verpflegung der Kinder auf. Wenn statt Inanspruchnahme der Kinderbetreuung von Schloss Dagstuhl eine Betreuungsperson mitreist, hat diese ebenfalls freien Aufenthalt.

Dagstuhls Angebot der Kinderbetreuung für Eltern wird weiterhin gut genutzt. Im Jahr 2017 wurden 14 Kinder durch eine Tagesmutter und 22 weitere durch Verwandte betreut. Insgesamt beherbergte Schloss Dagstuhl 36 Kinder von Teilnehmern an 28 Veranstaltungen während 24 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.

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 for the children. If an own nanny takes care for the children instead of Dagstuhls childcare service, also the cost for the accompanying person for room and board are absorbed by Dagstuhl.

Guests make good use of Dagstuhl's childcare offer for parents. In 2017, Dagstuhl hosted 36 children, 14 of whom were cared for by a nanny on site and 22 by relatives. Participants of 28 events in 24 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.

8.2 Geschäftsstelle in Saarbrücken

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.

8.2 *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.

8.3 Dagstuhl an der Universität Trier

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.

8.3 *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 2017 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 647 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. In 2017, 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 647 volumes during the year 2017 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.

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 (https://www.dagstuhl.de), die Informationen für Teilnehmer, zum Seminarprogramm usw. enthält,
- einen Server, auf dem sich DROPS befindet, Schloss Dagstuhls Publikationsplattform (http://drops.dagstuhl.de),
- den dblp-Server (https://dblp.dagstuhl.de und https://dblp.org).

Darüber hinaus stellt der IT-Service Tools für das gemeinschaftliche Arbeitsumfelds zur Verfügung und hält sie in Stand, z.B. *Sihot* (eine Software zur Organisation von Gastdaten), MySQL-Datenbanken, 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.

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[55] (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 https://www.dagstuhl.de, providing information for participants, information about the seminar program, etc.,
- a server hosting DROPS at http://drops.dagstuhl.de, Schloss Dagstuhl's publishing platform,
- the dblp server at https://dblp.dagstuhl.de and at https://dblp.org.

Furthermore, for internal work procedures, the IT service provides and maintains tools for a collaborative work environment, such as *Sihot* (a software for organizing guest data), MySQL data bases, TeamDrive (a cloud-based storage system), and several others.

[55] *eduroam* (education roaming) is a world-wide roaming access service developed for the international research and education community, see https://www.eduroam.org.

10 Kunst
Art

Dagstuhl als Galerie 10.1 *Dagstuhl as Art Gallery*

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des Flurs sowie die hervorragende Ausleuchtung mit starken Kontrasten zwischen Tag und Nacht bieten den Künstlern sehr gute Möglichkeiten, ihre Werke darzustellen. Die Kunstwerke an den Wänden des schmalen Gangs durchbrechen die Nüchternheit des Neubaus in anregender und angenehmer Weise. Die wechselnden Ausstellungen bieten einen erfrischenden und dynamischen Kontrast zu der ständigen Kunstsammlung von Schloss Dagstuhl.

Prof. Reinhard Wilhelm, ehemaliger wissenschaftlicher Direktor des Zentrums, fungierte nach seinem Eintritt in den Ruhestand im April 2014 weiterhin als Betreuer der Ausstellungsaktivitäten von Schloss Dagstuhl. Das Zentrum veranstaltet jährlich etwa drei bis vier Kunstausstellungen für jeweils zwei bis drei Monate.

Waren es bisher Künstler und einzelne Sammler, die ihre Werke ausstellten, so kam seit 2016 durch die Zusammenarbeit zwischen der Saarland-Sporttoto GmbH (kurz Saartoto), der Hochschule der Bildenden Künste Saar (kurz HBKsaar) und Schloss Dagstuhl die Sammlung von Saartoto als Reservoir für eine Ausstellungsserie hinzu. Als bedeutender Förderer von Künstlern besitzt Saartoto einen großen Bestand an Kunstwerken. Im Rahmen der Zusammenarbeit wird diese Kunstsammlung durch die HBKsaar erfasst und dokumentiert. Gleichzeitig wurden und sollen auch in Zukunft aus dem Saartoto-Fundus Ausstellungen für Schloss Dagstuhl zusammengestellt werden. Dabei werden die Kunstwerke aktuellen Werken von Künstlern der HBKsaar und aus der Großregion Saar-Lor-Lux gegenübergestellt. Die Galerie MediArt aus Luxemburg unterstützte das Projekt durch die Leihgabe von Bildern der Künstler aus der Großregion. Schloss Dagstuhl möchte an dieser Stelle allen beteiligten Personen danken, namentlich insbesondere Michael Burkert, Peter Jacoby und Josef Gros (Saartoto); Matthias Winzen und Nadine Brettar (HBKsaar); Paul Bertemes (MediArt); sowie Reinhard Wilhelm und Angelika Mueller-von Brochowski (Schloss Dagstuhl).

Die fünf Ausstellungen (siehe Fig. 10.1), die im Jahr 2017 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. The year 2016, however, saw the establishment of a cooperation between Saarland-Sporttoto GmbH (Saartoto for short), Hochschule für Bildende Künste Saar (university of art and design; HBKsaar for short), and Schloss Dagstuhl, which makes Saartotos collection accessible to Schloss Dagstuhl for a series of exhibitions. Being a major art sponsor, Saartoto is in possession of a substantial art collection. In the context of this collaboration, HBKsaar takes stock of and documents Saartoto's art collection. At the same time, there were, and will be, exhibitions at Schloss Dagstuhl where Saartoto artworks are contrasted with recent works by HBKsaar artists and artists from the greater region Saar-Lor-Lux (Saarland, Lorraine, and Luxembourg). The Luxembourg-based art gallery MediArt supported the project by loaning several paintings by artists from the greater region. Schloss Dagstuhl would like to thank everyone involved, especially Michael Burkert, Peter Jacoby, and Josef Gros (Saartoto); Matthias Winzen and Nadine Brettar (HBKsaar); Paul Bertemes (MediArt); as well as Reinhard Wilhelm and Angelika Mueller-von Brochowski (Schloss Dagstuhl).

The five exhibitions (cf. Fig. 10.1) hosted by Schloss Dagstuhl in 2017 are described below. Current exhibitions are open to the interested public upon request.

»Bodyart: Bilder auf menschlicher Haut«

Gesine Marwedel, die seit 2016 als freischaffende Künstlerin arbeitet, ist insbesondere für Ihre Tierbodypaintings bekannt und mehrfach preisgekrönt. Ausgestellt wurden in Dagstuhl zahlreiche großformatige Fotos ihrer vergänglichen Werke. In diesen ist der Mensch Mittelpunkt der gezeigten Bilder und doch ist er es nicht. Denn der

»Bodyart: Bilder auf menschlicher Haut«[56]

Gesine Marwedel, who is working as a freelance artist since 2016, is especially renowned for her *animal body paintings* and has won several awards. Numerous large photos of her ephemeral works have been presented in an exhibition at Dagstuhl. While people appear to be the focus of these images, in reality, they are not. The first

[56] engl. Bodyart: Paintings on Human Skin

erste Blick zeigt dem Betrachter eine durch Farbe auf der Haut erzeugte Illusion, die den Menschen verwandelt oder verschwinden lässt. Die Metamorphose in ein Tier, ein abstraktes Kunstwerk oder eine Landschaft ist gleichzeitig eine Hommage an die Schönheit und Einzigartigkeit jedes menschlichen Körpers.

Die Ausstellung in Schloss Dagstuhl fand in einem ungewohnt hohen Maße die Beachtung der regionalen Bevölkerung und Presse. Als ganz besonderen Höhepunkt konnte Schloss Dagstuhl zum Abschluss der Ausstellung die Künstlerin gewinnen, am 2. März 2017 im Schloss Dagstuhl eine Live-Bodypainting Aktion durchzuführen.

»S(ch)ichtwechsel: Das Saarland als Industrie- und Lebensraum«

Die geschichtsträchtige saarländische Bergbaulandschaft, die Auswirkungen der Montanindustrie auf die Menschen und deren Lebensraum sowie die Transformation der regionalen Landschaften stehen als künstlerisches Sujet im Fokus der zweiten Ausstellung der neuartigen Kooperation zwischen der Saarland Sporttoto GmbH, dem Leibniz Zentrum für Informatik Schloss Dagstuhl und der Kunsthochschule in Saarbrücken.

Geschichte und Gegenwart des durch den Steinkohleabbau und die Eisen- und Stahlindustrie geprägten Lebensraums spiegeln sich in Druckgrafiken, Grafiken, Fotografien, Installationen und Gemälden saarländischer Künstlerinnen und Künstler wider. Walter Bernsteins Milieuschilderungen und Erinnerungsbildern regionaler Industrielandschaft und Arbeitswelt seit den 50er Jahren begegnen später entstandene, assoziative Sichtweisen verschiedener im Saartotofundus vertretener Künstlerkollegen, sowie aktuelle Positionen von Studierenden und Gästen der Kunsthochschule in Saarbrücken.

Der konkreten Beobachtung von in den industriellen Arbeits- und Lebensraum eingebetteten Menschen und der Visualisierung von Kräften, Lichtszenarien, Materialität und Stimmungen steht die neue Bild- und Materialsprache der jungen Künstler gegenüber. Inspiriert vom ästhetischen Reiz der zum Relikt und Denkmal gewordenen Industrie-

glance shows an illusion produced by color on the skin that transforms the person or makes them fade into the background. The metamorphosis into an animal, abstract work of art, or a landscape is at the same time a homage to the beauty and uniqueness of each human body.

The exhibition in Schloss Dagstuhl attracted an unusually high amount of attention from the press and people in the region. As a very special highlight, Schloss Dagstuhl was able to win the artist for a live bodypainting performance at Schloss Dagstuhl at the end of the exhibition on March 2, 2017.

»S(ch)ichtwechsel: Das Saarland als Industrie- und Lebensraum«[57]

The historic Saarland mining landscape, the impact of the coal and steel industry on people and their habitat, as well as the transformation of regional landscapes are the focus of the second exhibition of the new cooperation between Saarland Sporttoto GmbH, the Leibniz Zentrum für Informatik Schloss Dagstuhl, and the HBKsaar in Saarbrücken.

The history and the present day of the living environment, which is characterized by coal mining and the iron and steel industry, is reflected in prints, graphics, photographs, installations, and paintings by artists from Saarland. Walter Bernstein's portrayals of the milieu and images of the memories of the regional industrial landscapes and the working environment since the 1950s meet the associative views of later artists from the Saartoto collection, as well as the current views of the students and guests of the HBKsaar.

The concrete observation of people embedded in the industrial working and living environment and the visualization of forces, lighting scenarios, materiality and moods are juxtaposed with the new visual and material language of the young artists. Inspired by the aesthetic charm of the industrial architecture, which has become a relic and monument, and seeking out the transformations of its surrounding natural environment, fragments from the exhi-

[57] engl. Change of View/Change of Shifts: Saarland as Industrial and Living Environment

»Bodyart: Bilder auf menschlicher Haut« Works from Gesine Marwedel | January 09 to March 3, 2017
»S(ch)ichtwechsel: Das Saarland als Industrie- und Lebensraum« Works from the art collection of Saartoto, of the artist Walter Bernstein, and of the HBKsaar exhibition »Land(e)scape« curated by Nadine Brettar and Maria Fernanda Palacios | March 13 to May 12, 2017
»a silent walk-fragmente« Works from Annette Marx | May 22 to July 26, 2017
»HAP Grieshaber und Emil Kiess« Works from a private art collection curated by Beate Kolodziej, M.A | August 28 to November 19, 2017
»Verneigung« Works from Emanuela Assenza | November 27, 2017 to January 18, 2018

Fig. 10.1
Art exhibitions in 2017.

triearchitektur, und die Transformationen des sie umgebenden Naturraums aufspürend, deuten Fragmente aus der 2017 von Nadine Brettar und der Masterabsolventin Maria Fernanda Palacios für die Galerie der HBKsaar kuratierten Ausstellung „Land(e)scape“ den künstlerischen Schichtwechsel an.

»a silent walk-fragmente«

Die in Saarbrücken lebende und arbeitende Künstlerin Annette Marx schloß zunächst ihr Diplom in Betriebswirtschaftslehre ab, ehe sie zwischen 1994 und 1998 eine künstlerische Ausbildungen an der Saarländischen Sommerakademie Wadgassen machte und sich anschließend bei verschiedenen Künstlern fortbildete. Seit 1998 setzt sie sich intensiv mit der abstrakten Malerei auseinander.

Die in ihrer Ausstellung in Schloss Dagstuhl gezeigten Arbeiten sind Teil einer Werkreihe vom Sommer 2016 bis Frühjahr 2017. In ihrem Mittelpunkt stehen fragmentarische Fotos, die die Künstlerin in der Völklinger Hütte aufgenommen hat. Fotografierte Gegenständlichkeit und malerische Abstraktion ergänzen einander zu einer spannenden Gratwanderung, wobei deren Grenzen unwichtig werden. Spielerisch gezeichnete Linien folgen Spuren von realer Verbindung oder der Zerstörung des Gegenstands im Foto. Diese Linien führen über das Foto hinaus in die Farbe, was teilweise surreal und irritierend anmutet. Kraft und Zartheit der Farbe wiederum lösen dies jedoch auf, fangen den Betrachter auf und nehmen ihn mit.

»HAP Grieshaber und Emil Kiess«

Bereits 2016 wurden in der Ausstellung »Farbe und Form, Abstraktion und Expression« Arbeiten aus einer baden-württembergischen Privatsammlung gezeigt, in der nicht nur regionale, sondern auch internationale Künstler vertreten sind. Diesem ersten Überblick folgt nun eine Präsentation, die den Fokus auf zwei Künstler richtet: HAP Grieshaber und Emil Kiess. Beide Namen sind mit der „Bernsteinschule“, einer privaten Kunstschule in Sulz am Neckar, verbunden. HAP Grieshaber gehört zu den bekanntesten Lehrern dieser Schule und prägte maßgeblich deren künstlerische Ausrichtung. Emil Kiess war einer seiner Schüler.

In der Ausstellung im Berichtsjahr wurden knapp 50 Gemälde und Grafiken von Kiess sowie Holzschnitte von Grieshaber gezeigt, die vor allem in den 50er und 60er Jahren entstanden sind. Wie bereits 2016 wurde auch diese Ausstellung von der Kunsthistorikerin Beate Kolodziej kuratiert.

HAP Grieshaber (1909 – 1981) war ein wichtiger Neuerer der Mediums Holzschnitt nach 1945 und ist vor allem bekannt für Arbeiten zu Mensch und Natur sowie für seine religiösen und mythologischen Motive. Emil Kiess (geb. 1930), der in Fürstenberg lebt und arbeitet, ist einer der renommiertesten Künstler im südwestdeutschen Raum. Er lotet die Möglichkeiten der Farbe zwischen Gegenständlichkeit und Abstraktion aus.

bition “Land(e)scape”, curated in 2017 by Nadine Brettar and the master’s graduate Maria Fernanda Palacio for the gallery of the HBKsaar, point to the artistic changeover.

»a silent walk-fragmente«

Annette Marx, an artist who lives and works in Saarbrücken, first completed her diploma in business administration before completing an artistic training at the Saarland Summer Academy in Wadgassen between 1994 and 1998 and then continued her training with various artists. Since 1998, she has been delving into abstract painting.

The works displayed in her exhibition at Schloss Dagstuhl are part of a series of works from summer 2016 to spring 2017, focusing on fragmentary photographs taken by the artist at the Völklinger Hütte. Photographed concreteness and abstraction in painting complement each other into an exciting tightrope walk, in which their boundaries become insignificant. Playfully drawn lines follow traces of the real connection or the destruction of the object in the photo. These lines lead beyond the photo into the color, which seems somewhat surreal and irritating. However, this is resolved by the power and tenderness of the color which pick the beholder up again.

»HAP Grieshaber and Emil Kiess«

As early as 2016, the exhibition »Farbe und Form, Abstraktion und Expression«[58] featured works from a private collection in Baden-Württemberg in which not only regional but also international artists are represented. This first overview is now followed by a presentation focusing on two artists: HAP Grieshaber and Emil Kiess. Both names are associated with the “Bernsteinschule”, a private art school in Sulz am Neckar. HAP Grieshaber is one of the most famous teachers of this school and was instrumental in shaping its artistic orientation. Emil Kiess was one of his students.

Almost 50 paintings and graphics by Kiess and woodcarvings by Grieshaber, mainly from the 1950s and 1960s, were displayed in the exhibition during 2017. Just like in 2016, the exhibition was curated by the art historian Beate Kolodziej.

HAP Grieshaber (1909 - 1981) was an important innovator of the medium of woodcarving after 1945 and is primarily known for his works on humans and nature as well as his religious and mythological motifs. Emil Kiess (born 1930), who lives and works in Fürstenberg, is one of the most renowned artists in southwestern Germany. He explores the potentials of color between abstract and figurative art.

[58] engl. "Color and Form, Abstraction and Expression"

»Verneigung«

„In der Farbe und Form liegt mehr als ich bin. Unendlich vielfältig und unerschöpflich sind die Bildmittel, ebenso wie die Wege, sie erscheinen zu lassen“ schreibt die in der Schweiz geborene Künstlerin Emanuela Assenza über ihre Arbeiten. Sie studierte Kunsttherapie und Kunstpädagogik an der an der Hochschule für Künste im Sozialen in Ottersberg und machte dort bei den Master of Fine Arts bei Prof. Jochen Stenschke.

Emanuela Assenza sucht in ihrer ungegenständlichen Malerei und Zeichnung die in den Bildmitteln selbst liegenden Kräfte zur Erscheinung zu bringen. Dieser Fokus auf eine autonome zeitgenössische Kunst führt zur Emanzipation von äußeren Bildinhalten. Wenn überhaupt, lässt sich eine stilistische Zuordnung ihrer Arbeiten unter einzelnen Aspekten zur Konkreten Kunst herstellen, die in Farbe und Form vollgültige Entitäten sieht.

»Verneigung«[59]

“There’s more to color and shape than I am. Infinitely diverse and inexhaustible is the visual medium, as is its way of expression.” writes Swiss-born artist Emanuela Assenza about her works. Assenza studied art therapy and art pedagogy at the Academy of Arts in Ottersberg, where she did her Master of Fine Arts under the supervision of Prof. Jochen Stenschke.

In her non-objective paintings and drawings, Emanuela Assenza seeks to make the inherent powers of the visual medium emerge. This focus on autonomous contemporary art leads to an emancipation from external image content. If at all, her works can be related in some aspects to Concrete Art, which sees color and shape as fully valid entities.

10.2 Kunstankauf durch Spenden

10.2 *Art Sponsorship and Donations*

Das Internetangebot von Schloss Dagstuhl enthält eine Seite, die es Teilnehmern, Einzelpersonen und Gruppen ermöglicht, Kunst für Dagstuhl zu stiften. Die Kunstobjekte werden über das Internet angeboten, dabei wird der Preis in kostengünstige Anteile aufgeteilt. Sobald alle Anteile eines Bilds gezeichnet sind, werden die Teilnehmer aufgefordert, den Gegenwert der bestellten Anteile als Spende einzuzahlen, wodurch dann das Objekt angekauft werden kann. Die Stifter werden sowohl in der virtuellen Internet-Galerie von Schloss Dagstuhl als auch an dem realen Objekt genannt. Dadurch ist es Schloss Dagstuhl möglich, Werke von Künstlern, die im Zentrum ausgestellt haben, anzukaufen und permanent auszustellen.

Im Jahr 2017 erhielt Schloss Dagstuhl insgesamt 861,15 € von verschiedenen Spendern. Wir möchten diese Stelle nutzen, allen Spendern, die 2017 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Nähere Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite[60] von Dagstuhl.

Dagstuhl’s website contains a page featuring an internet gallery enabling participants, individuals, and groups to make contributions to Dagstuhl for art donations. The works of art are featured online and donations are made by acquiring shares at affordable prices. Donors pay the value of their pledged shares as soon as a piece is fully subscribed for, thus allowing it to be purchased. Donors’ names appear in Dagstuhl’s online art gallery and also next to the art items themselves. In this way, Schloss Dagstuhl is able to purchase works of art from those who exhibit at the center, and add these works to its permanent art exhibition.

In 2017, Schloss Dagstuhl received a total of 861.15 € from various donors. We would like to thank all donors who contributed to Dagstuhl’s art collection in 2017.

For further information and current news about Dagstuhl’s art program, please visit Dagstuhls’s art webpage[60].

10.3 Dagstuhls permanente Kunstausstellung

10.3 *Dagstuhl’s Permanent Art Exhibition*

Die von Gästen immer wieder positiv hervorgehobene Kunstsammlung geht auf den Gründungsdirektor Professor Wilhelm zurück. Seine Idee war es, den 1995 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

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 180

[59] engl. Bow
[60] https://www.dagstuhl.de/art/

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.

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.

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 mit hervorragendem Ergebnis evaluiert. In der Stellungnahme des Senats der Leibniz-Gemeinschaft wurde das Veranstaltungsprogramm und die Beteiligung an der Literaturdatenbank dblp als „exzellent“ bewertet, während der Bereich Open Access (Publishing) als „sehr gut“ bewertet wurde.

Schloss Dagstuhl is operated as a non-profit organization by currently eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e. V.[61] (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[62] (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”[63] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[64] in 2008.

In July 2009, Schloss Dagstuhl was evaluated for the first time by the Leibniz Association. The March 2010 findings of the evaluation commission were very positive, and established that the center has shown outstanding commitment to its designated task of supporting the international computer science research community by providing a seminar center for academic events. In 2016, Schloss Dagstuhl has been evaluated again, with excellent results. In the Leibniz Association Senate report, the seminar program and the cooperation with the computer science bibliography dblp were rated as “excellent” whereas the Open Access Publishing was rated “very good.”

Organe der Gesellschaft 11.2 *Dagstuhl Organs*

Die drei Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, die stellvertretend für die Gesellschaft als juristische Person handeln, sind die folgenden:

- Gesellschafterversammlung
- Aufsichtsrat
- Geschäftsführung

Details zu den Organen sind den folgenden Abschnitten zu entnehmen.

The three organs of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, which act for the company as a legal entity, are the following:

- Shareholders’ Meeting
- Supervisory Board
- Management

Detailed information is given in the sections below.

[61] engl.: German Informatics Society
[62] engl.: Federal/State Government Commission for Educational Planning and Research Promotion
[63] engl.: International Conference and Research Center for Computer Science
[64] 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 für das vierte Geschäftsjahr. Die Vertreter der Universitäten in Darmstadt und Stuttgart wechseln im Allgemeinen Amtszeit für Amtszeit mit denen der Universitäten in Frankfurt und Trier ab.

Der Aufsichtsrat entscheidet über die Berufung und Abberufung der Geschäftsführer sowie der Mitglieder des Wissenschaftlichen Direktoriums, des Wissenschaftlichen Beirates und des Kuratoriums. Alle Beschlüsse, die die Finanzen oder das Vermögen der Firma betreffen, benötigen seine Zustimmung. Beschlüsse von forschungspolitischer Bedeutung und Beschlüsse mit erheblichen finanziellen Auswirkungen können nicht gegen die Stimmen der Vertreter des Bundes und der beiden Sitzländer gefasst werden. Der Aufsichtsrat entscheidet zudem über die Erteilung einer Prokura.

Die Geschäftsführung

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH hat zwei Geschäftsführer (siehe Fig. 11.3), die gemeinsam die Gesellschaft vertreten. Die Geschäftsführung besteht aus dem *Wissenschaftlichen Direktor* und dem *Technisch-administrativen Geschäftsführer*.

Der Wissenschaftliche Direktor ist verantwortlich für die wissenschaftlich-fachliche Zielsetzung und die Programmgestaltung, und ist zudem Mitglied und Vorsitzender des Wissenschaftlichen Direktoriums. Seit Mai 2014

Shareholders' Meeting

All changes to the company, in particular the inclusion of new associates, the revision of the Shareholders' agreement, and the dissolution of the company, are decided by the shareholders. Shareholders also confirm new members forwarded by them to the Supervisory Board and the appointment or recall of the managing directors. In accordance with their shares, all shareholders currently have the same number of votes except the Gesellschaft für Informatik, which has three times the number of votes of the other shareholders in proportion to its larger number of shares. Decisions are made in shareholders' meetings which take place at least once a year, or via a written vote.

Supervisory Board

The Supervisory Board is responsible for ensuring that the management complies with the center's objectives in a legally and economically meaningful manner. The board is involved in all essential matters with regard to research and financial planning.

The 12-member board (see Fig. 11.2) is composed of four representatives of the Gesellschaft für Informatik, one representative from each of the three founding universities, two representatives of the four universities that subsequently joined, and one representative from each of the German federal government and the two host state governments of Saarland and Rhineland-Palatinate. The Supervisory Board members typically hold office for at least four full fiscal years. The term of office ends with the approval for the fourth fiscal year. In general, representatives of the universities in Darmstadt and Stuttgart and of the universities in Frankfurt and Trier rotate after each term of office.

The Supervisory Board formally appoints and recalls the managing directors and members of the Scientific Directorate, Scientific Advisory Board, and Industrial Curatory Board. Furthermore, all decisions regarding financial issues and company assets must be approved by the Supervisory Board. Consent cannot be given against the votes of the represented (federal) state governments if the matter affects political issues in the area of science or has considerable financial weight. The Supervisory Board also holds decision power with respect to the granting of power of attorney.

Management

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH has two managing directors (see Fig. 11.3) who jointly represent the company. These are the *Scientific Director* and the *Technical Administrative Director*.

The Scientific Director is in charge of drafting the company's scientific goals and program planning, and is also a member and the chairperson of the Scientific Directorate. Since May 2014, Prof. Raimund Seidel, Ph.D., is the Scientific Director of Schloss Dagstuhl.

ist Prof. Raimund Seidel, Ph. D., der wissenschaftliche Direktor von Schloss Dagstuhl.

Der Wissenschaftliche Direktor wird dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Dieser Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Seit Juli 2014 hat Frau Heike Meißner diese Position inne.

Gremien der Gesellschaft 11.3

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.

Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 11.4) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Der Wissenschaftliche Direktor gehört dem Wissenschaftlichen Direktorium an. Er empfiehlt dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können das bestehende Wissenschaftliche Direktorium sowie der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt dem Beirat zusammen mit dem Wissenschaftlichen Direktor.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums – mit Ausnahme der des Wissenschaftlichen Direktors – beträgt drei Jahre. Sie beginnt am 1. November des Jahres ihrer Berufung und endet drei Jahre später am 31. Oktober. Eine Wiederberufung ist möglich.

Der Wissenschaftliche Beirat

Die Aufgaben des Wissenschaftlichen Beirats (siehe Fig. 11.5) werden nicht nur durch den Gesellschaftsvertrag festgelegt, sondern auch durch die Empfehlungen der Leibniz-Gemeinschaft. Im Sinne dieser wirkt der Wissen-

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, Ms Heike Meißner holds this position.

Dagstuhl Bodies

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.

Scientific Directorate

The Scientific Directorate (see Fig. 11.4) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In its biannual directorate meetings, the Scientific Directorate discusses the proposals and decides which of them to accept or reject.

The Scientific Director is member of the Scientific Directorate. He recommends to the Supervisory Board the number of Scientific Directorate members. Candidates for the Scientific Directorate may be suggested not only by the shareholders, but also by the Scientific Directorate and the Scientific Advisory Board. The selection of candidates, which are recommended to the Supervisory Board for appointment, is carried out by the Scientific Advisory Board together with the Scientific Director.

The term of office of Scientific Directorate members – with the exception of the Scientific Director – is three years. It begins on November 1 of the year of appointment and ends three years later on October 31. Reelections are possible.

Scientific Advisory Board

The tasks of the Scientific Advisory Board (see Fig. 11.5) are not only defined by the Shareholders' Agreement, but also by the recommendations of the Leibniz Association. The latter stipulates two different ways in

schaftliche Beirat auf zwei Wegen bei der Qualitätssicherung mit. Zum einen berät er die Leitung in Fragen der Forschungs- und Entwicklungsplanung, nimmt Stellung zu den Programmbudgets und gibt Empfehlungen zum Ressourceneinsatz. Er unterstützt weiterhin den Aufsichtsrat bei wichtigen Entscheidungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung (SAE) der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat sollte aus sechs bis zwölf international angesehenen, im Berufsleben stehenden Wissenschaftlern aus dem In- und Ausland bestehen. Die Amtszeit der Mitglieder beträgt vier Jahre, eine einmalige Wiederberufung ist möglich. Der Beirat wählt aus seiner Mitte einen Vorsitzenden. Der Wissenschaftliche Beirat tagt einmal im Jahr. Mitglieder des Beirats werden vom Aufsichtsrat auf Vorschlag des Beirats ernannt.

Das Kuratorium

Das Kuratorium (siehe Fig. 11.6) erfüllt eine Transmissionsfunktion zwischen Schloss Dagstuhl und den Forschungsabteilungen und Entwicklungslaboren der Industrie. Es hat die Aufgabe, die Akzeptanz des Zentrums in Verwaltung, Industrie und Wirtschaft abzusichern und als Förderungsorganisation die wirtschaftliche Basis des Zentrums zu verbreitern. Mitglieder des Kuratoriums werden vom Aufsichtsrat ernannt.

Nach seiner Geschäftsordnung hat das Kuratorium mindestens fünf Mitglieder, deren Amtszeit vier Jahre beträgt. Eine einmalige Wiederberufung ist möglich. Die Mitglieder des Kuratoriums unterstützen das Zentrum dabei, aktuelle Themen zu identifizieren und dazu geeignete zugkräftige Organisatoren aus der Industrie zu gewinnen. Sie werden ebenso gebeten, geeignete Personen aus der Industrie als Teilnehmer von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops zu benennen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

which the Scientific Advisory Board is involved in quality assurance. On the one hand, the board offers advice to the management with regard to research as well as development planning and issues comments on the program budget draft, making recommendations on the use of resources. It also assists the Supervisory Board in important decisions with regard to future development of the institute as well as the acquisition of management staff. On the other hand, it carries out an audit of the entire institute between two evaluations by the Senatsausschuss Evaluierung (SAE, Senate Committee Evaluation) of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board, and the SAE.

The Scientific Advisory Board should consist of six to twelve internationally reputable, well established scientists and academics from Germany and abroad. The term of office for members is four years and can be prolonged once. The Scientific Advisory Board members elect a chairperson from their midst. The board convenes once a year. Members are appointed by the Supervisory Board in accordance with the suggestions of the Scientific Advisory Board.

Industrial Curatory Board

The Industrial Curatory Board (see Fig. 11.6) performs a transmissional function between the center and the industrial R&D departments and laboratories. Its role is to secure acceptance of Schloss Dagstuhl within the business, industry and administrative communities, and as a promotional organization to broaden the economic basis of the center. Board members are appointed by the Supervisory Board.

According to its rules of procedure, the Industrial Curatory Board consists of at least five members whose term of office is four years. A one-off reappointment for a second term is possible. The board members help the center to identify current R&D topics for seminars and locate attractive organizers in industry. The Industrial Curatory Board is regularly called upon to propose suitable participants for Dagstuhl Seminars and Dagstuhl Perspectives Workshops known to it from its activities. It convenes once a year together with the Scientific Advisory Board.

Gesellschafter | Associates
Centrum Wiskunde & Informatica (CWI), The Netherlands
Gesellschaft für Informatik e. V., Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Karlsruher Institut für Technologie (KIT), Germany
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Berlin, Germany
Technische Universität Darmstadt, Germany
Technische Universität Kaiserslautern, Germany
Universität des Saarlandes, Germany
Universität Stuttgart, Germany
Universität Trier, Germany

Fig. 11.1
Associates.

Aufsichtsrat | Supervisory Board
Dr. Marc Brüser Ministerium für Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state | *tenure started in May 2017*
Prof. Alejandro P. Buchmann, Ph. D. Technische Universität Darmstadt, Germany | Representative of Technische Universität Darmstadt | *tenure ended in May 2017*
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 Lindenstruth Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany | Representative of Johann Wolfgang Goethe-Universität Frankfurt am Main | *tenure started in May 2017*
Prof. Dr. Volker Linneweber Universität des Saarlandes, Saarbrücken, Germany | Representative of Universität des Saarlandes | *tenure ended in October 2017*
Prof. Dr. Erhard Plödereder Universität Stuttgart, Germany | Representative of Universität Stuttgart | *tenure ended in May 2017*
Prof. Dr. Arnd Poetzsch-Heffter Technische Universität Kaiserslautern, Germany | Representative of Technische Universität Kaiserslautern
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | Representative of the Saarland
Prof. Dr. Ralph Schenkel Universität Trier | Representative of Universität Trier | *tenure started in May 2017*
Prof. Dr. Manfred J. Schmitt Universität des Saarlandes, Saarbrücken, Germany | Representative of Universität des Saarlandes | *tenure started in November 2017*
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
Marcus Wittrin Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government | *tenure ended in November 2017*
N.N. Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government

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. Daniel Cremers Technische Universität München, Germany | *tenure started in November 2017*
Prof. Dr. Stefan Diehl Universität Trier, Germany
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany
Prof. Dr. Hans Hagen Technische Universität Kaiserslautern, Germany | *tenure ended in October 2017*
Prof. Dr. Lynda Hardman Centrum Wiskunde & Informatica (CWI), Amsterdam and University of Utrecht, The Netherlands | *tenure started in November 2017*
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 | *tenure ended in October 2017*
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 Ludwig-Maximilians Universität München, Germany | *since October 2017 at LMU München; before at Universität Stuttgart, Germany*
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Dr. Emmanuel Thomé Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France | *tenure started in November 2017*
Prof. Dr. Ir. Arjen P. de Vries Centrum Wiskunde & Informatica (CWI), Amsterdam and University of Utrecht, The Netherlands | *tenure ended in October 2017*
Prof. Dr. Heike Wehrheim Universität Paderborn, Germany | *tenure started in November 2017*
Prof. Dr.-Ing. Klaus Wehrle Rheinisch-Westfälische Technische Hochschule Aachen, Germany | *tenure ended in October 2017*
Prof. Dr. Verena Wolf Universität des Saarlandes, Saarbrücken, Germany

Fig. 11.4
Scientific Directorate.

Wissenschaftlicher Beirat | Scientific Advisory Board
Prof. Dr. Christel Baier Technische Universität Dresden, Germany
Prof. Dr. Anja Feldmann Max-Planck-Institut für Informatik, Saarbrücken, Germany | *tenure started in January 2018*
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 EML European Media Laboratory GmbH, Heidelberg, Germany
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany

Fig. 11.5
Scientific Advisory Board.

Kuratorium | Industrial Curatory Board
Dr. Udo Bub EIT ICT Labs, Berlin, Germany
Dr.-Ing. Elmar Dorner SAP SE, Karlsruhe, Germany
Dr. Jo Ebergen Oracle Labs, Redwood Shores, United States | *tenure ended in May 2017*
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“

Holger Hermanns (Universität des Saarlandes, Germany)
Erich Reindel (Universität des Saarlandes, Germany)

Seit Mitte 2014 gibt es den Verein zur Förderung von Schloss Dagstuhl — Leibniz-Zentrum für Informatik e.V.. Der sehr technische und holprig klingende Name spiegelt dabei exakt den Vereinszweck wider: die Förderung von Wissenschaft und Forschung im Leibniz-Zentrum für Informatik in Schloss Dagstuhl. Für die Webpräsenz wurde allerdings ein wesentlich geschmeidigerer Name gewählt: „Friends of Dagstuhl“ (http://www.friends-of-dagstuhl.de).

Der Verein ist darauf ausgerichtet, finanzielle Mittel zur erfolgreichen Umsetzung des Vereinszwecks zu beschaffen und bereitzustellen sowie die ihm zu diesem Zweck anvertrauten Mittel treuhänderisch zu verwalten. Die Stiftung Informatikzentrum Schloss Dagstuhl wurde daher auch als nicht rechtsfähige Stiftung in den Verein überführt. Seit Ende 2014 vertreten nun die Freunde von Dagstuhl die Stiftung im Rechts- und Geschäftsverkehr und verwalten das Stiftungsvermögen. Der Verein wird von einem Vorstand (siehe Fig. 12.1 und Fig. 12.3) geleitet.

Nach sorgfältiger Prüfung aller Möglichkeiten und unter Einbeziehung des Stiftungsrates (siehe Fig. 12.2) konnten dann im Jahr 2016 die entscheidenden Schritte zur Anlage des Stiftungsvermögens vorgenommen werden. Trotz der andauernden Niedrigzinsphase wurde eine sichere, aber dennoch nicht ganz renditelose Anlageform gefunden. Hierzu wurden Verträge mit einer professionellen und auf Stiftungskapital spezialisierten Vermögensverwaltung geschlossen. Bisher hat die gewählte Anlageform die Erwartungen erfüllt.

Inzwischen gehören dem Verein 43 persönliche sowie bisher leider nur vier institutionelle Mitglieder an. Gerade im Hinblick auf die noch geringe Anzahl institutioneller Mitglieder wünschen sich die Freunde von Dagstuhl noch regen Zulauf.

Weitere Informationen zum Verein, aber auch Mitgliedschaftsanträge finden Sie unter http://www.friends-of-dagstuhl.de.

Association “Friends of Dagstuhl”

Since mid 2014, the registered association to support of Schloss Dagstuhl – Leibniz Center for Informatics (Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.) exists. This very technical and rather clumsy name nevertheless reflects the precise purpose of the association: the support of science and research at the Leibniz Center for Informatics at Schloss Dagstuhl. A significantly smoother name, i.e. “Friends of Dagstuhl”, was chosen for the website (http://www.friends-of-dagstuhl.de).

The association aims at acquiring and providing funds for the successful execution of its purpose, as well as holding these funds in trust. The Dagstuhl Foundation (Stiftung Informatikzentrum Schloss Dagstuhl) was therefore integrated into the association as a dependent foundation. Since late 2014, Friends of Dagstuhl represent the foundation in legal and business transactions and manage the foundation assets. The association is chaired by a board (see Fig. 12.1 and Fig. 12.3).

After a careful evaluation of all options and in cooperation with the foundation council (see Fig. 12.2), crucial steps were taken regarding the foundation assets in 2016. It was possible to invest the capital safely but not without return, in spite of the persistently low interest rates. Subsequently, contracts with an investment management company specialized in foundation assets were made. So far, the chosen investment form has met the expectations.

Currently, the association has 43 individual members but unfortunately only 4 institutional members. Especially with regard to the small number of institutional members, the Friends of Dagstuhl desire to welcome new members.

Further information about the association as well as the membership application form can be found at http://www.friends-of-dagstuhl.de.

Vorstand des Vereins | Chair of the association
Prof. Dr. Holger Hermanns (Vorstandsvorsitzender | First deputy chairperson) Universität des Saarlandes, Saarbrücken, Germany
Angelika Müller-von Brochowski (Schriftführerin | Secretary)
Erich Reindel (Schatzmeister | Treasurer) Universität des Saarlandes, Saarbrücken, Germany

Fig. 12.1
Der Vorstand des Vereins “Friends of Dagstuhl”
The chair of the association “Friends of Dagstuhl”

Stiftungsrat | Foundation council
Prof. Dr. Holger Hermanns (Vorstandsvorsitzender des Vereins “Friends of Dagstuhl” | First deputy chairperson of the association “Friends of Dagstuhl”) Universität des Saarlandes, Saarbrücken, Germany
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

13.1 Statistiken zu Seminaren und 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.

13.1 *Statistics on Seminars and Workshops*

This section provides statistical data about the scientific program and the composition of program participants. Charts and tables in this chapter may be outlined as follows.

Proposal-related data: Fig. 13.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, as well as acceptance rates for recent years. The size and duration of accepted seminars and workshops are displayed in Fig. 13.2.

Event-related data: Fig. 13.3 shows the number and the fraction of invited seminar participants who accepted or declined the invitation. The distribution of the rate is given in Fig. 13.4. In contrast, Fig. 13.5 visualizes how much of the reserved space was actually used by seminar participants. Data related to the number of seminars held in the last years together with their sizes and durations are given in Fig. 13.6. Fig. 13.7 shows the distribution of different types of events at Dagstuhl.

Participant-related data: Fig. 13.8 shows the number of participants according to event type. Fig. 13.9 shows the distribution of country affiliations.

Survey-related data: In this section we present data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results of the participants survey for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Fig. 13.10. Fig. 13.11 displays how often participants have attended seminars in the past. Fig. 13.12 gives data on the seniority of participants. While Dagstuhl Seminars and Dagstuhl Perspectives Workshops are mainly oriented towards academic researchers, also researchers and developers from industry are welcome. The distribution of their ratio compared to all participants of a seminar is shown in Fig. 13.13.

Utilization-related data: Finally, Fig. 13.14 states the number of overnight stays – separated by event type – hosted at Schloss Dagstuhl as well as their distribution about the weeks.

Gender-related data: Fig. 13.15 shows mixed-gender data with respective to organizer teams of Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In contrast Fig. 13.16 presents this data with respect to proposed seminar events. In Fig. 13.17 and Fig. 13.18 data is given with regard to female participants and invitees, respectively. The distribution of the rate of female participants by seminar and year is displayed in Fig. 13.19.

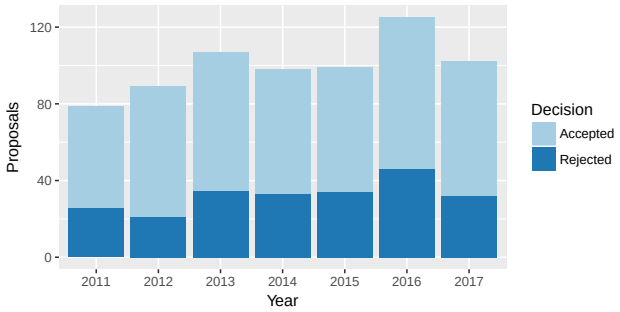

(a) Chart for 2011–2017

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
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
2017	102	70	68.6	32	31.4

(b) Detailed numbers for 2011–2017

Fig. 13.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

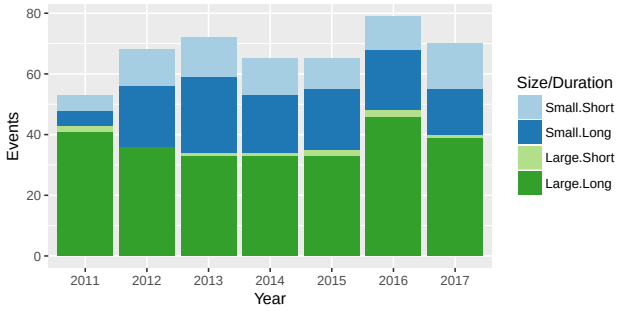

(a) Chart for 2011–2017

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
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
2017	15	15	1	39	70

(b) Detailed numbers for 2011–2017

Fig. 13.2
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2011–2017. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

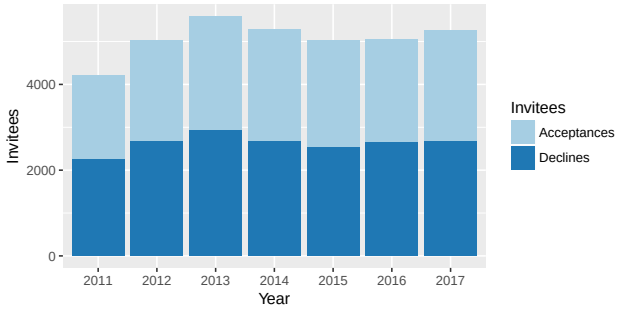

(a) Chart for 2011–2017

Year	Invitees	Acceptances		Declines	
	#	#	%	#	%
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
2017	5267	2572	48.8	2695	51.2

(b) Detailed numbers for 2011–2017

Fig. 13.3
Total number of invitees, acceptances, and declines for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

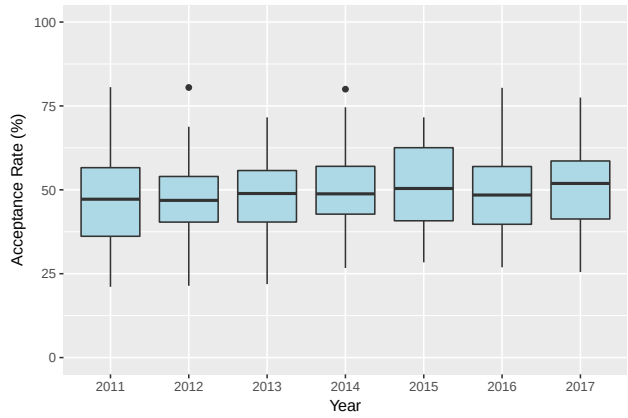

(a) Chart for 2011–2017

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2017	25.5	77.5	50.3	12.4

(b) Detailed numbers for 2011–2017

Fig. 13.4
Distribution of the acceptance rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2011–2017. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

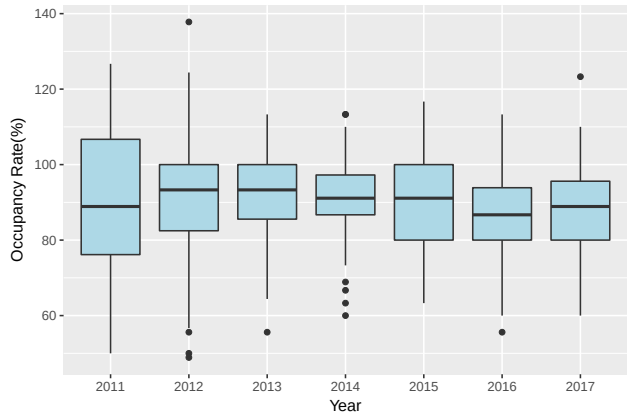

(a) Chart for 2011–2017

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2017	60.0	123.3	87.3	12.3

(b) Detailed numbers for 2011–2017

Fig. 13.5
Distribution of the occupancy rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2011–2017. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

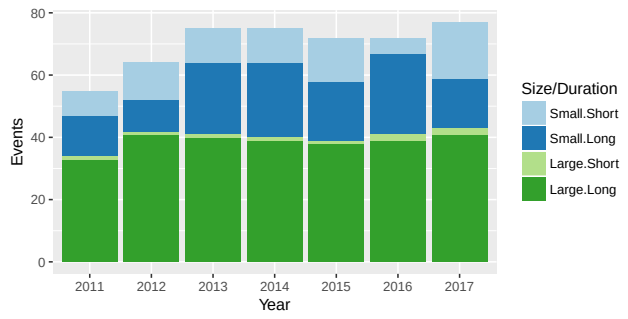

(a) Chart for 2011–2017

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
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
2017	18	16	2	41	77

(b) Detailed numbers for 2011–2017

Fig. 13.6
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2011–2017. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

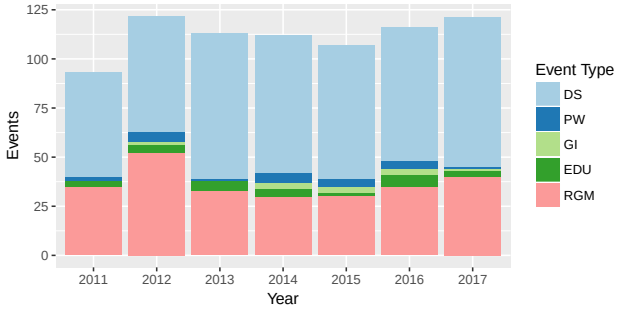

(a) Chart for 2011–2017

Year	DS	PW	GI	EDU	RGM	Total
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
2017	76	1	1	3	40	121

(b) Detailed numbers for 2011–2017

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.

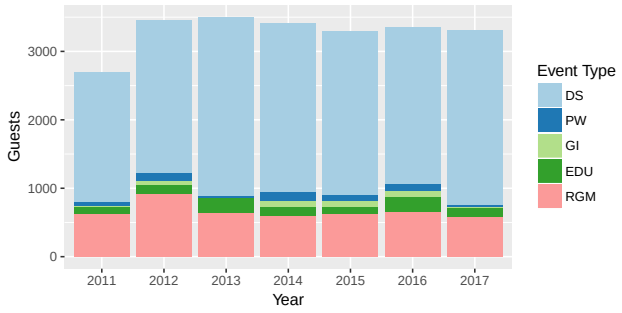

(a) Chart for 2011–2017

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	#
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
2017	2551	77.1	21	0.6	21	0.6	131	4.0	584	17.7	3308

(b) Detailed numbers for 2011–2017

Fig. 13.8
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

Country	A	B	Total	Country	A	B	Total
Germany	600	557	1157	Republic of Korea	6	2	8
United States	581	31	612	Russian Federation	8	0	8
United Kingdom	256	12	268	Slovenia	8	0	8
France	183	9	192	Hong Kong	7	0	7
Netherlands	111	11	122	Ireland	5	2	7
Switzerland	82	13	95	Turkey	5	1	6
Italy	68	8	76	Malta	5	0	5
Canada	58	4	62	Serbia	1	3	4
Sweden	56	6	62	Estonia	3	0	3
Austria	50	9	59	Iceland	3	0	3
Japan	54	0	54	Croatia	0	2	2
Israel	53	0	53	Cyprus	2	0	2
Denmark	50	0	50	Iran	1	1	2
Spain	34	8	42	Saudi Arabia	2	0	2
Australia	39	1	40	Algeria	1	0	1
Luxembourg	9	28	37	Argentina	1	0	1
Belgium	34	0	34	Bulgaria	1	0	1
Finland	26	4	30	Colombia	1	0	1
China	22	2	24	Kuwait	1	0	1
Norway	20	4	24	Latvia	1	0	1
India	19	1	20	Namibia	1	0	1
Czech Republic	18	0	18	Palestinian Territory	1	0	1
Poland	18	0	18	Puerto Rico	1	0	1
Singapore	14	1	15	Qatar	1	0	1
Portugal	14	0	14	Romania	1	0	1
Greece	7	6	13	South Africa	1	0	1
Brazil	6	4	10	Taiwan	1	0	1
New Zealand	9	0	9	United Arab Emirates	1	0	1
Chile	2	6	8	Venezuela	1	0	1
Hungary	8	0	8	Total	2572	736	3308

(a) Details for 2017 by country

Region	A		B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	1087	42.3	126	17.1	1213	36.7
Germany	600	23.3	557	75.7	1157	35
North America	640	24.9	35	4.8	675	20.4
Asia	183	7.1	7	1	190	5.7
Australia	48	1.9	1	0.1	49	1.5
South America	11	0.4	10	1.4	21	0.6
Africa	3	0.1	0	0	3	0.1
Total	2572	100	736	100	3308	100

(b) Details for 2017 by region

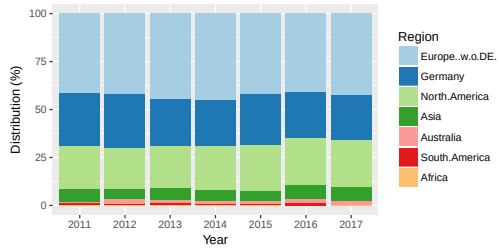

(c) Graphical distribution of seminar type A in 2011–2017 by year and region

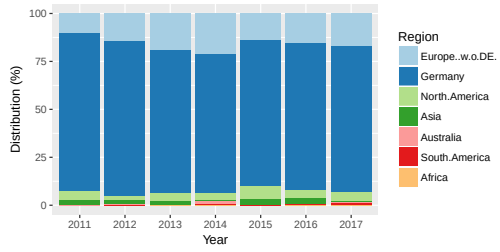

(d) Graphical distribution of seminar type B in 2011–2017 by year and region

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).

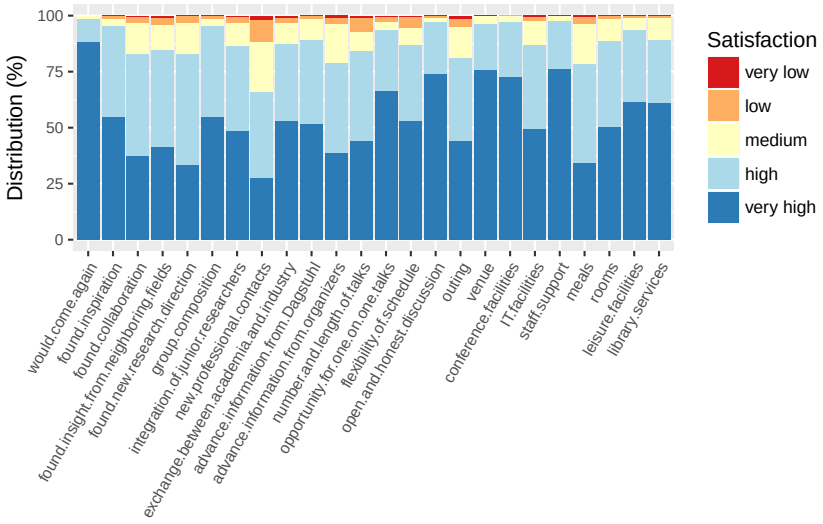

(a) Graphical distribution for 2017

	2011	2012	2013	2014	2015	2016	2017	2017 – Detailed Numbers					
	ø	ø	ø	ø	ø	ø	ø	1	2	3	4	5	total
would come again	4.9	4.9	4.9	4.9	4.9	4.9	4.9	1	0	18	148	1278	1445
found inspiration	4.4	4.4	4.5	4.4	4.5	4.5	4.5	3	15	51	589	805	1463
found collaboration	4.0	4.0	4.0	4.1	4.1	4.1	4.2	5	44	198	658	541	1446
found insight from neighboring fields	4.1	4.1	4.1	4.2	4.3	4.2	4.2	12	48	158	630	604	1452
found new research direction	4.0	4.0	4.0	4.0	4.1	4.1	4.1	4	42	198	716	481	1441
group composition	4.4	4.4	4.4	4.4	4.5	4.5	4.5	2	14	51	592	803	1462
integration of junior researchers	4.2	4.2	4.2	4.2	4.2	4.3	4.3	5	45	147	545	701	1443
new professional contacts	3.7	3.6	3.7	3.7	3.6	3.7	3.8	25	139	323	546	391	1424
exchange between academia and industry	4.1	4.3	4.2	4.2	4.3	4.3	4.4	8	22	94	338	521	983
advance information from Dagstuhl	4.3	4.4	4.4	4.4	4.4	4.4	4.4	2	19	137	541	750	1449
advance information from organizers	4.2	4.1	4.1	4.1	4.1	4.2	4.1	13	36	246	567	548	1410
number and length of talks	4.1	4.1	4.2	4.1	4.2	4.3	4.2	11	93	121	580	632	1437
opportunity for one on one talks	4.5	4.4	4.5	4.5	4.5	4.6	4.6	5	34	55	390	964	1448
flexibility of schedule	4.2	4.2	4.2	4.3	4.3	4.4	4.3	6	71	112	490	766	1445
open and honest discussion	4.7	4.6	4.7	4.7	4.7	4.7	4.7	3	8	27	340	1064	1442
outing	4.2	4.1	4.1	4.1	4.1	4.2	4.2	14	38	150	400	477	1079
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	2	0	49	305	1102	1458
conference facilities	4.8	4.7	4.6	4.7	4.6	4.7	4.7	0	2	33	356	1057	1448
IT facilities	4.6	4.4	4.4	4.4	4.3	4.4	4.3	6	24	136	484	637	1287
staff support	4.7	4.7	4.7	4.7	4.7	4.7	4.7	1	0	30	300	1066	1397
meals	4.3	4.2	4.1	4.1	4.1	4.1	4.1	11	39	259	641	493	1443
rooms	4.5	4.4	4.4	4.4	4.4	4.4	4.4	1	20	138	562	728	1449
leisure facilities	4.6	4.6	4.6	4.6	4.6	4.5	4.5	4	10	70	432	820	1336
library services	4.5	4.5	4.5	4.5	4.5	4.5	4.5	1	3	63	183	391	641

(b) Averages for 2011–2017 and detailed numbers for 2017: 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.

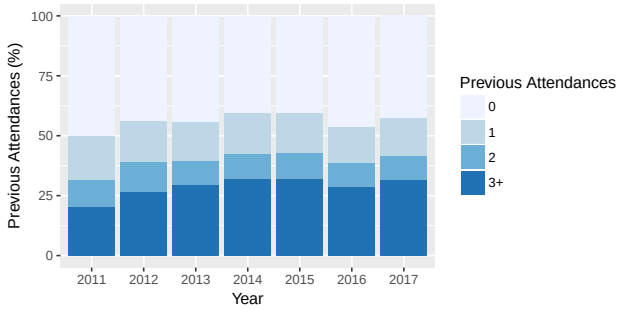

(a) Graphical distribution for 2011–2017

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
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
2017	607	43	222	16	148	10	446	31	1423

(b) Detailed numbers for 2011–2017

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.

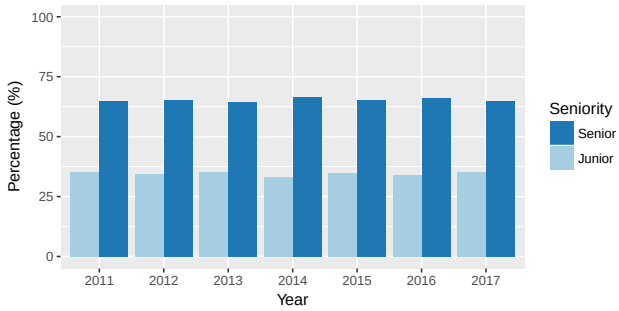

(a) Chart for 2011–2017

Year	Junior		Senior		Total
	#	%	#	%	#
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
2017	422	35.2	778	64.8	1200

(b) Detailed numbers for 2011–2017

Fig. 13.12

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

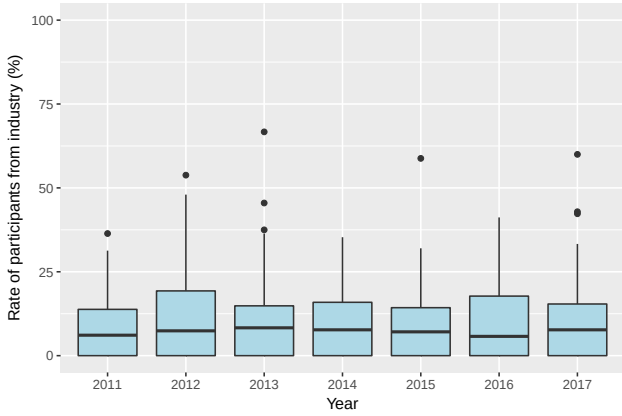

(a) Chart for 2011–2017

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2017	0.0	60.0	10.9	11.6

(b) Detailed numbers for 2011–2017

Fig. 13.13

Distribution of the rate of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2011–2017, 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”.

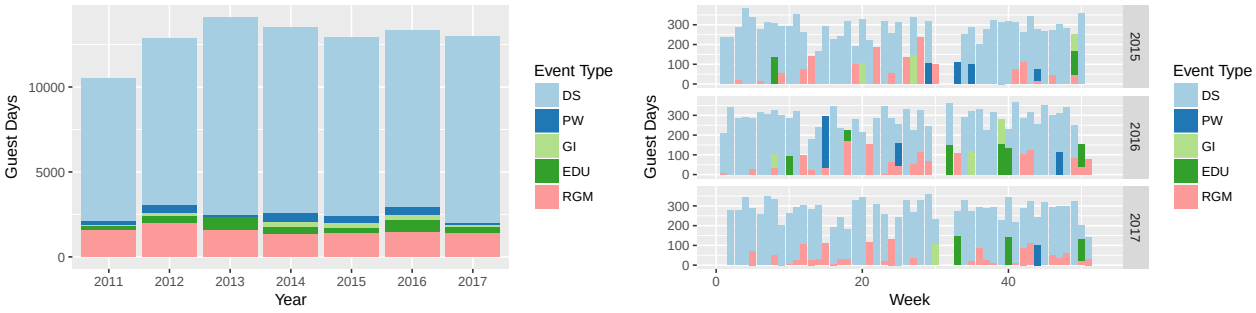

(a) Chart for 2011–2017

(b) Graphical distribution for 2015–2017 by week

Year	DS	PW	GI	EDU	RGM	Total
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
2017	10989	102	105	401	1391	12988

(c) Detailed numbers for 2011–2017

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.

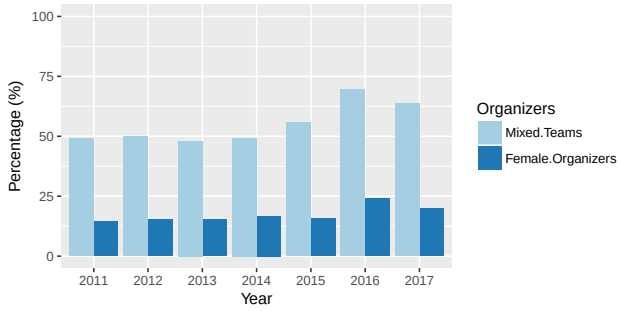

(a) Chart for 2011–2017

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
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
2017	77	49	63.6	303	60	19.8

(b) Detailed numbers for 2011–2017

Fig. 13.15

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

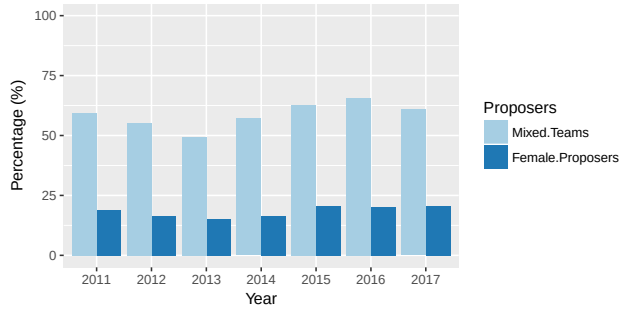

(a) Chart for 2011–2017

	Proposer Teams			Proposers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
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
2017	102	62	60.8	394	81	20.6

(b) Detailed numbers for 2011–2017

Fig. 13.16
Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.

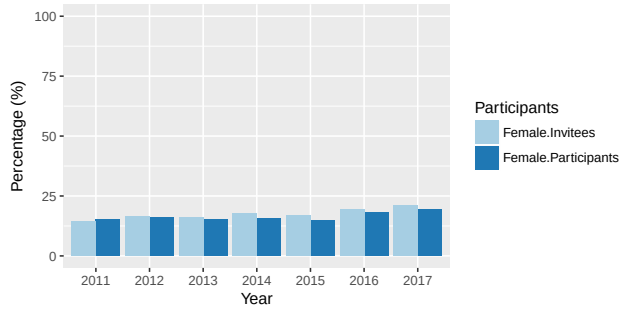

(a) Chart for 2011–2017

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2011	4223	606	14.3	1958	295	15.1
2012	5033	822	16.3	2346	378	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
2017	5267	1110	21.1	2572	495	19.2

(b) Detailed numbers for 2011–2017

Fig. 13.17
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

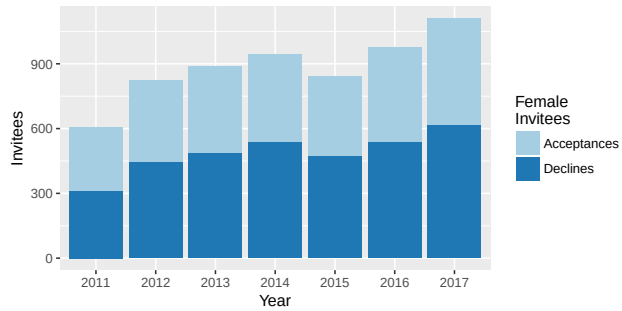

(a) Chart for 2011–2017

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2011	606	295	48.7	311	51.3
2012	822	378	46.0	444	54.0
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
2017	1110	495	44.6	615	55.4

(b) Detailed numbers for 2011–2017

Fig. 13.18
Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.

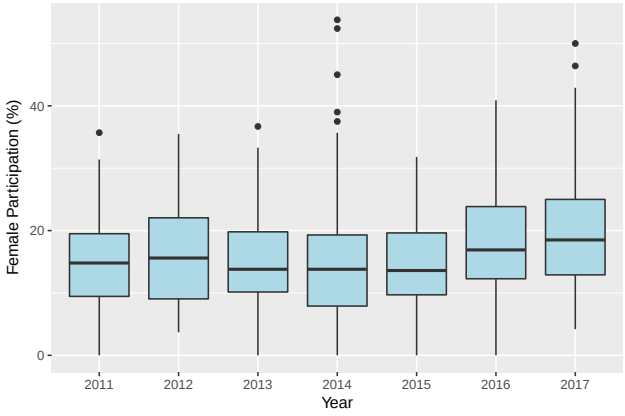

(a) Chart for 2011–2017

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2011	0.0	35.7	14.7	7.4
2012	3.7	35.5	16.1	7.8
2013	0.0	36.7	15.1	7.3
2014	0.0	53.8	15.9	11.1
2015	0.0	31.8	14.8	7.7
2016	0.0	40.9	18.3	9.1
2017	4.2	50.0	19.7	9.8

(b) Detailed numbers for 2011–2017

Fig. 13.19

Distribution of female participants rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2011–2017. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

Statistiken zur Bibliographiedatenbank dblp

13.2

Statistics of the dblp computer science bibliography

Dieser Abschnitt enthält statistische Daten zur Bibliographiedatenbank dblp. Fig. 13.20 listet die durchschnittlichen Nutzungszahlen der letzten Jahre. Ein Überblick über die Entwicklung des dblp Datenbestandes kann Fig. 13.21 und Fig. 13.22 entnommen werden. Fig. 13.23–13.25 geben Auskunft über die kontinuierliche Datenkuration und -anreicherung des Bestandes.

This section provides statistical data about the dblp computer science bibliography. Fig. 13.20 shows the average usage statistics of the dblp servers in the past years. An overview of the development of the dblp database can be found in Fig. 13.21 and Fig. 13.22. Information about the continuous data curation and enrichment of existing records can be found in Fig. 13.23–13.25.

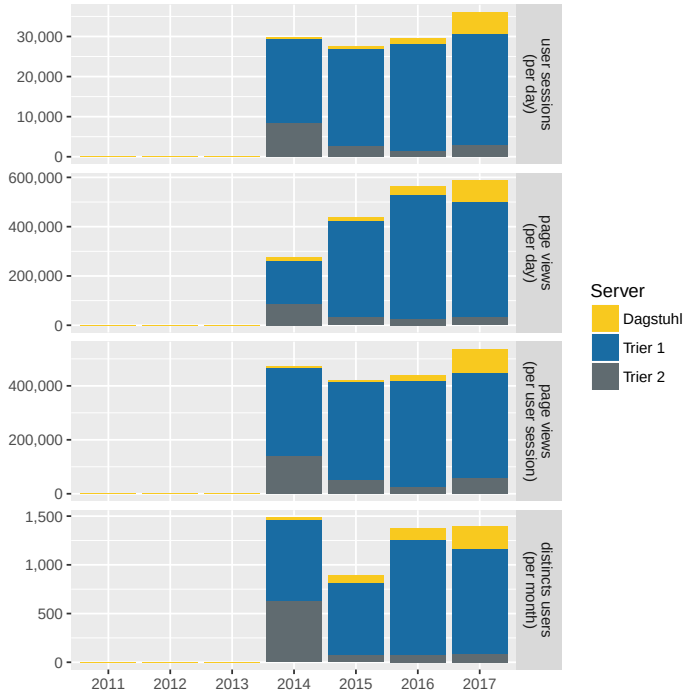

(a) Chart for 2014–2017

	Trier 1		Trier 2		Dagstuhl		Total	
	2016	2017	2016	2017	2016	2017	2016	2017
user sessions (visits) per day	26,911	27,931	1,427	2,836	1,254	5,366	29,593	36,133
page views per day	501,208	466,989	26,354	35,140	35,406	85,537	562,969	587,668
page views per user session	18.6	16.7	18.5	12.4	28.2	15.9	19.0	16.3
distinct users (IPs) per month	393,273	390,886	25,249	58,975	20,416	86,985	438,938	536,847
data served per month	1,187.6 GB	1,082.3 GB	72.7 GB	82.8 GB	120.7 GB	235.0 GB	1,381.0 GB	1,400.1 GB

(b) Detailed numbers for the past two years

Fig. 13.20

Average usage of the three dblp servers. Trier 1 = dblp.uni-trier.de, Trier 2 = dblp2.uni-trier.de, Dagstuhl = dblp.dagstuhl.de. All figures exclude traffic caused by recognized bots and web crawlers. Usage data has not been collected before 2014. In 2015, changes have been made in the server setup in order to shift traffic from development server Trier 2 to the more powerful server Trier 1. Since 2017, server Dagstuhl has been promoted to play a more prominent role under the domain dblp.org

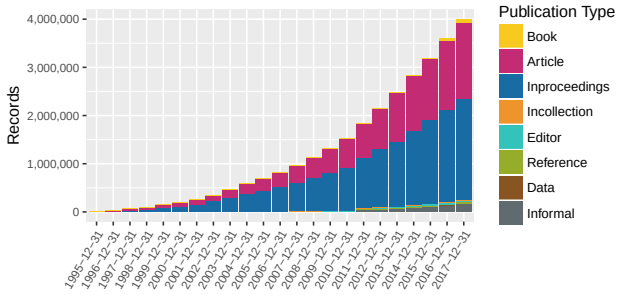

(a) Chart for 1996–2017

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2011	9,356	0.5	705,734	38.3	1,048,662	57.0	7,453	0.4	17,151	0.9	12,207	0.7	0	0.0	39,925	2.2	1,840,488
2012	16,037	0.7	850,603	39.3	1,199,892	55.5	9,631	0.4	20,154	0.9	13,125	0.6	0	0.0	53,810	2.5	2,163,252
2013	16,819	0.7	997,820	40.2	1,350,713	54.4	12,797	0.5	22,771	0.9	13,125	0.5	0	0.0	69,905	2.8	2,483,950
2014	17,533	0.6	1,129,231	39.8	1,545,065	54.5	14,470	0.5	26,137	0.9	14,690	0.5	0	0.0	88,217	3.1	2,835,343
2015	18,318	0.6	1,281,245	40.0	1,724,262	53.9	16,288	0.5	30,044	0.9	19,103	0.6	12	0.0	110,974	3.5	3,200,246
2016	51,070	1.4	1,429,427	39.7	1,912,895	53.1	19,774	0.5	33,782	0.9	20,174	0.6	26	0.0	134,354	3.7	3,601,502
2017	77,408	1.9	1,576,972	39.4	2,091,486	52.2	23,101	0.6	37,049	0.9	23,089	0.6	49	0.0	174,723	4.4	4,003,877

(b) Detailed numbers for 2011–2017

Fig. 13.21
Development of the total size of the dblp database.

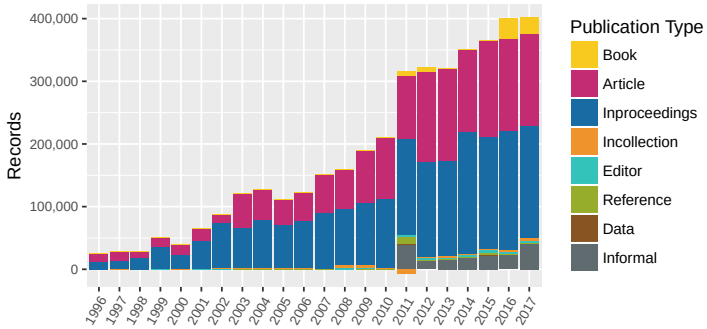

(a) Chart for 1996–2017

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2011	7,920	2.6	98,981	32.1	154,705	50.2	−7,664	-2.5	2,442	0.8	12,207	4.0	0	0.0	39,878	12.9	308,469
2012	6,681	2.1	144,869	44.9	151,230	46.9	2,178	0.7	3,003	0.9	918	0.3	0	0.0	13,885	4.3	322,764
2013	782	0.2	147,217	45.9	150,821	47.0	3,166	1.0	2,617	0.8	0	0.0	0	0.0	16,095	5.0	320,698
2014	714	0.2	131,411	37.4	194,352	55.3	1,673	0.5	3,366	1.0	1,565	0.4	0	0.0	18,312	5.2	351,393
2015	785	0.2	152,014	41.7	179,197	49.1	1,818	0.5	3,907	1.1	4,413	1.2	12	0.0	22,757	6.2	364,903
2016	32,752	8.2	148,182	36.9	188,633	47.0	3,486	0.9	3,738	0.9	1,071	0.3	14	0.0	23,380	5.8	401,256
2017	26,338	6.5	147,545	36.7	178,591	44.4	3,327	0.8	3,267	0.8	2,915	0.7	23	0.0	40,369	10.0	402,375

(b) Detailed numbers for 2011–2017

Fig. 13.22
Development 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. Similarly, in the same year, several thousand *Articles* and *Inproceedings* records have been labeled as *Informal*.

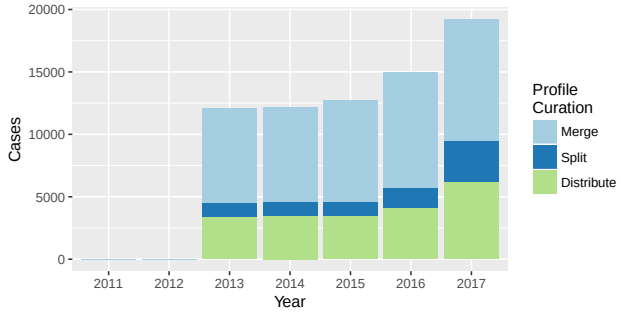

(a) Chart for 2013–2017

Year	Merge		Split		Distribute		Total
	#	%	#	%	#	%	#
2013	7,559	62.7	1,105	9.2	3,399	28.2	12,063
2014	7,576	62.4	1,107	9.1	3,467	28.5	12,150
2015	8,159	63.9	1,130	8.8	3,481	27.3	12,770
2016	9,249	61.7	1,626	10.8	4,120	27.5	14,995
2017	9,731	50.7	3,254	16.9	6,213	32.4	19,198

(b) Detailed numbers for 2013–2017

Fig. 13.23
Data curation of existing dblp person profiles. No curation data has been tracked before 2013. We distinguish between three curation cases: *Merge* = Two or more synonymous profiles have been merged into a single profile. *Split* = A single, homonymous profile has been split into two or more profiles. *Distribute* = A mixed case where records from two or more profiles have been redistributed between two or more profiles.

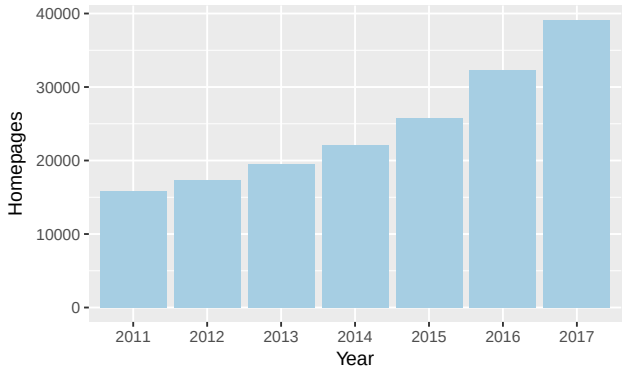

(a) Chart for 2011–2017

Year	Homepages
2011	15,797
2012	17,384
2013	19,511
2014	22,101
2015	25,814
2016	32,249
2017	39,134

(b) Detailed numbers for 2011–2017

Fig. 13.24
Linked academic homepages in dblp person profiles. A single person profile may be linked to multiple academic homepages.

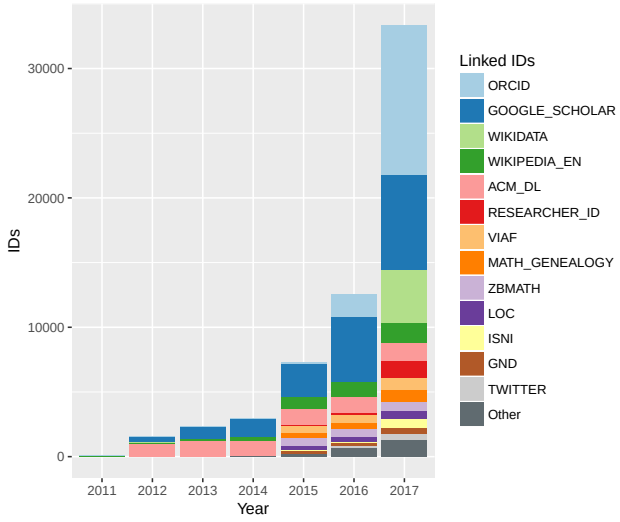

(a) Chart for 2011–2017

Year	ORCID		Google Scholar		Wikidata		Wikipedia (en)		ACM DL		ResearcherID		VIAF	
	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2011	0	0.0	0	0.0	0	0.0	79	100.0	0	0.0	0	0.0	0	0.0
2012	6	0.4	459	29.3	0	0.0	116	7.4	977	62.3	1	0.1	0	0.0
2013	14	0.6	947	39.9	0	0.0	207	8.7	1,169	49.2	2	0.1	2	0.1
2014	24	0.8	1,379	46.9	0	0.0	271	9.2	1,173	39.9	8	0.3	3	0.1
2015	89	1.2	2,510	34.6	0	0.0	1,002	13.8	1,225	16.9	34	0.5	573	7.9
2016	1,717	13.7	4,992	39.9	4	0.0	1,236	9.9	1,236	9.9	154	1.2	585	4.7
2017	11,584	34.7	7,323	22.0	4,050	12.1	1,551	4.7	1,425	4.3	1,296	3.9	950	2.8

Year	Math Genealogy		Zentralblatt MATH		LOC		ISNI		GND		Twitter		Other		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2011	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	79
2012	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	9	0.6	1,568
2013	0	0.0	0	0.0	0	0.0	0	0.0	0	0.0	1	0.0	32	1.3	2,374
2014	0	0.0	4	0.1	0	0.0	1	0.0	0	0.0	13	0.4	64	2.2	2,940
2015	341	4.7	622	8.6	357	4.9	1	0.0	240	3.3	67	0.9	200	2.8	7,261
2016	420	3.4	687	5.5	357	2.9	7	0.1	243	1.9	184	1.5	700	5.6	12,522
2017	879	2.6	701	2.1	677	2.0	671	2.0	481	1.4	454	1.4	1,301	3.9	33,343

(b) Detailed numbers for 2011–2017

Fig. 13.25
Linked external person IDs in dblp person profiles. A single person profile may be linked to multiple external IDs.

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. 13.26 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen, Fig. 13.27 die der Publikationen in der Reihe Dagstuhl Manifestos und schließlich Fig. 13.28 die der veröffentlichten Bände in der Reihe Dagstuhl Follow-Ups.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 13.29 fasst die Daten in der Reihe OASIcs und Fig. 13.30 die der Reihe LIPIcs zusammen.

Fig. 13.31 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-focused series: Fig. 13.26 summarizes the data of the periodical Dagstuhl Reports, Fig. 13.27 the data of the Dagstuhl Manifestos series, and, finally, Fig. 13.28 that of the volumes published in the Dagstuhl Follow-Ups series.

The statistical data to the service-focused series are presented afterwards. Fig. 13.29 presents numbers related to OASIcs and Fig. 13.30 numbers related to LIPIcs.

We summarize the publications of the journal LITES in Fig. 13.31.

Please note that the publication series were established in different years in the period between 2009 and 2015. However, we always consider this complete period.

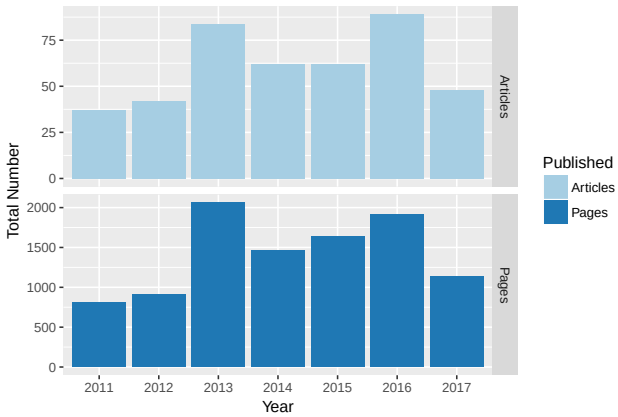

(a) Graphical distribution for 2011–2017

Year	Articles	Pages
2011	37	806
2012	42	913
2013	84	2059
2014	62	1464
2015	62	1636
2016	89	1910
2017	48	1138

(b) Detailed numbers for 2011–2017

Fig. 13.26
Statistics about Dagstuhl Reports published between 2011 to 2017.

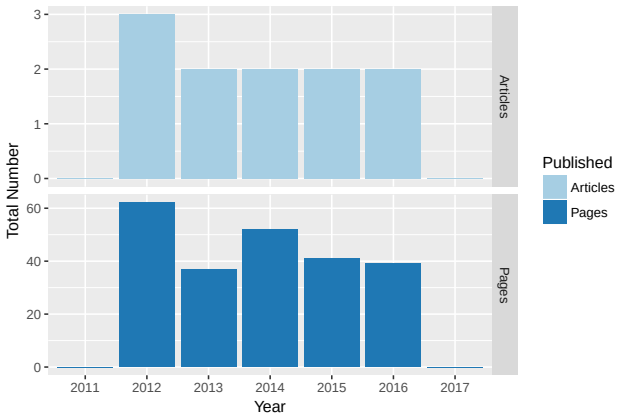

(a) Graphical distribution for 2012–2017

Year	Articles	Pages
2011	0	0
2012	3	62
2013	2	37
2014	2	52
2015	2	41
2016	2	39
2017	0	0

(b) Detailed numbers for 2012–2017

Fig. 13.27
Statistics about Dagstuhl Manifestos published between 2012 to 2017.

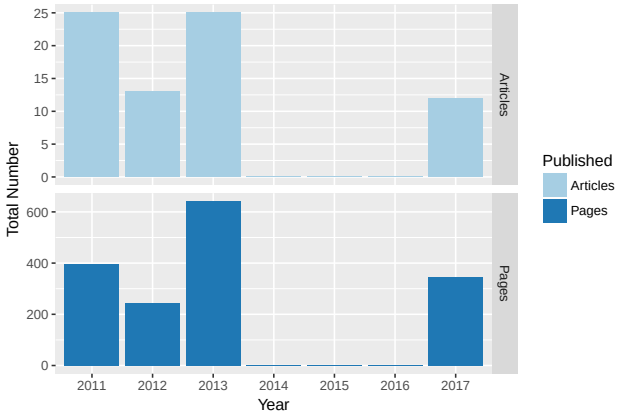

(a) Graphical distribution for 2010–2017

Year	Volumes	Articles	Pages
2011	1	25	395
2012	1	13	246
2013	3	25	641
2014	0	0	0
2015	0	0	0
2016	0	0	0
2017	1	12	346

(b) Detailed numbers for 2010–2017

Fig. 13.28
Statistics about Dagstuhl Follow-Ups volumes published between 2010 to 2017.

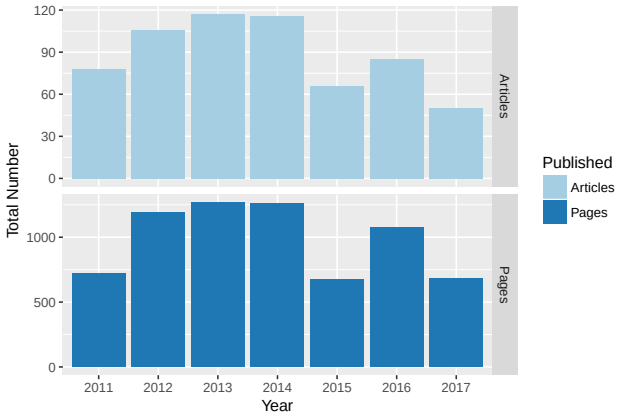

(a) Graphical distribution for 2011–2017

Year	Volumes	Articles	Pages
2011	5	78	717
2012	8	106	1192
2013	7	117	1265
2014	8	116	1264
2015	6	66	674
2016	6	85	1078
2017	3	50	684

(b) Detailed numbers for 2011–2017

Fig. 13.29
Statistics about OASIcs volumes published between 2011 to 2017.

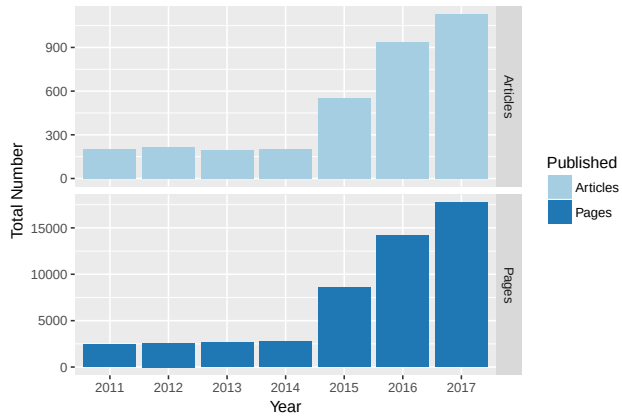

(a) Graphical distribution for 2011–2017

Year	Volumes	Articles	Pages
2011	5	205	2439
2012	5	215	2591
2013	6	195	2607
2014	5	204	2752
2015	16	553	8565
2016	19	939	14222
2017	25	1127	17687

(b) Detailed numbers for 2011–2017

Fig. 13.30
Statistics about LIPIcs volumes published between 2011 to 2017.

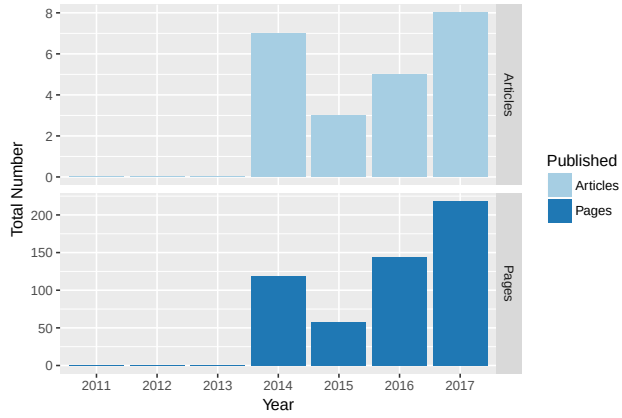

(a) Graphical distribution for 2014–2017

Year	Articles	Pages
2011	0	0
2012	0	0
2013	0	0
2014	7	119
2015	3	58
2016	5	144
2017	8	218

(b) Detailed numbers for 2014–2017

Fig. 13.31
Statistics about LITES articles published between 2014 to 2017.

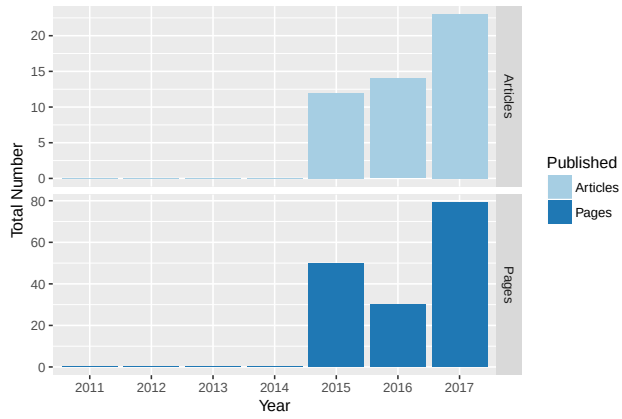

(a) Graphical distribution for 2011–2017

Year	Articles	Pages
2011	0	0
2012	0	0
2013	0	0
2014	0	0
2015	12	50
2016	14	30
2017	23	79

(b) Detailed numbers for 2011–2017

Fig. 13.32
Statistics about DARTS artifacts published between 2011 to 2017.

14 Veranstaltungen 2017
Schedule of Events 2017

Dagstuhl-Seminare 14.1 *Dagstuhl Seminars*

17021 – Functoriality in Geometric Data
Mirela Ben-Chen (Technion – Haifa, IL), Frédéric Chazal (INRIA Saclay – Île-de-France, FR), Leonidas J. Guibas (Stanford University, US), Maks Ovsjanikov (Ecole Polytechnique – Palaiseau, FR)
January 8–13, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17021

17022 – Automated Program Repair
Sunghun Kim (HKUST – Kowloon, HK), Claire Le Goues (Carnegie Mellon University – Pittsburgh, US), Michael Pradel (TU Darmstadt, DE), Abhik Roychoudhury (National University of Singapore, SG)
January 8–13, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17022

17031 – Planning and Robotics
Malik Ghallab (LAAS – Toulouse, FR), Nick Hawes (University of Birmingham, GB), Daniele Magazzeni (King's College London, GB), Andrea Orlandini (CNR – Rome, IT), Brian C. Williams (MIT – Cambridge, US)
January 15–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17031

17032 – Network Function Virtualization in Software Defined Infrastructures
David Hausheer (TU Darmstadt, DE), Oliver Hohlfeld (RWTH Aachen, DE), Diego R. Lopez (Telefonica I+D – Seville, ES), Bruce MacDowell Maggs (Duke University – Durham, US), Costin Raiciu (University Politehnica of Bucharest, RO)
January 15–18, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17032

17041 – Randomization in Parameterized Complexity
Marek Cygan (University of Warsaw, PL), Fedor V. Fomin (University of Bergen, NO), Danny Hermelin (Ben Gurion University of the Negev – Beer Sheva, IL), Magnus Wahlström (Royal Holloway University of London, GB)
January 22–27, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17041

17042 – From Characters to Understanding Natural Language (C2NLU): Robust End-to-End Deep Learning for NLP
Phil Blunsom (University of Oxford, GB), Kyunghyun Cho (New York University, US), Chris Dyer (Carnegie Mellon University – Pittsburgh, US), Hinrich Schütze (LMU München, DE)
January 22–27, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17042

17051 – Theory and Applications of Behavioural Types
Simon Gay (University of Glasgow, GB), Vasco T. Vasconcelos (University of Lisbon, PT), Philip Wadler (University of Edinburgh, GB), Nobuko Yoshida (Imperial College London, GB)
January 29 to February 3, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17051

17061 – Wildly Heterogeneous Post-CMOS Technologies Meet Software
Jerónimo Castrillón-Mazo (TU Dresden, DE), Tei-Wei Kuo (National Taiwan University – Taipei, TW), Heike E. Riel (IBM Research Zurich, CH), Sayeef Salahuddin (University of California – Berkeley, US)
February 5–10, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17061

17062 – Beyond VR and AR: Reimagining Experience Sharing and Skill Transfer Towards an Internet of Abilities
Anind K. Dey (Carnegie Mellon University – Pittsburgh, US), Jonna Häkkilä (University of Lapland – Rovaniemi, FI), Niels Henze (Universität Stuttgart, DE), Kai Kunze (Keio University – Yokohama, JP), Jun Rekimoto (University of Tokyo, JP)
February 5–10, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17062

17071 – Computer-Assisted Engineering for Robotics and Autonomous Systems
Erika Abraham (RWTH Aachen, DE), Hadas Kress-Gazit (Cornell University – Ithaca, US), Lorenzo Natale (Italian Institute of Technology – Genova, IT), Armando Tacchella (University of Genova, IT)
February 12–17, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17071

17072 – Applications of Topology to the Analysis of 1-Dimensional Objects
Benjamin Burton (The University of Queensland, AU), Maarten Löffler (Utrecht University, NL), Carola Wenk (Tulane University, US), Erin Moriarty Wolf Chambers (St. Louis University, US)
February 12–17, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17072

17081 – Computability Theory
Klaus Ambos-Spies (Universität Heidelberg, DE), Vasco Brattka (Universität der Bundeswehr – München, DE), Rodney Downey (Victoria University – Wellington, NZ), Steffen Lempp (University of Wisconsin – Madison, US)
February 19–24, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17081

17082 – Shape-Changing Interfaces
Jason Alexander (Lancaster University, GB), Sean Follmer (Stanford University, US), Kasper Hornbaek (University of Copenhagen, DK), Anne Roudaut (University of Bristol, GB)
February 19–22, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17082

17091 – Computer Science Meets Ecology
Benjamin Adams (University of Auckland, NZ), Gustau Camps-Valls (University of Valencia, ES), Thomas Hickler (Senckenberg Research Centre – Frankfurt am Main, DE), Birgitta König-Ries (Universität Jena, DE)
February 26 to March 3, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17091

17092 – Search as Learning
Kevyn Collins-Thompson (University of Michigan – Ann Arbor, US), Preben Hansen (Stockholm University, SE), Claudia Hauff (TU Delft, NL), Claus-Peter Klas (GESIS – Köln, DE)
February 26 to March 1, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17092

17101 – Databases on Future Hardware
Gustavo Alonso (ETH Zürich, CH), Michaela Blott (Xilinx – Dublin, IE), Jens Teubner (TU Dortmund, DE)
March 5–10, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17101

17102 – Rethinking Productivity in Software Engineering
Thomas Fritz (Universität Zürich, CH), Gloria Mark (University of California – Irvine, US), Gail C. Murphy (University of British Columbia – Vancouver, CA), Thomas Zimmermann (Microsoft Corporation – Redmond, US)
March 5–8, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17102

17111 – Game Theory in AI, Logic, and Algorithms
Swarat Chaudhuri (Rice University – Houston, US), Sampath Kannan (University of Pennsylvania – Philadelphia, US), Rupak Majumdar (MPI-SWS – Kaiserslautern, DE), Michael J. Wooldridge (University of Oxford, GB)
March 12–17, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17111

17112 – Using Networks to Teach About Networks
Timur Friedman (UPMC – Paris, FR), Aiko Pras (University of Twente, NL), Jürgen Schönwälder (Jacobs University Bremen, DE)
March 12–15, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17112

17121 – Computational Complexity of Discrete Problems
Anna Gál (University of Texas – Austin, US), Michal Koucký (Charles University – Prague, CZ), Oded Regev (New York University, US), Till Tantau (Universität zu Lübeck, DE)
March 19–24, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17121

17131 – Mixed Criticality on Multicore / Manycore Platforms
Sanjoy K. Baruah (University of North Carolina at Chapel Hill, US), Liliana Cucu-Grosjean (INRIA – Paris, FR), Robert Davis (University of York, GB), Zoë Stephenson (Rapita Systems Ltd. – York, GB)
March 26–31, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17131

17132 – Opportunities and Risks of Blockchain Technologies
Roman Beck (IT University of Copenhagen, DK), Christian Becker (Universität Mannheim, DE), Juho Lindman (University of Göteborg, SE), Matti Rossi (Aalto University, FI)
March 26–29, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17132

17141 – Probabilistic Methods in the Design and Analysis of Algorithms
Bodo Manthey (University of Twente, NL), Claire Mathieu (ENS – Paris, FR), Heiko Röglin (Universität Bonn, DE), Eli Upfal (Brown University – Providence, US)
April 2–7, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17141

17142 – Formal Methods of Transformations
Emmanuel Filiot (Free University of Brussels, BE), Sebastian Maneth (University of Edinburgh, GB), Helmut Seidl (TU München, DE)
April 2–5, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17142

17161 – Ambient Notification Environments
Lewis Chuang (MPI für biologische Kybernetik – Tübingen, DE), Sven Gehring (DFKI – Saarbrücken, DE), Judy Kay (The University of Sydney, AU), Albrecht Schmidt (Universität Stuttgart, DE)
April 17–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17161

17162 – Online Privacy and Web Transparency
Nataliia Bielova (INRIA Sophia Antipolis, FR), Nikolaos Laoutaris (Telefónica Research – Barcelona, ES), Arvind Narayanan (Princeton University, US), Nick Nikiforakis (Stony Brook University, US)
April 17–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17162

17171 – Computational Geometry
Otfried Cheong (KAIST – Daejeon, KR), Anne Driemel (TU Eindhoven, NL), Jeff Erickson (University of Illinois – Urbana-Champaign, US)
April 23–28, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17171

17181 – Theory and Applications of Hashing
Martin Aumüller (IT University of Copenhagen, DK), Martin Dietzfelbinger (TU Ilmenau, DE), Michael Mitzenmacher (Harvard University – Cambridge, US), Rasmus Pagh (IT University of Copenhagen, DK), David P. Woodruff (IBM Almaden Center – San Jose, US)
May 1–5, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17181

17191 – Theory of Randomized Optimization Heuristics
Carola Doerr (CNRS & UPMC, Paris, FR), Christian Igel (University of Copenhagen, DK), Lothar Thiele (ETH Zürich, CH), Xin Yao (University of Birmingham, GB)
May 7–12, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17191

17192 – Human-Like Neural-Symbolic Computing
Tarek R. Besold (Universität Bremen, DE), Artur d'Avila Garcez (City – University of London, GB), Ramanathan V. Guha (Los Altos Hills, US), Luis C. Lamb (Federal University of Rio Grande do Sul, BR)
May 7–12, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17192

17201 – Formal Synthesis of Cyber-Physical Systems
Calin A. Belta (Boston University – Brookline, US), Rupak Majumdar (MPI-SWS – Kaiserslautern, DE), Matthias Rungger (TU München, DE), Majid Zamani (TU München, DE)
May 14–19, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17201

17202 – Challenges and Opportunities of User-Level File Systems for HPC
André Brinkmann (Universität Mainz, DE), Kathryn Mohror (LLNL – Livermore, US), Weikuan Yu (Florida State University – Tallahassee, US)
May 14–19, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17202

17221 – Geometric Modelling, Interoperability and New Challenges
Falai Chen (Univ. of Science & Technology of China – Anhui, CN), Tor Dokken (SINTEF – Oslo, NO), Thomas A. Grandine (The Boeing Company – Seattle, US), Géraldine Morin (University of Toulouse, FR)
May 28 to June 2, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17221

17222 – Robust Performance in Database Query Processing
Renata Borovica-Gajic (The University of Melbourne, AU), Goetz Graefe (Google – Madison, US), Allison Lee (Snowflake Computing Inc. – San Mateo, US)
May 28 to June 2, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17222

17231 – Epistemic Planning
Chitta Baral (Arizona State University – Tempe, US), Thomas Bolander (Technical University of Denmark – Lyngby, DK), Sheila McIlraith (University of Toronto, CA), Hans Van Ditmarsch (LORIA – Nancy, FR)
June 5–9, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17231

17232 – Computational Interactivity
Xiaojun Bi (Stony Brook University, US), Otmar Hilliges (ETH Zürich, CH), Takeo Igarashi (University of Tokyo, JP), Antti Oulasvirta (Aalto University, FI)
June 5–8, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17232

17251 – Game Theory Meets Computational Learning Theory
Maria-Florina Balcan (Carnegie Mellon University – Pittsburgh, US), Paul Dütting (London School of Economics, GB), Paul W. Goldberg (University of Oxford, GB), Michael J. Kearns (University of Pennsylvania – Philadelphia, US), Yishay Mansour (Tel Aviv University, IL)
June 18–23, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17251

17252 – Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction
Rolf Backofen (Universität Freiburg, DE), Yael Mandel-Gutfreund (Technion – Haifa, IL), Uwe Ohler (Max-Delbrück-Centrum – Berlin, DE), Gabriele Varani (University of Washington – Seattle, US)
June 18–21, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17252

17261 – Voting: Beyond Simple Majorities and Single-Winner Elections
Dorothea Baumeister (Heinrich-Heine-Universität Düsseldorf, DE), Piotr Faliszewski (AGH University of Science & Technology – Krakow, PL), Annick Laruelle (University of the Basque Country – Bilbao, ES), Toby Walsh (TU Berlin, DE)
June 25–30, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17261

17262 – Federated Semantic Data Management
Johann-Christoph Freytag (HU Berlin, DE), Olaf Hartig (Linköping University, SE), Maria-Esther Vidal (Universidad S. Bolivar – Caracas, VE)
June 25–30, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17262

17271 – Foundations of Wireless Networking
Christina Fragouli (University of California at Los Angeles, US), Magnus M. Halldorsson (Reykjavik University, IS), Kyle Jamieson (Princeton University, US & University College London, GB), Bhaskar Krishnamachari (USC – Los Angeles, US)
July 2–7, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17271

17272 – Citizen Science: Design and Engagement
Irene Celino (CEFRIEL – Milan, IT), Oscar Corcho (Polytechnic University of Madrid, ES), Franz Hölker (IGB – Berlin, DE), Elena Simperl (University of Southampton, GB)
July 2–5, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17272

17281 – Malware Analysis: From Large-Scale Data Triage to Targeted Attack Recognition
Saumya K. Debray (University of Arizona – Tucson, US), Thomas Dullien (Google Switzerland – Zürich, CH), Arun Lakhotia (University of Louisiana – Lafayette, US), Sarah Zennou (Airbus Group – Suresnes, FR)
July 9–14, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17281

17282 – From Observations to Prediction of Movement
Mark Birkin (University of Leeds, GB), Somayeh Dodge (University of Minnesota – Minneapolis, US), Brittany Terese Fasy (Montana State University – Bozeman, US), Richard Philip Mann (University of Leeds, GB)
July 9–14, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17282

17291 – Resource Bound Analysis
Marco Gaboardi (University at Buffalo, US), Jan Hoffmann (Carnegie Mellon University – Pittsburgh, US), Reinhard Wilhelm (Universität des Saarlandes, DE), Florian Zuleger (TU Wien, AT)
July 16–21, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17291

17292 – Topology, Computation and Data Analysis
Hamish Carr (University of Leeds, GB), Michael Kerber (TU Graz, AT), Bei Wang (University of Utah – Salt Lake City, US)
July 16–21, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17292

17301 – User-Generated Content in Social Media
Tat-Seng Chua (National University of Singapore, SG), Norbert Fuhr (Universität Duisburg-Essen, DE), Gregory Grefenstette (IHMC – Paris, FR), Kalervo Järvelin (University of Tampere, FI), Jaakko Peltonen (Aalto University, FI)
July 23–28, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17301

17332 – Scalable Set Visualizations
Yifan Hu (Yahoo! Research – New York, US), Luana Micallef (Aalto University, FI), Martin Nöllenburg (TU Wien, AT), Peter Rodgers (University of Kent – Canterbury, GB)
August 13–18, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17332

17341 – Computational Counting
Ivona Bezáková (Rochester Institute of Technology, US), Leslie Ann Goldberg (University of Oxford, GB), Mark R. Jerrum (Queen Mary University of London, GB)
August 20–25, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17341

17342 – SLEBOK: The Software Language Engineering Body of Knowledge
Benoit Combemale (IRISA – Rennes, FR), Ralf Lämmel (Universität Koblenz-Landau, DE), Eric Van Wyk (University of Minnesota – Minneapolis, US)
August 20–25, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17342

17351 – Machine Learning and Formal Methods
Susmit Jha (SRI – Menlo Park, US), Andreas Krause (ETH Zürich, CH), Sanjit A. Seshia (University of California – Berkeley, US), Xiaojin Zhu (University of Wisconsin – Madison, US)
August 27 to September 1, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17351

17352 – Analysis and Synthesis of Floating-point Programs
Eva Darulova (MPI-SWS – Saarbrücken, DE), Alastair F. Donaldson (Imperial College London, GB), Zvonimir Rakamaric (University of Utah – Salt Lake City, US), Cindy Rubio-Gonzalez (University of California – Davis, US)
August 27–30, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17352

17361 – Finite and Algorithmic Model Theory
Anuj Dawar (University of Cambridge, GB), Erich Grädel (RWTH Aachen, DE), Phokion G. Kolaitis (University of California – Santa Cruz, US), Thomas Schwentick (TU Dortmund, DE)
September 3–8, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17361

17371 – Deduction Beyond First-Order Logic
Jasmin Christian Blanchette (VU University of Amsterdam, NL), Carsten Fuhs (Birkbeck, University of London, GB), Viorica Sofronie-Stokkermans (Universität Koblenz-Landau, DE), Cesare Tinelli (University of Iowa – Iowa City, US)
September 10–15, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17371

17372 – Cybersafety in Modern Online Social Networks
Jeremy Blackburn (University of Alabama at Birmingham, US), Emiliano De Cristofaro (University College London, GB), Michael Sirivianos (Cyprus University of Technology – Lemesos, CY), Thorsten Strufe (TU Dresden, DE)
September 10–13, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17372

17381 – Recent Trends in Knowledge Compilation
Adnan Darwiche (UCLA, US), Pierre Marquis (Artois University – Lens, FR), Dan Suciu (University of Washington – Seattle, US), Stefan Szeider (TU Wien, AT)
September 17–22, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17381

17382 – Approaches and Applications of Inductive Programming
Stephen H. Muggleton (Imperial College London, GB), Ute Schmid (Universität Bamberg, DE), Rishabh Singh (Microsoft Research – Redmond, US)
September 17–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17382

17391 – Deep Learning for Computer Vision
Daniel Cremers (TU München, DE), Laura Leal-Taixé (TU München, DE), Ian Reid (University of Adelaide, AU), René Vidal (Johns Hopkins University – Baltimore, US)
September 24–29, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17391

17392 – Body-Centric Computing
Steve Benford (University of Nottingham, GB), Kristina Höök (KTH Royal Institute of Technology – Stockholm, SE), Florian Mueller (RMIT University – Melbourne, AU), Dag Svanes (NTNU – Trondheim, NO)
September 24–29, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17392

17401 – Quantum Cryptanalysis
Michele Mosca (University of Waterloo, CA), Nicolas Sendrier (INRIA – Paris, FR), Rainer Steinwandt (Florida Atlantic University – Boca Raton, US), Krysta Svore (Microsoft Corporation – Redmond, US)
October 1–6, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17401

17411 – Hyperspectral, Multispectral, and Multimodal (HMM) Imaging: Acquisition, Algorithms, and Applications
Gonzalo R. Arce (University of Delaware, US), Richard Bamler (DLR – Oberpfaffenhofen, DE), Jon Yngve Hardeberg (Norwegian Univ. of Science & Technology – Gjøvik, NO), Andreas Kolb (Universität Siegen, DE), Shida Beigpour (MPI für Informatik – Saarbrücken, DE)
October 8–13, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17411

17412 – Internet of People
Elizabeth M. Belding (University of California – Santa Barbara, US), Jörg Ott (TU München, DE), Andrea Passarella (CNR – Pisa, IT), Peter Reichl (Universität Wien, AT)
October 8–11, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17412

17421 – Computational Proteomics
Bernhard Küster (TU München, DE), Kathryn Lilley (University of Cambridge, GB), Lennart Martens (Ghent University, BE)
October 15–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17421

17431 – Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions
Anshu Dubey (Argonne National Laboratory, US), Paul H. J. Kelly (Imperial College London, GB), Bernd Mohr (Jülich Supercomputing Centre, DE), Jeffrey S. Vetter (Oak Ridge National Laboratory, US)
October 22–27, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17431

17441 – Big Stream Processing Systems
Martin Hirzel (IBM TJ Watson Research Center – Yorktown Heights, US), Tilmann Rabl (TU Berlin, DE), Sherif Sakr (KSAU – Riyadh, SA)
October 29 to November 3, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17441

17451 – New Challenges in Parallelism
Annette Bieniusa (TU Kaiserslautern, DE), Hans-J. Boehm (Google – Palo Alto, US), Maurice Herlihy (Brown University – Providence, US), Erez Petrank (Technion – Haifa, IL)
November 5–10, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17451

17452 – Algorithmic Cheminformatics
Jakob Lykke Andersen (Tokyo Institute of Technology, JP), Christoph Flamm (Universität Wien, AT), Daniel Merkle (University of Southern Denmark – Odense, DK), Peter F. Stadler (Universität Leipzig, DE)
November 5–10, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17452

17461 – Connecting Visualization and Data Management Research
Remco Chang (Tufts University – Medford, US), Jean-Daniel Fekete (INRIA Saclay – Orsay, FR), Juliana Freire (New York University, US), Carlos E. Scheidegger (University of Arizona – Tucson, US)
November 12–17, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17461

17462 – A Shared Challenge in Behavioural Specification
Klaus Havelund (NASA – Pasadena, US), Martin Leucker (Universität Lübeck, DE), Giles Reger (University of Manchester, GB), Volker Stolz (West. Norway Univ. of Applied Sciences – Bergen, NO)
November 12–15, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17462

17471 – Artificial and Computational Intelligence in Games: AI-Driven Game Design
Elisabeth André (Universität Augsburg, DE), Michael Cook (University of London, GB), Mike Preuß (Universität Münster, DE), Pieter Spronck (Tilburg University, NL)
November 19–24, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17471

17472 – Addressing the Computational Challenges of Personalized Medicine
Niko Beerenwinkel (ETH Zürich – Basel, CH), Holger Fröhlich (UCB Biosciences GmbH – Monheim, DE), Franziska Michor (Harvard Medical School – Boston, US), Susan A. Murphy (Harvard University – Cambridge, US)
November 19–22, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17472

17481 – Reliable Computation and Complexity on the Reals
Norbert T. Müller (Universität Trier, DE), Siegfried M. Rump (TU Hamburg-Harburg, DE), Klaus Weihrauch (FernUniversität in Hagen, DE), Martin Ziegler (KAIST – Daejeon, KR)
November 26 to December 1, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17481

17491 – Computational Metabolomics: Identification, Interpretation, Imaging
Theodore Alexandrov (EMBL Heidelberg, DE), Sebastian Böcker (Universität Jena, DE), Pieter Dorrestein (UC – San Diego, US), Emma Schymanski (University of Luxembourg, LU)
December 3–8, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17491

17492 – Multi-Level Modelling
Joao Paulo Almeida (Federal University of Espírito Santo – Vitória, BR), Colin Atkinson (Universität Mannheim, DE), Ulrich Frank (Universität Duisburg-Essen, DE), Thomas Kühne (Victoria University – Wellington, NZ)
December 3–8, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17492

17502 – Testing and Verification of Compilers
Junjie Chen (Peking University, CN), Alastair F. Donaldson (Imperial College London, GB), Andreas Zeller (Universität des Saarlandes, DE), Hongyu Zhang (University of Newcastle, AU)
December 10–13, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17502

17511 – The Critical Internet Infrastructure Revisited
Georg Carle (TU München, DE), Thomas C. Schmidt (HAW – Hamburg, DE), Steve Uhlig (Queen Mary University of London, GB), Matthias Wählisch (FU Berlin, DE), Walter Willinger (Niksun – Princeton, US)
December 17–20, 2017 | Dagstuhl Seminar | https://www.dagstuhl.de/17511

Dagstuhl-Perspektiven-Workshops 14.2 *Dagstuhl Perspectives Workshops*

17442 – Towards Cross-Domain Performance Modeling and Prediction: IR/RecSys/NLP
Nicola Ferro (University of Padova, IT), Norbert Fuhr (Universität Duisburg-Essen, DE), Gregory Grefenstette (IHMC – Paris, FR), Joseph Konstan (University of Minnesota – Minneapolis, US)
October 29 to November 3, 2017 | Dagstuhl Perspectives Workshop | https://www.dagstuhl.de/17442

GI-Dagstuhl-Seminare 14.3 *GI-Dagstuhl Seminars*

17303 – Internet of Things Hackathon: From Research to Practice
Alexander Afanasyev (UCLA, US), Oliver Hahm (Zühlke Engineering GmbH, DE), Matthias Wählisch (FU Berlin, DE)
July 23–28, 2017 | GI-Dagstuhl Seminar | https://www.dagstuhl.de/17303

Lehrveranstaltungen 14.4 *Educational Events*

17333 – Sommerschule "Data Management Techniques"
Goetz Graefe (Google – Madison, US)
August 13–18, 2017 | Educational Event | https://www.dagstuhl.de/17333

17403 – Autumn School 2017 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 1–6, 2017 | Educational Event | https://www.dagstuhl.de/17403

17503 – Lehrerfortbildung in Informatik
Manuel Garcia Mateos (LPM Saarbrücken, DE), Michael Gerke (Schloss Dagstuhl – Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 13–15, 2017 | Educational Event | https://www.dagstuhl.de/17503

Forschungsgruppentreffen 14.5 *Research Group Meetings*

17053 – Intelligent Multimodal Interaction
Elisabeth André (Universität Augsburg, DE), Leo Wanner (ICREA and Universitat Pompeu Fabra – Barcelona, ES)
January 29 to February 1, 2017 | Research Group Meeting | https://www.dagstuhl.de/17053

17083 – Klausurtagung Telematik Karlsruhe
Robert Bauer (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)
February 22–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17083

17084 – Digitalisierte Mobilität
Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE)
February 22–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17084

17109 – Forschungsaufenthalt
Kay Hamacher (TU Darmstadt, DE)
March 9–14, 2017 | Research Group Meeting | https://www.dagstuhl.de/17109

17113 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)
Michaela Huhn (Ostfalia Hochschule – Wolfenbüttel, DE), Matthias Riebisch (Universität Hamburg, DE)
March 15–17, 2017 | Research Group Meeting | https://www.dagstuhl.de/17113

17123 – Lehrstuhltreffen "Embedded Intelligence"
Bernhard Sick (Universität Kassel, DE)
March 19–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17123

17124 – Klausurtagung Large Scale Ontology Learning
Klaus-Peter Scherer (KIT – Karlsruher Institut für Technologie, DE)
March 20–22, 2017 | Research Group Meeting | https://www.dagstuhl.de/17124

17125 – Yearly Research Retreat of the Software Technology Group from TU Kaiserslautern
Arnd Poetzsch-Heffter (TU Kaiserslautern, DE)
March 23–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17125

17133 – Lehrstuhltreffen AG Zeller
Andreas Zeller (Universität des Saarlandes, DE)
March 29–31, 2017 | Research Group Meeting | https://www.dagstuhl.de/17133

17143 – PROTASIS Marie Sklodowska-Curie IoT Workshop
Herbert Bos (Free University Amsterdam, NL), Thorsten Holz (Ruhr-Universität Bochum, DE), Evangelos Markatos (FORTH – Heraklion, GR)
April 5–7, 2017 | Research Group Meeting | https://www.dagstuhl.de/17143

17153 – GIBU 2017: GI-Beirat der Universitätsprofessoren
Gregor Snelting (KIT – Karlsruher Institut für Technologie, DE)
April 9–11, 2017 | Research Group Meeting | https://www.dagstuhl.de/17153

17154 – Klausurtagung "LST Halang"
Wolfgang A. Halang (FernUniversität in Hagen, DE)
April 9–12, 2017 | Research Group Meeting | https://www.dagstuhl.de/17154

17163 – DBLP-Treffen
Florian Reitz (Schloss Dagstuhl – Trier, DE)
April 20–21, 2017 | Research Group Meeting | https://www.dagstuhl.de/17163

17174 – Arbeitstreffen Text-Technology Lab
Alexander Mehler (Goethe-Universität – Frankfurt a. M., DE)
April 23–26, 2017 | Research Group Meeting | https://www.dagstuhl.de/17174

17185 – Fakultätsklausur der Fakultät für Informatik – SRH Hochschule Heidelberg
Barbara Sprick (SRH Hochschule Heidelberg, DE)
May 2–3, 2017 | Research Group Meeting | https://www.dagstuhl.de/17185

17189 – Forschungsaufenthalt
Carla Manni (University of Rome "Tor Vergata", IT)
May 1–6, 2017 | Research Group Meeting | https://www.dagstuhl.de/17189

17213 – Kolloquium zum GI Dissertationspreis 2016
Steffen Hölldobler (TU Dresden, DE)
May 21–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17213

17215 – Klausurtagung
Christian Rossow (Universität des Saarlandes, DE)
May 22–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17215

17229 – Forschungsaufenthalt
Paul Buitelaar (National University of Ireland – Galway, IE)
May 28 to June 2, 2017 | Research Group Meeting | https://www.dagstuhl.de/17229

17233 – GI-Tagung der Arbeitsgruppe Robotersysteme
Karsten Berns (TU Kaiserslautern, DE)
June 8–9, 2017 | Research Group Meeting | https://www.dagstuhl.de/17233

17243 – Gemeinsamer Workshop der Graduiertenkollegs – 11th Joint Workshop of the German Research Training in Computer Science: GRKs 1994 AIPHES and 2050 PAT
Thomas Arnold (TU Darmstadt, DE), Tim Grube (TU Darmstadt, DE), Iryna Gurevych (TU Darmstadt, DE), Benjamin Heinzerling (TU Darmstadt, DE), Max Jakob Maaß (TU Darmstadt, DE), Max Mühlhäuser (TU Darmstadt, DE), Christopher Tauchmann (TU Darmstadt, DE), Nora Wessels (TU Darmstadt, DE)
June 11–14, 2017 | Research Group Meeting | https://www.dagstuhl.de/17243

17273 – Workshop Buchprojekt "Corporate Semantic Web"
Thomas Hoppe (Datenlabor Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE), Anatol Reibold (OntoPort UG – Darmstadt, DE)
July 5–7, 2017 | Research Group Meeting | https://www.dagstuhl.de/17273

17279 – Forschungsaufenthalt
Carola Doerr (CNRS & UPMC, Paris, FR)
July 2–7, 2017 | Research Group Meeting | https://www.dagstuhl.de/17279

17353 – Klausurtagung "AG Goesele"
Michael Goesele (TU Darmstadt, DE)
August 30 to September 1, 2017 | Research Group Meeting | https://www.dagstuhl.de/17353

17363 – Klausurtagung
Felix Freiling (Universität Erlangen-Nürnberg, DE)
September 3–6, 2017 | Research Group Meeting | https://www.dagstuhl.de/17363

17364 – Propelling Business Process Management
Martin Matzner (Universität Erlangen-Nürnberg, DE), Nadine Ogonek (Universität Münster, DE), Armin Stein (Universität Münster, DE)
September 6–8, 2017 | Research Group Meeting | https://www.dagstuhl.de/17364

17373 – Digitalisierte Mobilität
Karl-Heinz Krempels (Fraunhofer FIT – Sankt Augustin, DE)
September 13–15, 2017 | Research Group Meeting | https://www.dagstuhl.de/17373

17374 – Lehrstuhltreffen Rechtsinformatik
Christoph Sorge (Universität des Saarlandes, DE)
September 13–14, 2017 | Research Group Meeting | https://www.dagstuhl.de/17374

17375 – Arbeitstreffen Prof. Slusallek
Philipp Slusallek (DFKI – Saarbrücken, DE)
September 14–15, 2017 | Research Group Meeting | https://www.dagstuhl.de/17375

17383 – Klausurtagung Large Scale Ontology Learning (II)
Klaus-Peter Scherer (KIT – Karlsruher Institut für Technologie, DE)
September 20–22, 2017 | Research Group Meeting | https://www.dagstuhl.de/17383

17413 – Klausurtagung AG Schneider
Klaus Schneider (TU Kaiserslautern, DE)
October 11–13, 2017 | Research Group Meeting | https://www.dagstuhl.de/17413

17423 – DDI Moving Forward: Production Framework and Bindings
Kelly Chatain (University of Michigan – Ann Arbor, US), Larry Hoyle (The University of Kansas – Lawrence, US), Jon Johnson (UK Data Service – Colchester, GB), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 15–20, 2017 | Research Group Meeting | https://www.dagstuhl.de/17423

17433 – DDI Moving Forward: Integration of Core Components / DDI-based Infrastructure Vision
Kelly Chatain (University of Michigan – Ann Arbor, US), Larry Hoyle (The University of Kansas – Lawrence, US), Jon Johnson (UK Data Service – Colchester, GB), Steven McEachern (Australian National University – Canberra, AU), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 22–27, 2017 | Research Group Meeting | https://www.dagstuhl.de/17433

17463 – Retreat for the Research Training Group "Adaptive Information Preparation from Heterogeneous Sources (AIPHES)"
Johannes Fürnkranz (TU Darmstadt, DE)
November 15–17, 2017 | Research Group Meeting | https://www.dagstuhl.de/17463

17473 – Klausurtagung "LST Schmeck"
Fabian Rigoll (FZI – Karlsruhe, DE), Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)
November 22–24, 2017 | Research Group Meeting | https://www.dagstuhl.de/17473

17483 – LDBC graphQL Writing Workshop
Peter A. Boncz (CWI – Amsterdam, NL)
November 26 to December 1, 2017 | Research Group Meeting | https://www.dagstuhl.de/17483

17489 – Forschungsaufenthalt
Carla Manni (University of Rome "Tor Vergata", IT)
November 26 to December 2, 2017 | Research Group Meeting | https://www.dagstuhl.de/17489

17505 – Redaktions-Sitzung "Informatik in der Schule"
Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 12, 2017 | Research Group Meeting | https://www.dagstuhl.de/17505

17509 – Forschungsaufenthalt
Gunnar Karlsson (KTH Royal Institute of Technology – Stockholm, SE)
December 10–20, 2017 | Research Group Meeting | https://www.dagstuhl.de/17509

17513 – Secan Lab Meeting
Thomas Engel (University of Luxembourg, LU)
December 18–19, 2017 | Research Group Meeting | https://www.dagstuhl.de/17513

Leibniz
Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2017 | ISSN 2199-1995
http://www.dagstuhl.de